

WHY (SOME) KNOWLEDGE IS THE PROPERTY OF A COMMUNITY AND POSSIBLY NONE OF ITS MEMBERS

BY BOAZ MILLER¹

(FORTHCOMING IN *THE PHILOSOPHICAL QUARTERLY*)

Mainstream analytic epistemology regards knowledge as the property of individuals, rather than groups. Drawing on insights from the reality of knowledge production and dissemination in the sciences, I argue, from within the analytic framework, that this view is wrong. I defend the thesis of ‘knowledge-level justification communalism’, which states that at least some knowledge, typically knowledge obtained from expert testimony, is the property of a community and possibly none of its individual members, in that only the community or some members of it collectively possesses knowledge-level justification for its individual members’ beliefs. I address several objections that individuals, *qua* individuals, have or are able to acquire knowledge-level justification for all the beliefs they obtain from expert testimony. I argue that the problem I identify with individualism is invariant under any specific account of justification, internalist or externalist.

Keywords: testimony; expertise; extendedness hypothesis; epistemic communities; distributed knowledge

I. INTRODUCTION

Orthodox analytic epistemology views knowledge as individualistic. It holds that though an individual may depend on others for acquiring beliefs through testimony, if she ends up acquiring knowledge from testimony, she will have satisfied by herself the substantive conditions for knowledge. This *inter alia* means that she personally possesses the evidence justifying her belief, or that the processes that confer justification on her belief are confined to her cognitive system.

I argue that individualistic conceptions of knowledge cannot characterize as knowledge many of a normal adult’s true beliefs, acquired or inferred from experts’ testimonies, although we normally and legitimately regard them as knowledge. Individualism fails because the justifying elements of many of an individual’s expert-obtained beliefs are distributed among members of her epistemic community. For a wide class of beliefs, only the justification collectively possessed by different members of a subject’s epistemic community amounts to knowledge-level justification. Only a communalist conception of knowledge adequately characterizes such true beliefs as knowledge.

Section II defines and characterizes the strand of epistemic communalism I defend. Section III situates it within the epistemology-of-testimony debate. Section IV outlines the argument. Section V and Section VI deal with objections. Until Section VI, an evidentialist account of justification is assumed. Section VI extends the argument to reliabilism.

II. KNOWLEDGE-LEVEL JUSTIFICATION COMMUNALISM

Individualism is usually not shared by epistemologists working within traditions other than the analytic tradition, but it has also been under attack by philosophers *from within the orthodox analytic camp*. They argue that socially-extended conceptions of knowledge better account for the generation and dissemination of knowledge (Giere 2006; Goldberg 2007; 2010; Bird 2010; Shieber 2011b; Palermos and Pritchard 2013; Adam 2013; 2014; Kerr and Gelfert 2014). They mostly give positive arguments that stress the merits of socially-extended accounts of

knowledge. My argument is negative. I argue that individualism cannot correctly characterize as knowledge a wide class of beliefs. I do not offer a full-fledged communalist account of knowledge, but defend a minimal thesis any analysis of knowledge in the orthodox tradition should accept.

A well-discussed challenge to individualism is Hardwig (1985). The replies to Hardwig misconstrue his challenge. I start by reviewing Hardwig's argument and clarifying the problem it identifies. Hardwig observes that he acquired many of his true beliefs that are normally regarded knowledge from experts' testimonies. However, he does not possess evidence for many of them, or only weak evidence, because he, *qua* individual, lacks the expertise to evaluate the evidence for them. Even when he has or is able to acquire expertise, reviewing the available evidence to justify *all* such beliefs takes more than a lifetime. Therefore, he cannot possess the evidence required for satisfying the justification condition for knowledge for many of his true beliefs.

Call the true beliefs for which a subject lacks, *qua* individual, sufficient evidence for satisfying the justification condition for knowledge '*suspect beliefs*'. Hardwig argues that the evidence required for satisfying the justification condition for knowledge for his suspect beliefs is distributed among members of his epistemic community. Consequently, he poses a dilemma: If the propositions of suspect beliefs are known, either only an epistemic community collectively knows them, in that only the community collectively possesses the evidence that justifies them, or individuals know them vicariously by trusting others.

Hardwig is interpreted as stating that individuals have weak or no evidence for their suspect beliefs, and hold them on blind or partly blind trust; hence, they are not knowledge according to standard analyses of knowledge. Individualists reply that a normal subject personally possesses or is able to acquire *indirect evidence* for her suspect beliefs. While indirect evidence is not the same evidence experts possess, it is still evidence, hence the subject's suspect beliefs are justified rather than held on blind trust, avoiding Hardwig's dilemma (Schmitt 1988; Adler 1994; Goldman 2001, Almassi 2007).

These replies, however, misconstrue Hardwig's challenge. The distinction between believing on blind trust and believing on evidence is a red herring. The relevant distinction is between having mere doxastic justification (or none), and having *knowledge-level justification*, i.e. justification sufficient to grant its possessor knowledge. To avoid Hardwig's dilemma, it is insufficient to show that individuals have some evidence for their suspect beliefs. Such evidence should reach knowledge-level justification. If it does not, Hardwig's dilemma remains, because these beliefs fall short of knowledge.

Let me clarify what 'knowledge-level justification' means. A standard assumption in epistemology is that knowledge is binary, i.e. either *S* knows that *p* or *S* does not know that *p*. By contrast, justification is a matter of degree, i.e. *S*'s belief that *p* can be more or less justified. *S*'s belief may be strongly, moderately, or weakly supported by her evidence. Similarly, *S*'s vision may generate strongly-justified beliefs in good visibility, or weakly-justified beliefs in poor visibility. Thus, to constitute knowledge, a belief must be justified to a sufficient degree: 'there is a certain threshold of justification that must be equalled or exceeded if knowledge is to be obtained' (Dretske 1981: p. 363). Define this threshold as follows:

KL-justification

If p is true and S is not gettiered, then S has KL-justification for S 's belief that p if and only if S knows that p .¹

I argue that subjects typically lack, *qua* individuals, KL-justification for their suspect beliefs (I characterize this class of beliefs in Section IV). Rather, the KL-justification of a normal subject's suspect beliefs is typically possessed by other members of her epistemic community:

KLJ-communalism

For a normal adult subject S , whether some of the true beliefs that S holds are sufficiently justified to amount to knowledge; i.e. have KL-justification, depends on evidence (or other building blocks of which epistemic justification consists) which S does not possess or are not situated within S 's own cognitive system, but are possessed by, or situated within the cognitive systems of other relevant members of S 's epistemic community.

According to KLJ-communalism, a believed proposition is known only if it is true and the overall available evidence within the believing subject's epistemic community sufficiently supports it (I stress that this is a necessary but insufficient condition). It is *not* necessary for knowledge that the evidence the believing subject personally possesses sufficiently supports the believed proposition. If the overall available evidence in the community it is too weak or flawed, the believed proposition is unknown. The communal evidence may be dispersed among many members, e.g. members of a collaborative research project, or concentrated in the hands of few. When the subject possesses this evidence alone, this is a special case where the class of relevant members consists of only one person.

Following Green (1991), we may distinguish two forms of epistemic communalism: *collective-agent* communalism and *collective-good* communalism. According to *collective-agent* communalism, social collectives may be proper subjects of beliefs and knowledge just like individuals.² *Collective-agent* communalism does *not* challenge the traditional individualistic analysis of knowledge, but supplements it. It does not deny that individuals, *qua* individuals, have knowledge. Rather, it adds another ontological level of analysis – the group level. My KLJ-communalism is *not* a form of *collective-agent* communalism (though it is compatible with it).

By contrast, *collective-good* communalism states that some or all of a subject's *ordinary* knowledge or some of its constituents are not hers, but a joint property of many. Knowledge, like language, constitutes a *participatory good*: a good that can only be produced and held by many (Réaume 1988). *Collective-good* communalism need *not* attribute agency to groups or deny that individuals know, only insist that 'the knowing we do as individuals is derivative, that your knowing or mine depends on *our* knowing, for some 'we'' (Nelson 1993: p. 123; emphasis in the origin). *KLJ-communalism is a form of collective-good communalism.*

Many existing *collective-good* accounts of knowledge (e.g. Kusch 2002) subscribe to schools remote from orthodox analytic epistemology. They typically challenge standard orthodox assumptions, e.g. that knowledge is true, factive, rational, propositional, or doxastic.³ I do *not* challenge any orthodox assumption besides individualism. I argue that the orthodox

¹ To remove ambiguity, the logical structure of *KL-justification* is this: $(\phi \wedge \psi) \rightarrow (\phi \leftrightarrow \chi)$.

² E.g. Gilbert (1987), Corlett (1996), Tuomela (2011), List and Pettit (2011).

³ For the assumptions underpinning the orthodox and non-orthodox conceptions of knowledge, see Kitcher (1994).

analysis of knowledge itself entails that sometimes, an individual's true beliefs rise to knowledge only because other members of her epistemic community possess the justificatory building blocks due to which these beliefs reach KL- justification.

III. KLJ-COMMUNALISM AND THE EPISTEMOLOGY OF TESTIMONY

I situated KLJ-communalism within possible communalist accounts of knowledge as a form of collective-good communalism consistent with most orthodox assumptions in analytic epistemology. This section clarifies the relations between KLJ-communalism and the epistemology of testimony. KLJ-communalism is distinct from testimonial reductionism and non-reductionism.

The debate between testimonial reductionism and non-reductionism concerns testimony as a putative source of knowledge. Non-reductionists regard testimony a knowledge source on par with perception and memory; reductionists regard it a secondary source derivative of perception and memory. Both agree that when a hearer possesses defeaters, she does *not* have a presumptive epistemic right to trust the speaker. They disagree whether a hearer has such a right, derivable perhaps from higher epistemic principles, to believe a speaker when the hearer has no positive or negative reasons to trust him. Non-reductionists say she does, like the presumptive right she allegedly has to trust her senses; reductionists say she does not.⁴

KLJ-communalism is either an alternative to reductionism and non-reductionism or consistent with both. Fricker (2002) claims, by way of an argument against non-reductionism, that communalism is an alternative to reductionism and non-reductionism. Fricker grants for the sake of the argument that an *a priori* presumptive right to trust exists, as most non-reductionists claim, but argues that it is too weak to explain the trust we extend to testimony. If our right to trust is like a heavy slab supported by many veridical columns, it is like one column, which cannot carry the weight on its own. Thus, Fricker poses a dilemma to non-reductionists: Either accept that testimonial knowledge 'is the property not of individuals but of communities' (Fricker 2002: p. 372), because only the collectively-held evidence can carry this weight, or opt for reductionism. Fricker argues that communalism is resistible because reductionists can explain testimonial trust by appealing to an individual's rich set of background beliefs by which he can evaluate, on a reductionist model, another person's testimony, trust it, and acquire knowledge from it when appropriate. I refute Fricker's argument in Section VI. Here I stress that for Fricker, epistemic communalism is a *third* position in the testimony debate.

As an internalist, Fricker construes the right to trust in terms of the subject's evidence for believing a testimony. Externalists about justification construe this right in terms of the *de facto* reliability of testimony; namely, its truth conduciveness. For instance, someone an externalist *and* non-reductionist may argue that testimony is *de facto* sufficiently reliable to establish a presumptive right to trust. In Section VII, I argue against this position, yet, unlike Fricker, I do not think that this necessarily has a bearing on the reductionism/non-reductionism debate. Whatever grants a hearer epistemic right to trust a speaker may not give him, *qua* individual, KL-justification. For example, if a presumptive right to trust exists, then everyone arguably has it, including Truman in the movie *The Truman Show*, although most of the

⁴ Fricker (e.g. 2002) is the major contemporary defender of reductionism. Contemporary defenders of non-reductionism include Coady (1992), Burge (1993), and Goldberg (2007; 2010). The characterization of the debate in terms of a hearer's epistemic right follows Fricker. There are similar, yet not equivalent characterizations, which stress other aspects of it (Goldberg 2007: pp. 147-153). My point about the lack of a necessary connection between KLJ-communalism and the reductionism/non-reductionism debate stands under these similar formulations.

testimonies Truman encounters are false.⁵ Put differently, I think the minimal threshold level of justification required for an individual's *justified acceptance* of testimony may be lower than KL-justification. When the two thresholds are debated separately, KLJ-communalism is consistent with reductionism and non-reductionism.

This analysis is disputable. Goldberg (2007: pp. 146-148) argues that non-reductionism, which he endorses, entails that one who justifiably accepts a testimony that p typically also acquires KL-justification for believing that p . Goldberg's argument explicitly rests, however, on a false empirical assumption. I discuss Goldberg's view in Section 6, but as a preview, my argument lends support to Goldberg's account, and the fact that it rests on a false empirical assumption is not detrimental to it.

KLJ-communalism collapses to neither reductionism nor non-reductionism. I leave open whether KLJ-communalism is an alternative to them or consistent with both. Having clarified this, I present my argument in the next section.

IV. AN OUTLINE OF THE ARGUMENT

I argue that a normal subject lacks KL-justification for her suspect beliefs, and this KL-justification is dispersed among members of her epistemic community. Why is that so? Consider the following example. Einstein's General Relativity Theory (GR) entails that massive bodies, such as the sun, deflect light. During a total solar eclipse, this deflection can be measured by a change in the observed position of stars around the sun. In 1918, Eddington designed an experiment to test Einstein's theory against Newton's. Two expeditions went out to record the sun during a total solar eclipse, one to West Africa, another to Brazil. On November 9, 1919, in a meeting of the Royal Society, Eddington announced that the experiment confirmed Einstein's theory and refuted Newton's. The next day, a *New York Times* headline announced 'Einstein's theory triumphs'.

Earman and Glymour (1980) argue that Eddington discarded without scientific rationale eighteen photographic plates from the Brazil expedition. With the Brazil data, the experimental results *disconfirm* Einstein's theory, and support Newton's. Eddington had supported Einstein's theory. His motivation to confirm it arguably influenced his decision to discard the plates. Eddington's claim that Einstein's theory had been confirmed was not justified, because it was unsupported by the evidence (cf. Collins and Pinch 1993: pp. 43-54; Waller 2002: Ch. 3).⁶

Suppose that on November 10, 1919, Hyde, a normal scientifically-lay subject, reads the *New York Times*, and forms the (true) belief that GR is correct. In a nearby possible world, Jekyll, Hyde's counterpart, reads the same headline and forms the same belief. The difference between the worlds is that in Jekyll's world, unlike Hyde's (our) world, Eddington's crew performed the experiment competently and correctly, Eddington did not discard data, and arrived at the right conclusion, confirming Einstein's theory.

My argument is this:

⁵ The movie is about Truman, who, unbeknownst to him, is the star of a reality show about his life, where everyone he encounters is an actor playing a part.

⁶ Kennefick (2009) infers from Eddington's correspondence that Frank Dyson, who was not biased toward Einstein's theory, calculated the Brazil results. Kennefick argues that there is a scientific reason to exclude the eighteen plates. Kennefick's argument, however, is speculative. His evidence that Dyson calculated the Brazil results is suggestive. Kennefick gives *no evidence* that the scientific rationale he suggests for discarding the plates was either Dyson's or Eddington's. I believe Earman and Glymour's historical evidence is compelling, but in the next section I argue that similar cases are common, thus my argument does not depend on the historical accuracy of the example. This is not an argument against Eddington. I use this example because it illustrates my point, and has already been discussed in the context of the epistemology of expertise.

- (1) *Qua* individuals, upon reading Eddington's testimony in the paper, Hyde and Jekyll possess the same evidence for *q*.
- (2) Hyde and Jekyll both believe that *q* based on the evidence they individually possess.
- (3) Therefore (from 1,2), the evidence Jekyll/Hyde individually possesses for *q* is insensitive to the truth of *q*: On this evidence, Jekyll/Hyde believes that *q* whether *q* is true or false.
- (4) Jekyll knows that *q*,
- (5) Therefore (from 4 and *KL-justification*), Jekyll's belief that *q* has KL-justification.
- (6) Hyde does not know that *q*.
- (7) Hyde is not gettiered with respect to his belief that *q*.
- (8) Therefore (from 6,7 and *KL-justification*), Hyde's belief that *q* lacks KL-justification.
- (9) In Hyde's (our) failed-experiment world, the evidence collectively possessed by the relevant members of Hyde's epistemic community on whom Hyde epistemically relies (Eddington, his team, the *New York Times* reporter, etc.) does not amount to KL-justification with respect to Hyde's belief that *q*.
- (10) In Jekyll's successful-experiment world, the evidence collectively possessed by the same relevant members amounts to KL-justification.
- (11) By supposition, the only differences between Jekyll's and Hyde's worlds are the ones described in (9) and (10), their causes, and their outcomes, inasmuch as there are any.
- (12) Therefore (from 1,3,5,8,9,10,11), Jekyll's belief that *q* has KL-justification due to evidence possessed by other members of his epistemic community.

Individualism cannot explain the difference between Jekyll and Hyde. *Qua* individuals, Jekyll and Hyde arguably possess the same evidence for *q*, and form the belief that *q* in the same way. How come Jekyll knows and has KL-justification, while Hyde neither knows nor has KL-Justification? According to KLJ-Communalism, the difference between the cases is the overall available evidence (or other building blocks of epistemic justification) distributed among different relevant members of Jekyll/Hyde's epistemic communities. In Jekyll's world, the overall distributed evidence supports *q*, and amounts to KL-justification. In Hyde's world, it does not. One may immediately object, however, to premises (4) or (7). The next section addresses these objections.

V. THE IMPLICATIONS OF THE PERVASIVENESS OF SCIENTIFIC CORNER-CUTTING TO KL-JUSTIFICATION

Jekyll and Hyde form a true belief that *q* reading a newspaper report containing Eddington's testimony. While in Jekyll's world, the experiment was successful and supports *q*, in Hyde's (our) world, the experiment failed, and Eddington fudged his data to reach his desired conclusion. I claimed that Jekyll knows that *q*, while Hyde does not. One may deny that Jekyll knows that *q*, or alternatively argue that Hyde is gettiered, rather than lacks KL-justification. If either objection is true, my argument is blocked. I address them here.

I first argue that Jekyll knows that *q*. To support this claim, I will show, drawing on historical and empirical evidence about science, that Jekyll-like cases are common. This means that for *many* true beliefs we form or derive from expert testimonies and *normally regard as knowledge*, the experts possess supporting evidence, but it could have relatively easily been the case that the experts would not have had this evidence, but would still have made the same testimonies, and we would still have believed them.

I will give historical and empirical evidence that shows that *Hyde*-like cases are relatively common in science: Due to institutional and psychological factors, scientists sometimes make a testimony that *p*, when their evidence for *p* falls short of KL-justification in that it does not rule out an alternative *plausible* hypothesis that successfully accounts for the evidence too. From the

prevalence of *Hyde*-like case, we may infer the prevalence of *Jekyll*-like cases; namely, that there is a non-negligible number of cases where scientists possess good evidence for their testimony, but would still make the same testimony even without it. The reasoning is simple: The same institutional and psychological factors that operate in *Hyde*-like cases operate in *Jekyll*-like cases, and would push scientists to give unsupported testimonies if they did not obtain the evidence they have. This is similar to inferring that there are many near car-accidents from that there are, in fact, many car-accidents.

The prevalence of *Jekyll*-like cases helps counter the objection that denies that *Jekyll* knows by highlighting the high unwarranted sceptical price of the objection. If *Jekyll* does not know and *Jekyll*-like cases are common, then we all do not know all those things we think we know, which we learned in *Jekyll*-like cases. Even worse, since we do not know which of our expert-obtained beliefs are tainted beliefs acquired in *Jekyll*-like cases, we have *prima facie* good reasons to doubt *all* our beliefs from experts for which we are unable to evaluate the evidence ourselves, i.e. a large portion of our ordinary knowledge. Granting this objection would mean that our normal practice of regarding all these beliefs as knowledge is mistaken. Denying that *Jekyll* knows is not merely denying a philosophical intuition, but amounts to vast scepticism about much of our ordinary knowledge. It means rejecting the normal epistemic practices of classifying it as knowledge, the very same practices our theory of knowledge purports to explain. A defensible argument for paying this price is hard to imagine. Even if one is willing to pay this price, the resulting theory of knowledge will be empirically inadequate, and we will still need a theory that accounts for our actual ordinary knowledge, to which my argument will still apply.

What is the evidence for the prevalence of *Hyde*-like cases from which we infer the prevalence of *Jekyll*-like cases? Historical studies show that Eddington is not an isolated event. Waller reviews six eminent figures from the history of science, including Eddington, Millikan, and Pasteur. Waller (2002: p. 12) writes:

each of the six major scientists examined manipulated their experimental data to fit their preconceived notions of how things really are. Then, to win the scientific battles in which they were engaged, they exploited their powers of obfuscation and deception, their friends in high places, and their reputations as reliable witnesses. All six have been fortunate in the fact that because they were advancing major ideas that now enjoy, at the very least, widespread support, posterity has been largely blind to the equivocal nature of the evidence they presented.

Waller (2002: p. 99) adds: '[i]n all probability, manipulation of experimental data is not just the sin of a few great men who have somehow managed to slip through the net'.

Systematic empirical evidence suggests that such phenomena are indeed prevalent. An anonymous survey of about 3,200 scientists published in *Nature* (Martinson et al. 2005) yields stunning findings. While only 0.3% of surveyed scientists admit to 'falsifying or 'cooking' research data', 6% report 'failing to present data that contradict one's own previous research', 12.5% report 'overlooking others' use of flawed data or questionable interpretation of data', 13.5% report 'using inadequate or inappropriate research designs', and 15.5%(!) report 'dropping observations or data points from analyses based on a gut feeling that they were inaccurate'.

Several institutional and psychological factors are responsible for the prevalence of data manipulation and unreliable testimonies in science. As physicist David Goodstein notes, scientists are rewarded for their results, not effort. Frequently, scientists who commit misconduct are under career pressure, and think they know what their experiment outcomes would be if done properly. They self-justify their data manipulation as a mere shortcut. Goodstein adds that in practice, but as opposed to common perception, experiments are rarely replicated. Additionally, peer-review is not intended to detect fraud, manipulation, or

fabrication, but only to evaluate whether the conclusions stem from the data (Goodstein 2010: pp. 3-17).

The growing role of private industry in research intensifies sceptical worries about scientific knowledge. Pharmaceutical companies are accused of giving clinicians economic incentives to skew results, and of repressing the publication of negative trials. In the survey for *Nature*, 15.5% of scientists admit to ‘changing the design, methodology or results of a study in response to pressure from a funding source’. Clinical trials are becoming increasingly large in the number of participants, decision-makers, and researchers, whose interests are not always clear. The questionable phenomenon of ‘ghost-writing’ is increasing (Sismondo 2009). Marcia Angell, a Harvard researcher and former editor-in-chief of the *New England Journal of Medicine* (2009), argues – among others – that due to commercialization and corruption, much and even most clinical medical research is untrustworthy. Such sceptical worries arise exactly because there is usually no one individual involved in a trial who knows, in the sense of possessing KL-justification, that the drug is efficacious. I do not advocate sweeping scepticism about science, but offer an analysis of knowledge that excludes unsupported research results, while retaining good ones, even if identifying which is which is not always trivial.

This situation is not unique to commercialized biomedicine. Consider mathematics – a field hardly affected by commercial interests. Contra common perception, papers in mathematical journals often contain errors, which are not caught in the peer-review process. Some mathematicians complain that the rate of errors is too high, which endangers the epistemic integrity of the field (Grcar 2010; Nathanson 2008). Social epistemologists have consequently challenged the widely held view that testimony in mathematical journals is a reliable source of knowledge and called for tightening the conditions under which mathematical knowledge is assumed to be acquired from testimony (Geist et al. 2010; Frans and Kosolovsky 2014). Most published theorems, however, are still true although their proofs contain errors. *This means that Hyde-like cases are de facto common in mathematics, which gives us good reasons to think that Jekyll-like cases are also common.* KLJ-communalism correctly analyzes this situation without resorting to scepticism: It grants knowledge status only to true theorems whose proofs are not flawed.

This discussion leads to the second objection. One may grant that Hyde does not know, but deny premise (7); namely, argue that Hyde is in a Gettier case. Thus, Hyde’s failure to know is *not* a justification failure, and Hyde does individually lack KL-justification. Orthodox individualism explains why Jekyll knows while Hyde does not, and KLJ-communalism is redundant. Against this, I argue that Hyde is not gettiered for two reasons. First, Hyde’s case lacks an essential characteristic of Gettier cases; namely, their coincidental nature. ‘Knowledge must somehow not depend on coincidence or luck. This was just the point of the Gettier counter-examples’ (Dancy 1985: p. 134). There is no coincidence in Jekyll/Hyde-like cases, because scientists in them are disposed to report and be believed that p , whether they have evidence for p or not. Moreover, coincidences are rare. Unlike Hyde, in most nearby possible worlds, a gettiered subject has a false belief rather than a true one (Pritchard 2005: p. 145-152). Hyde-like cases are common. They are not mere flukes, but a systematic by-product of the institutional design of science and scientists’ psychological makeup.

Second, claiming that Hyde-like cases are Gettier cases leads to a wrong normative appraisal thereof. In Gettier cases, nobody is blameworthy, or they can be reconstructed such that nobody is blameworthy, e.g. such that nobody has given incompetent or insincere testimony or engaged in illegitimate inference. A gettiered subject is a victim of pure misfortune, whereas in Hyde-like cases, scientists are epistemically responsible for sloppy or fraudulent research and unreliable reporting.

Correctly assigning epistemic praise and blame is necessary for reaching truth, avoiding error, and resisting sceptical challenges when scepticism is inappropriate. Correct epistemic normative evaluation is required for enacting epistemic checks, such as double-blind review,

and mandatory disclosure of conflicts of interest. We must know how inquiry may go wrong for setting it right. Alluding to such safety measures is required for responding to sceptical challenges. For example, one may claim to know that p by stating that one read that p in peer-reviewed journal with a mandatory disclosure policy. If Hyde-like cases are treated as Gettier cases, we cannot detect epistemic failures, thus we cannot implement epistemic measures to prevent them, thus we cannot respond to sceptical challenges in cases we should be able to have knowledge. KLJ-communalism avoids such scepticism by correctly assigning epistemic praise and blame to relevant members of the epistemic community. Denying that Jekyll knows that q or claiming that Hyde is gettiered, then, amounts to scepticism and other unhappy consequences. The next section addresses the objection that Hyde and Jekyll may differ in their respective beliefs or the evidence they individually possess.

VI. WHY INDIVIDUALS TYPICALLY DO NOT HAVE AND ARE UNABLE TO ACQUIRE KL-JUSTIFICATION FOR THEIR SUSPECT BELIEFS

This section addresses two lines of objection to KLJ-communalism. According to the first, when S forms a belief that p based on an expert's testimony that p at time t , typically S has indirect evidence at t about the expert's sincerity and competence that allows S to justifiably believe or disbelieve that p even without access to the expert's evidence. This objection amounts to denying premise (1), insisting that Jekyll and Hyde differ in the evidence they individually possess just after encountering Eddington's testimony, and while Jekyll's evidence inclines him to believe that q , Hyde's evidence does not incline him to believe that q .

According to the second line of objection, even if S does not have evidence at t to justifiably believe or disbelieve that p , with reasonable effort, S can typically gather further evidence that will eventually allow S to form a KL-justified belief. Namely, S 's not knowing is a temporary state, and in the long run, S is able to, and typically will personally acquire enough evidence that gives S KL-justification for believing or disbelieving the original testimony. Such evidence consists of multiple confirmations of the original testimony and independent indicators of the expert's reliability. This objection amounts to claiming that while at first, Hyde may believe that q without personally having KL-justification, this is a temporary state, which is not epistemically alarming.

I will now argue that in Jekyll-like cases, evidence subjects have at t or are able to acquire later is typically insufficient for giving them KL-justification *qua* individuals. This is because the evidence is significantly insensitive to falsehood: Based on it, a subject would believe that p even in a likely scenario in which p were false.

What evidence, then, does an individual have at the time he hears an expert testimony, which allegedly allows him to form a KL-justified belief? As mentioned in Section III, Fricker (2002) argues that KL-communalism can be resisted because individuals, *qua* individuals, have sufficient *indirect* evidence for their suspect beliefs that amounts to KL-justification at the time they acquire them. Fricker focuses on how individual scientists may justify their trust in their peers. Fricker argues that a scientist has plenty of evidence about his peers' competence and trustworthiness. Such evidence draws on his personal acquaintance with his peers, knowledge of their role and status in society, and their commitment to norms of competence and truth-telling. This evidence is indirect, because it is about the testifiers, rather than the content of their testimonies.

Can such indirect evidence give individuals KL-justification for their suspect beliefs in Jekyll-like cases? Fricker assumes that scientists' appearing trustworthy, e.g. having relevant credentials, is correlated with them actually being so. But how can a person have KL-justification for believing that the two are correlated? How can scientists know that their seemingly respectable and trustworthy colleagues are actually so? They must individually verify at least some of their peers' claims to establish their reliability.

How practicable is this option? Let us return to the Eddington example. When Eddington announced his result, his colleagues took his word for it, and did not examine the evidence themselves, as the following quote by Sir Joseph Thomson, president of the Royal Society, demonstrates:

It is difficult for the audience to weigh fully the meaning of the figures that have been put before us, but the Astronomer Royal [Dyson] and Professor Eddington have studied the material carefully, and they regard the evidence as decisively in favour of the larger value for the displacement. (quoted in Earman and Glymour 1980: p. 77)

Almassi (2009) attempts to defuse this example by arguing that Eddington's colleagues had the required expertise to evaluate his evidence; hence, they did not blindly trust his testimony. Almassi's response, however, misses its target. First, unlike Eddington's colleagues, researchers usually do not have access to their peers' raw data. Their reliable sense of the reliability of a peer's outcomes depends on the largely unverifiable assumption that he has not infringed accepted evidential standards (Wilholt 2009). But these standards are exactly those fraudulent or sloppy researchers infringe in Hyde-like cases.

Second, scientists, Eddington's colleagues included, cannot verify *every* scientific claim they encounter. Because scientists' time and resources are limited, they must still rely on proxies, such as a researcher's reputation, to decide whom to trust. Jekyll-like cases do not raise any suspicion that the researcher would have made the same testimony even without evidence. Thus, scientists would likely still rely on social proxies in Jekyll-like cases. But, in Jekyll-like cases, social proxies are insufficient to accord individuals KL-justification. They are even misleading: Scientists with good reputation are more tempted to make unsubstantiated testimonies exactly because their good reputation makes it less likely that someone would check on them.

I argued, contra Fricker, that indirect evidence individuals have when they encounter an expert testimony in Jekyll-like cases is typically insufficient for KL-justification. Will evidence subjects acquire later amount to KL-justification at some future point? I now argue that for suspect beliefs acquired in Jekyll-like cases, the answer is negative. I focus on two types of evidence: multiple testimonial confirmation, and independent reliability indicators.

Regarding multiple testimonial confirmations, Adler (1994) replies to Hardwig that within a reasonable time, a person is able to possess sufficient justification by gaining multiple confirmations of the same report, as they tend to converge on the truth in the long run. My reply to Adler will be threefold. First, the prevalence of multiply-confirmed false testimonies, including in science, significantly reduces their justificatory value. Second, empirical evidence militates against the convergence-on-truth thesis. Third, in many Jekyll-like cases, multiple confirmations are useless because they originate from the same suspect source.

First, many *false* testimonies enjoy multiple confirmations. Distinguishing them from reliable multiple confirmations is hard. This severely hinders their justificatory value over time. Urban legends, for example, are impossible to kill. There is always someone who swears that one of them actually happened to a friend's friend. They persistently pop up in mainstream news sources as true stories (Coleman 1979; Best and Horiuchi 1985; Smith 1992). Webb (2004) convincingly argues that formulating a practicable method individuals can employ for distinguishing between multiply-confirmed true and false report seems hopeless. He examines several proposed principles for discerning reliable from unreliable testimonies (e.g. 'never trust a story about a friend's friend'; 'never trust 'too good a story') and shows that they are all too permissive or too restrictive.

Science is hardly immune to urban legends. According to a study that tracked enduring false stories in scientific papers,

Many of the messages presented in respectable scientific publications are, in fact, based on various forms of rumors. Some of these rumors appear so frequently, and in such complex, colorful, and entertaining ways that we can think of them as academic urban legends (Rekdal 2014: p. 638).

The study finds that once a colourful false story is reported as true in a respectable journal, it tends to be continually reported in subsequent publications, and even trickle to popular media. The publications that continue to disseminate it tend *not* to cite the secondary source from which they picked it up, but the alleged original source without consulting it. 'The explanation for this phenomenon is usually that authors have lazily, sloppily, or fraudulently employed sources, and peer reviewers and editors have not discovered these weaknesses in the manuscripts during evaluation'. This is 'an unfortunate side effect of publication pressure and competition for academic positions and scarce resources' (Rekdal 2014: p. 638). Having instructed philosophy-of-science courses, I am personally struck by the difficulty of eradicating debunked myths about the history of science. They persistently pop up in science and engineering courses, *Wikipedia* articles, and students' exams, even after learning the truth. I suspect that some of them endure not just because of poor scholarship, as Rekdal suggests, but because they play an ideological role in scientists' identity.

A second problem with Adler's argument is that empirical research militates against the convergence-on-truth thesis. Lewenstein (1995) tracks testimonies, emails, faxes, and news stories physicists received during the 1989 Cold Fusion affair, which he regards a typical case of information dissemination under uncertainty. He argues that the numbers of confirming and disconfirming reports fluctuate with time. At some point, a claim seems confirmed and later disconfirmed. Sometimes people lose interest in a story and stop talking about it. Successive reports do not always converge on the truth.⁷

Because of the ubiquity of false multiple confirmations, subjects' limited ability to distinguish them from true confirmations, and the fact that multiple confirmations need not converge on the truth in the long run, multiple confirmations are much less epistemically valuable for justifying subjects' beliefs in testimonies over time than Adler assumes. One may argue, however, that multiple confirmations still give subjects, *qua* individuals, KL-justification over time in Jekyll-like cases. In science, so this objection goes, the accumulation of empirical research over time gives KL-justification to a belief in a scientific hypothesis. My argument overplays the significance of a single experiment for reaching KL-justification, and falsely states that Jekyll knows merely by reading Eddington's testimony in the paper. If ever, Jekyll is KL-justified believing that q only after GR is put to further empirical tests, or so this objection goes.⁸

⁷ Elsewhere (Miller 2009), I disagree with Lewenstein's claim that relativism about scientific knowledge follows from his example.

⁸ I thank an anonymous reviewer for pressing me on this point. A common view in current philosophy of science is that sceptical worries from underdetermination of theory by evidence are inflated, because even if two theories logically accommodate the same evidence, they do not necessarily enjoy the same inductive warrant, thus a rational decision between them is possible. It is also argued that finding a genuine alternative to a theory, which is not just a parasitic variant of an existing theory is not trivial (Laudan 1990; Laudan and Leplin 1991; Norton 2008; Galison 1997: pp. 14-19). The notion of a crucial experiment has also been rehabilitated. While it is no longer believed that an experiment can decide, once and for all, between two competing theories, it is argued that an experiment can still decide between two well-defined available models (Giere 1999: pp. 123-128). We may also interpret Jekyll/Hyde's belief that GR is correct as stating that GR is empirically adequate or approximately true, which sets a lower bar for reaching KL-justification than the belief that GR is true simpliciter.

As I will argue, this objection assumes an unrealistic model of science. Many important experiments or proofs are performed once, and scientists routinely and legitimately make knowledge claims based on them, including in many Jekyll-like cases. Thus, in many Jekyll-like cases, all seemingly independent subsequent confirmations, if any, originate from the same testimonial source. Yet subsequent confirmations carry justificatory weight only insofar as they are independent (Goldman 2001: pp. 99-102; Miller 2013: pp. 1304-1311). As Wittgenstein (1967: Section 265) puts this, relying on subsequent confirmations in many Jekyll-like cases is like buying several copies of the morning paper to assure ourselves that what it said was true.

Quite a few scientific hypotheses are supported by evidence from a single experiment or proof due to the phenomenon of ‘one-shot science’, analyzed by Brown (2010). Brown draws two distinctions regarding the justification of scientific claims. The first is between narrow and wide criteria of evaluation. Some claims are supported by multiple lines of evidence, e.g. an event witnessed by many people, filmed, and recorded. Others have only one line of evidence, e.g. a mathematical theorem with a single proof. Brown’s second distinction is between low-cost and high-cost evaluation. The cost of evaluating a one-page proof is low, while evaluating a complex computer proof that constructs and tests thousands of cases is costly. ‘One-shot science’ covers cases where the evaluation criteria are narrow and the evaluation cost is high. In one-shot science, a claim is proved ‘once and never again’ (2010: p. 96). Examples of one-shot science include complex mathematical computer proofs, large clinical trials, experiments for detecting elementary particles, and experiments for detecting new quantum-phase transitions, where experiments are sometimes not even run once(!) but replaced with a computer simulation due to cost (Tal 2011).

The above distinctions are in degree. The evidence for a claim may be wider or narrower, costlier or cheaper. They are also relative to the state of science and technology at a given time. For example, with the improvement of computer technology, some computer proofs may become easier to implement and replicate. They are also agent relative. Some evidence is accessible to some people and not others, and some techniques are affordable to some bodies and not others. It might therefore be objected that in one-shot science, we should wait until the conditions change and allow a hypothesis to be multiply confirmed before the belief in it can gain KL-justification.

This objection fails for two reasons. First, cases at the end of the continuum of these distinctions are *in principle* one-shot science. Most computer proofs are too long for any mathematician to go over and verify, and some experiments or proofs are too expensive to replicate even for the wealthiest bodies. Second, this objection cuts both ways: It stresses the relevance of actual contingent conditions to justification, which are far from ideal. In principle, multiply confirming many hypotheses is possible and affordable. Yet *de facto*, because researchers’ time and money are limited, reconfirming them is never high enough on anybody’s priorities, and they end up being confirmed once. For example, in 2013, an economics graduate student found major calculation errors in an influential study by two Harvard economists, according to which high public debt inhibits economic growth. That one graduate student found these errors is the exception that proves the rule. None of the *thousands* of economists who cited the paper, including in World Bank and national banks reports, checked the data or noticed the errors (Roose 2013). *De facto* one-shot science is wider than in-principle one-shot science, and the *de facto* class matters for the justification of our actual beliefs.

It might be objected that the right conclusion is scepticism toward all one-shot science. But while scepticism may be warranted toward some one-shot science, such as portions of economics or industry-funded clinical trials, one-shot science exceeds these domains. Should we also adopt scepticism toward much of mathematics, high-energy physics, and all other experiments that are *de facto* performed once, even if adequately and rigorously? I argued in Section V that sweeping scepticism about scientific knowledge is uncalled for. Scepticism toward one-shot science is similarly unwarranted. It raises the question of why one-shot science is

performed to begin with. Why has the European Union invested billions of Euros building the Large Hadron Collider if it cannot give us knowledge of the existence of the Higgs boson?⁹ Are physicists hopelessly misguided when they claim to know it exists? Sweeping scepticism is too quick, and does not take seriously scientists' claims to knowledge.

By contrast, KLJ-communalism takes scientists' claims seriously, and passes a correct verdict on them. It grants them the status of knowledge only if they are true and have KL-justification. According to KLJ-communalism, a belief based on the results of a single experiment is knowledge only when it is true and the evidence produced by the experiment is sufficient to rule out all reasonable alternative hypotheses, and non-knowledge when the evidence is not. Such evidence may be dispersed among many researchers, or concentrated in the hands of few, but need not be confined to the believing subject. Thus, the Harvard economists' claim that high public debt inhibits economic growth is unknown, at least by the evidence they presented in their debunked paper, while the existence of the Higgs boson is known, at least if the evidence for it is as strong as physicists claim.

I argued that multiple confirmations are not available in one-shot-science cases, yet scientists legitimately make knowledge claims in them. For showing that the objection from multiple confirmations fails, however, I need to show that there is a significant overlap between one-shot-science and Jekyll-like cases; namely, that in many one-shot-science cases, the produced evidence supports the researchers' claims, but they would make the same claim even if the evidence were flawed or weak. In such overlapping cases, multiple confirmations would all ultimately originate from the same source, and would not give a lay subject KL-justification.

Indeed, we should expect a significant overlap between one-shot-science and Jekyll-like cases. In Jekyll-like cases, scientists would probably testify to a hypothesis even if their evidence were shaky, because they already believe or have a strong interest in confirming it. In one-shot science, scientists usually know their procedures will probably not be replicated. People are more tempted to misbehave when they think they can get away with it. If we add these three facts together, it follows that in one-shot-science cases in which the researchers are already inclined toward a certain hypnosis they would be more tempted to report it even without good evidence than in cases where they think their results may be checked. One-shot science lacks a major deterrent (the possibility of replication) that guards other cases from becoming Jekyll/Hyde-like cases.

Is there other evidence a subject can gather in the long run, which can give her KL-justification for her suspect beliefs? Goldman thinks there is. Replying to Hardwig, he argues that while laypersons cannot evaluate experts' evidence, they are still able to obtain justification for believing them, as some experts' statements are independently verifiable by laypersons or may become so over time. By following experts' truth track record, a layperson is able to establish their reliability and justifiably believe their testimony.

Goldman lists several examples of reliability indicators. When an astronomer predicts an eclipse, laypersons cannot evaluate his prediction, but once the eclipse happens (or does not happen) it becomes verifiable by laypersons. Similarly, while laypersons cannot evaluate the reasons for a forecaster's rain predictions, they can see for themselves if it rains. Other examples are successful medical treatments and mechanical repairs. Goldman argues that physicians and car mechanics are successful mainly because they have true beliefs. Their success records give laypersons justification to believe their testimonies (1999: pp. 79-82; 270-271; 2001: pp. 106-108).

⁹ To slightly overcome this problem, the Higgs boson detection claim was based on the results of two experiments performed by two independent teams. Both teams still used the same collider (Overbye 2013).

As Goldman admits in a footnote, his reply to Hardwig is lacking, since undisputable reliability indicators are scarce:

It must be conceded that solving this fundamental conundrum is not a panacea for problems of authority attribution, especially at the frontiers of science. At the frontiers, it is very difficult to pinpoint an outcome that can serve as an uncontroversial 'verification' (Goldman 1999: pp. 271).

Goldman's concession takes the wind out of his suggestion, since for defending individualism, it is insufficient to show that individuals can *in principle* gain reliability indicators, but that they are able in practice to obtain reliability indicators for suspect beliefs acquired in Jekyll-like cases. Goldman does not establish that.

The major reason uncontroversial reliability indicators are indeed scarce in theoretical science is that on their own, successful empirical results or material interventions often do not uniquely indicate the truth of the theory being tested or relied upon. A theory may contain just a grain of truth, which is responsible alone for its successful application; e.g. until the nineteenth century, excellent steel was manufactured by combining magical exorcism with the operations of the craft. Accuracy requirements in practice may fall far below the true values, allowing false values to suffice for application. Successful results may be obtained at the ontological level of reality to which a theory is applied, without fully knowing the deep structures and relations at the ontological levels it purports to describe.¹⁰ In practice, several measures are usually attempted together and other potential influences remain present; therefore, if the outcome is satisfactory, it remains unknown which measure, if any, was successful (Bunge 1998: pp. 140-143).

While a successful single experiment *may* give a subject KL-justification for believing a hypothesis, it *must* be a careful experiment that overcomes the above problems. Knowing whether a reliability indicator supports a theory to KL-justification requires knowing it was attained under the right background assumptions and adequate controls. In Jekyll-like cases, laypeople typically lack such knowledge. Thus, reliability indicators typically cannot give subjects, *qua* individuals, KL-justification in Jekyll-like cases.

It might be objected that reliability indicators need not establish the specific hypothesis in question, but the expert's track record. For example, an astronomer's predictions of independently verifiable astronomical events establish his trustworthiness regarding non-independently verifiable claims. But the Eddington case teaches us better. Relying on his excellent reputation as an experimenter is exactly what allowed Eddington to get away with making an unsubstantiated testimony. This phenomenon is common in many domains. For example, a major problem Internet reputation systems must handle is that users with high reputation allow themselves occasional trust violations because they know this would hardly affect their reputation (Simpson 2011: p. 36).

VII. RELIABILISM TO THE RESCUE?

I cashed out my argument in evidentialist terms. One may thus falsely think that the individualism/communalism debate is an in-house debate between internalists, because it assumes that knowledge must be supported by reasons accessible to subjects by introspection. Under this requirement, only the community as a whole has such access, therefore individualism fails. Since externalists make no such requirement, they seemingly avoid the communalist

¹⁰ These claims resonate with anti-realist arguments in philosophy of science that doubt the connection between the truth of a scientific theory and its empirical success.

dilemma (Palermos 2012: pp. 118-119). I will now argue, however, that the problem with individualism is invariant under any internalist or externalist justification account. The truth-yielding factors externalists regard responsible for a normal subject's beliefs' having KL-justification sometimes extend beyond her bodily boundaries into other members of her epistemic community.

My argument lends support to Goldberg's (2010) extended reliabilism, but takes it a further step away from individualism. Goldberg argues that orthodox reliabilism misrepresents the extent and nature of our epistemic reliance on others. Goldberg identifies an unwarranted assumption of process reliabilism, which is that the cognitive processes involved in belief formation never extend beyond the individual believer's bodily boundaries. Therefore, factors outside her bodily boundaries are considered part of the background environment, rather than relevant factors to a belief's reliability. As an alternative, Goldberg makes the *extendedness hypothesis*, according to which a belief-forming process may extend beyond a person's cognitive system into other people's cognitive systems, such that salient elements responsible for its reliability may be located at temporally and spatially remote ends of the process. Such extended processes include long testimonial chains, whose reliability is determined by the competence of the person who initiated them. My argument goes a step further away from individualism than Goldberg's. Goldberg still accords a privileged role to the belief-forming subject at the end of an extended cognitive process, whereas I hold that this role is not always maintainable.

Recall the Eddington example to understand why orthodox process reliabilism (and without loss of generality, any individualist version of reliabilism, such as virtue reliabilism) faces a parallel problem to evidentialism. How can orthodox process-reliabilists explain the difference between Jekyll and Hyde? Recall from Section V that they cannot claim that Hyde is gettiered or deny that Jekyll knows. Process-reliabilists may explain this difference by appealing to the environment, just like the different environments explain why a subject in Fake-Barn Country, who believes there is a barn on the hill, and happens to look at the only real barn around, does not know, while a similar subject who forms the same belief in Real-Barn Country knows. This attempt, however, will fail. The only difference between Jekyll and Hyde's worlds is that in Jekyll's world, Eddington's team conducted the experiment successfully and reached correct results, while in Hyde's (our) world, the experiment failed, Eddington fudged the data, and reported unsupported results. This difference does not result in any noticeable difference in the immediate environment to which Jekyll and Hyde's cognitive systems respond when they read Eddington's testimony in the newspaper, and form the belief that *q*.

Individualistically-minded process-reliabilists may insist that Jekyll and Hyde's testimonial belief-forming processes *are* able to reliably discern trustworthy and untrustworthy testimonies, perhaps not immediately on reading Eddington's testimony in the newspaper, but over time, when other testimonies accumulate. They may argue that social cues in the situation might help subjects in Jekyll/Hyde-like cases reliably discern trustworthy and untrustworthy expert testimonies. Presumably, such cues would be multiple confirmations, indirect social evidence, and external reliability indicators.

As I argued, however, such indications are insufficient for giving subjects' suspect beliefs KL-justification in Jekyll-like cases because *qua* evidence they are insensitive to a genuine possibility of the beliefs' falsehood. But if they are insufficient as evidence, they are just as insufficient as input data for subjects' cognitive processes for reliably distinguishing between true and false testimonies. It does not matter whether such indications are thought of as evidence or as sensory input data for subjects' belief-forming processes. The information contained in them is typically an insufficient basis for reliably discerning true and false expert testimonies in Jekyll/Hyde-like cases. Therefore, process-reliabilists must concede to KLJ-communalism.

This is where my account diverges from Goldberg's. Because Goldberg wants to keep individual epistemic responsibility, he still accords a special role to the recipient of testimony at

the end of an extended cognitive process. Goldberg requires that she filter testimonies based on their credibility and accept only those that amount to knowledge; i.e. form testimonial beliefs only when they reach KL-justification. This requirement is relative to the individual's environment. Young children, who are protected by their parents from misleading reports or members of an ideal community where extreme sincerity and competence are the norm, need not filter reports to the same degree as normal adults in normal environments (2010: pp. 110-111; 2007: pp. 209-226).

Goldberg's requirement explicitly rests on a dubious empirical assumption, which is that testimonial belief-forming cognitive processes

are in fact highly reliable in that, in cognitively mature and well-functioning adults, they issue in acceptance in most of the cases where the testimony is both true and reliable, and they issue in non-acceptance in most cases where the testimony is false or otherwise unreliable (2007: p. 148).

My argument entails that this assumption is false for suspect beliefs subjects acquire from expert testimony, and are unable to directly evaluate themselves, at least in Jekyll/Hyde-like cases. Normal subjects in normal circumstances cannot always meet Goldberg's requirement.

This failure does not severely hinder Goldberg's framework. People's reliability in accepting testimony is a contingent empirical matter, which cannot be decided from the armchair.¹¹ Goldberg can require, without relinquishing his overall point, that subjects only do their best to accept a testimony only when it amounts to knowledge. But at least regarding suspect beliefs acquired in Jekyll-like cases, Goldberg needs to resort to KLJ-communalism in that the knowledge in question is the property of a community but none of its members.¹²

Externalists should not resist KLJ-communalism, because they regard justification external to the subject's internally-accessible mental world, and a belief's causal history relevant to its justificatory status. It is a contingent empirical matter which elements in a belief's causal history have the greatest influence its justificatory status. It cannot be decided *a priori* that the causal history that warrants philosophical attention is confined to the believing subject's own cognitive system.

VIII. CONCLUSION

Knowledge is the property of a community but none of its members in that sometimes only the community possesses knowledge-level justification for it. The metaphysical building blocks of doxastic justification, such as evidence or segments of a reliable cognitive process, extend beyond the believing subject's bodily or mental boundaries into other members of her epistemic community. Sometimes, only because these members possess these justificatory building blocks does the subject's belief reach knowledge-level justification. Only knowledge-level communalism draws the right demarcation between knowledge and non-knowledge for a large class of expert-obtained beliefs. If epistemology is to deal with live sceptical doubts about our

¹¹ For a sustained attack on the monitoring-of-testimony requirement, which draws on empirical evidence from experimental psychology that demonstrates subjects' *de facto* unreliability in accepting testimony, see Shieber (2011).

¹² Goldberg (2013) relaxes the monitoring-of-testimony requirement with respect to young children because a stringent monitoring requirement would not allow them to acquire knowledge from adults' testimonies. If my argument is right, the requirement should be relaxed regarding adults too. .

actual beliefs it should give up its individualistic dogmatism and adopt knowledge-level communalism.¹³

REFERENCES

- Adler, J. (1994) 'Testimony, Trust, Knowing', *Journal of Philosophy*, 91(5), 264-275.
- Almassi, B. (2007) 'Experts, Evidence, and Epistemic Independence', *Spontaneous Generations*, 1(1), 58-66.
- . (2009) 'Trust in Expert Testimony: Eddington's 1919 Eclipse Expedition and the British Response to General Relativity', *Studies in History and Philosophy of Modern Physics*, 40(1), 57-67.
- Angell, M. (2009) 'Drug Companies and Doctors: A Story of Corruption', *New York Review of Books*, 56(1), 8-13.
- Best, J. and Horiuchi, G. (1985) 'The Razor Blade in the Apple: The Social Construction of Urban Legends', *Social Problems* 32(5), 488-499.
- Bird, A. (2010) 'Social Knowing: The Social Sense of "Scientific Knowledge"', *Philosophical Perspectives*, 24, 23-56.
- Brown, J.R. (2010) 'One-Shot Science', in H. Radder (ed.) *The Commodification of Academic Research*, 90-109. Pittsburgh: University of Pittsburgh Press.
- Bunge, M. (1998) *Philosophy of Science: From Explanation to Justification*. New Brunswick, NJ: Transaction.
- Burge, T. (1993) 'Content Preservation', *Philosophical Review*, 102(4), 457-488.
- Coady, C.A.J. (1992) *Testimony: A Philosophical Study*. Oxford: OUP.
- Coleman, L. (1979) 'Alligators-in-the-Sewers: A Journalistic Origin', *Journal of American Folklore*, 92(365), 335-338.
- Collins, H. and Pinch, T. (1993) *The Golem: What Everyone Should Know About Science*. Cambridge: CUP.
- Corlett, A.J. (1996) *Analyzing Social Knowledge*. Lanham, MD: Rowman and Littlefield.
- Dancy, J. (1985) *An Introduction to Contemporary Epistemology*. Oxford: Blackwell.
- Dretske, F. (1981) 'The Pragmatic Dimension of Knowledge', *Philosophical Studies*, 40(3), 363-378.
- Earman, J. and Glymour C. 1980. 'Relativity and Eclipses: The British Eclipse Expeditions of 1919 and their Predecessors', *Historical Studies in the Physical Sciences*, 11(1), 49-85.

¹³ I thank Anjan Chakravartty, Jim Brown, Joseph Berkovitz, Sandy Goldberg, Yves Gingras, Eran Tal, Arnon Keren, Levi Spectre, Hagit Benbaji, Ruth Weintraub, Ellie Louson, and the audience members at the many events I presented versions of this paper for helpful comments and criticisms. I am grateful to the Azrieli Foundation for an award of an Azrieli Fellowship. I thank the Edmond J. Safra Center for Ethics, Tel Aviv University, the Dan David Foundation, Tel Aviv University, the Cohn Institute for the History and Philosophy of Science and Ideas, Tel Aviv University, and the Sidney M. Edelstein Center for the History and Philosophy of Science, Technology, and Medicine, the Hebrew University of Jerusalem, for postdoctoral fellowships.

- Frans, J. and Kosolosky L. (2014) 'Mathematical Proofs in Practice: Revisiting the Reliability of Published Mathematical Proofs', *Theoria* 81, 345-360.
- Fricker, E. (2002) 'Trusting Others in the Sciences: *a priori* or Empirical Warrant?', *Studies in History and Philosophy of Science*, 33(2), 373-383.
- Galison, P. (1997) *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press.
- Geist, C., Löwe B. and Van Kerkhove, B. (2010) 'Peer Review and Knowledge by Testimony in Mathematics', in B. Löwe and T. Müller (eds.) *Philosophy of Mathematics: Sociological Aspects and Mathematical Practice*, 155-178. London: College Publications.
- Giere, R. (1999) *Science without Laws*. Chicago: University of Chicago Press.
- . (2006) *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Gilbert, M. (1987) 'Modeling Collective Belief', *Synthese*, 73(1), 185–204.
- Goldberg, S. (2007) *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: CUP.
- . (2010) *Relying on Others: An Essay in Epistemology*. Oxford: OUP.
- . (2013) 'Epistemic Dependence in Testimonial Belief, in the Classroom and Beyond', *Journal of Philosophy of Education*, 47(2), 168-186.
- Goldman, A.I. (1999) *Knowledge in a Social World*. Oxford: OUP.
- . (2001) 'Experts: Which Ones Should You Trust?', *Philosophy and Phenomenological Research*, 63(1), 85-111.
- Goodstein, D. (2010) *On Fact and Fraud: Cautionary Tales from the Lines of Science*. Princeton: Princeton UP.
- Grcar, J.F. (2010) 'Errors and Corrections in Mathematics Literature', *Notices of the American Mathematical Society*, 60(4), 418-425.
- Green, A. (2013) 'Deficient Testimony is Deficient Teamwork', *Episteme*, 11(2), 213-227.
- . (2014). 'Evaluating Distributed Cognition', *Synthese*, 191(1), 79-95.
- Green, L. (1991) 'Two Views of Collective Rights', *The Canadian Journal of Law and Jurisprudence*, 4(2), 315-327.
- Hardwig, J. (1985) 'Epistemic Dependence', *Journal of Philosophy*, 82(7), 335-349.
- Kennefick, D. (2009) 'Testing Relativity from the 1919 Eclipse – A Question of Bias', *Physics Today*, 62(3), 37-42.
- Kerr, E. and Gelfert A. (2014) 'The "Extendedness" of Scientific Evidence', *Philosophical Issues*, 24, 253-281.
- Kitcher, P. (1994) 'Contrasting Conceptions of Social Epistemology', in F.F. Schmitt (ed.) *Socializing Epistemology: The Social Dimensions of Knowledge*, 11-134. Lanham, MD: Rowman and Littlefield.
- Kusch, Martin. (2002) *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: OUP.
- Laudan, L. and Leplin J. (1991) 'Empirical Equivalence and Underdetermination', *Journal of Philosophy*, 88(9), 449-472.
- Laudan, L. (1990) 'Demystifying Underdetermination', in C.W. Savage (ed.) *Scientific Theories*, 267-297. Minneapolis: University of Minnesota Press

- Lewenstein, B. (1995) 'From Fax to Facts: Communication in the Cold Fusion Saga', *Social Studies of Science*, 25, 403-436.
- List, C. and Pettit P. (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: OUP.
- Longino, H. (2002) *The Fate of Knowledge*. Princeton: Princeton UP.
- Martinson, B., Anderson M. and de Vries R. (2005) 'Scientists Behaving Badly', *Nature*, 435(9), 737-738.
- Miller, B. (2009) 'What Does It Mean that PRIMES is in P? Popularization and Distortion Revisited', *Social Studies of Science* 39(2), 257-288.
- . (2013) 'When Is Consensus Knowledge Based? Distinguishing Shared Knowledge from Mere Agreement', *Synthese*, 190(7), 1293-1316.
- Nathanson, M. (2008) 'Desperately Seeking Mathematical Truth', *Notices of the American Mathematical Society*, 55(7), 773.
- Nelson, L.H. (1993) 'Epistemological Communities', in L. Alcoff and E. Potter (eds.) *Feminist Epistemologies*, 121-159. London: Routledge.
- Norton, J (2008) 'Must Evidence Underdetermine Theory?', in M. Carrier, D.A. Howard and J. Kourany (eds.) *The Challenge of the Social and the Pressure of Practice: Science and Values revisited*, 17-44. Pittsburgh: University of Pittsburgh Press
- Overbye, D. (2013) 'Chasing the Higgs', *New York Times*, March 5, D1.
- Palermos, O.S. and Pritchard D. (2013) 'Extended Knowledge and Social Epistemology', *Social Epistemology Review and Reply Collective*, 2(8), 105-120.
- Palermos, O.S. (2012) *Extending Cognition in Epistemology: Toward an Individualistic Social Epistemology*. Ph.D. thesis, University of Edinburgh.
- Pritchard, D. (2005) *Epistemic Luck*. Oxford: OUP.
- Réaume, D. (1988). 'Individuals, Groups, and Rights to Public Goods', *University of Toronto Law Journal*, 38, 1-27.
- Rekdal, O.B. (2014) 'Academic Urban Legends', *Social Studies of Science*, 44(4), 638-654.
- Roose, K. (2013) 'Meet the 28-Year-Old Grad Student Who Just Shook the Global Austerity Movement', *New York Times Magazine*, April 18, <<<http://nymag.com/daily/intelligencer/2013/04/grad-student-who-shook-global-austerity-movement.html>>>
- Schmitt, F. (1988) 'On the Road to Social Epistemic Interdependence', *Social Epistemology*, 2(4), 297-326.
- Shieber, J. (2011a) 'Against Credibility', *Australasian Journal of Philosophy*, 90 (1), 1-18.
- . (2011b) 'Toward a Truly Social Epistemology: Babbage, the Division of Mental Labor, and the Possibility of Socially Distributed Warrant', *Philosophy and Phenomenological Research*, 86(2), 266-294.
- Simpson, T. (2011) 'e-Trust and Reputation', *Ethics and Information Technology*, 13(1), 29-38.
- Sismondo, S. (2009) 'Ghosts in the Machine: Publication Planning in the Medical Sciences', *Social Studies of Science*, 39(2), 171-198.
- Smith, P. (1992) 'Read All About It! Elvis Eaten by Drug Crazy Alligators: Contemporary Legend and the Popular Press', *Contemporary Legend*, 2, 41-70.

- Tal, E. (2011) 'From Data to Phenomena and Back Again: Computer-Simulated Signatures', *Synthese*, 182(1), 117-129.
- Tuomela, R. (2011) 'An Account of Group Knowledge', in H. Schmid, M. Weber, and D. Sirteseds (eds.) *Collective Epistemology*, 75-117. Heusenstamm: Ontos.
- Waller, J. (2002) *Fabulous Science: Fact and Fiction in the History of Scientific Discovery*. Oxford: OUP.
- Webb, M. (2004) 'Can Epistemology Help? The Problem of the Kentucky-Fried Rats', *Social Epistemology*, 18(1), 51-58.
- Wilholt, T. (2009) 'Bias and Values in Scientific Research', *Studies in History and Philosophy of Science*, 40(1), 92-101.
- Wittgenstein, L. (1967) *Philosophical Investigations*, 3rd edition. Oxford: Blackwell.
-

¹ *The Hebrew University of Jerusalem, and Bar Ilan University, Israel*