

Knowledge and Evidence You Should Have Had

Matthew A. Benton

1 Introduction

Epistemologists focus primarily on cases of knowledge, belief, or credence where the evidence which one possesses, or on which one is relying, plays a fundamental role in the epistemic or normative status of one's doxastic state. Thus one's belief might be rational, or justified, or amount to knowledge (if true and other conditions are met); or in a partial belief framework, one's credence might be rational or optimal given one's prior probabilities, and it might also be scored as more or less "accurate" given the truth or falsity of the propositions to which one gives credences.

Recent work however goes beyond this evidence to consider the relevance for such statuses of evidence which one does *not* possess,¹ particularly when there is a sense in which one *should have had* some evidence. My focus here will be Sanford Goldberg's ([forthcominga](#), [forthcomingb](#)) ambitious approach, but the discussion should be of interest to anyone who is partial to what is sometimes called "normative defeat."

Goldberg talks at length of a subject's belief being "justified." His interest is in how the evidence one should have had with respect to some proposition p , particularly when one has higher-order evidence to the effect that there is further evidence one should acquire, can affect the epistemic justification one has for one's belief that p even if one does not acquire that further evidence. Goldberg tells us that what he means by "jus-

¹See for example Kornblith 1983, Lackey 2008, 45–46, Ballantyne 2015, Goldberg [forthcomingb](#), [forthcominga](#).

tification” is doxastic justification, that it comes in degrees, and that it has to do with evidence (though he does not tell us what exactly “evidence” is²). On Goldberg’s approach then, justification often seems to be a notion which encodes epistemic *permission* with respect to outright belief, and he acknowledges that he wants to think of such epistemic justification as necessary for knowledge.

The gradability of justification given the changes to one’s evidence makes it seem best modelled in a probabilistic framework. Because Goldberg’s principles discussed below make appeals to a belief being *more* or *less* justified had a subject acquired some further evidence, and because a prominent way of modelling the evidence-for relation is itself done in a probabilistic framework, I shall begin (§2) by attempting to elucidate some of his ideas using such a framework. Though such elucidations may ultimately reveal that a probabilistic framework is not well-suited for modelling the kind of normative defeat in which Goldberg is interested, it will be instructive to see why. §3 considers Goldberg’s account of what grounds the “should” or “ought” in play here, and introduces some difficulties for the conditions under which this obligation is supposed to be generated. Finally, §4 outlines some broader concerns about the notion of normative defeat.

2 Epistemic ceiling principles

Crucial to Goldberg’s approach is the notion of “should have had” or “ought to have had” (which he discusses a bit more fully in Goldberg forthcomingb). A more precise characterization of the “ought” here will matter for some points I shall make later in §3; but for now I shall proceed

² In particular, can evidence, understood propositionally, be false (itself a contentious position)? If so, some cases of evidence one ought to have had will expect one to have updated on a falsehood. But for such cases it is unclear why we should accept that one loses justification just because one ought to have updated on a falsehood: after all, updating on falsehoods is generally epistemically bad, and one should be able to preserve one’s justification by *avoiding* adding a falsehood to one’s evidence. I shall discuss an additional worry about false evidence in §3.

as if we have an adequate grasp of this notion for evaluating some cases and some epistemic principles.

Goldberg (*forthcoming*) argues for two principles of evidence where “ E_{SSH} ” abbreviates “evidence S should have had”:

- (**) S’s (total) evidence E establishes a *justificatory ceiling* for the justification of S’s belief that p , and does so in a way that is independent of the epistemic significance of the evidence S should have had (E_{SSH}).

- (++) S’s belief that p can be *no more justified* than it would have been had she acquired E_{SSH} (and updated in a rational fashion³).

I would like to begin by noting several worries about (++) in particular.

If we think of justification in terms of evidential support, such that justification, at least in these cases, is conferred by evidence, then a plausible principle concerning evidential support links it to probability-raising. That is, some E evidentially supports a proposition p (for some S) just in case E raises the probability of p (for S); and E raises the probability of p (for S) just in case $\Pr(p | E) > \Pr(p)$. And $\Pr(p | E) > \Pr(p)$ will hold just in case $\Pr(E | p) > \Pr(E | \neg p)$. If we understand justification in these terms, a belief being more or less justified amounts to one’s evidential probability being (made) higher or lower, and talk of “justificatory ceilings” in the above principles can be construed as probability ceilings. For example, (**) may be reconstrued as the principle that one’s credence in p should not exceed the probability of p on one’s evidence, that is, $\Pr(p | E)$; exceeding it would be to believe (have a degree of belief) beyond that which one’s evidence supports.

On this set up, (++) says that, where E is S’s current evidence, S’s $\Pr(p | E)$ cannot (*should* not?) be greater than it would have been had she

³I shall henceforth suppress this clause about rational updating, though it is assumed in all of what follows.

acquired E_{SSH} and rationally updated on it. That is, $\Pr(p | E) \leq \Pr(p | E \& E_{SSH})$. But given this set up, there are cases where $(++)$ is false. Surely there can be evidence I should have had, but which I don't yet have, which would lower my probability once I acquire and update on it. Any case in which I know there to be additional evidence I ought to have concerning some p which I believe given my present E , but where my higher-order evidence is neutral on whether that additional evidence will confirm or disconfirm p , will be one where, by my current evidence, I might currently be more justified or have a higher probability given E : for once I gain the evidence I ought to have, my probability for p should go down should E_{SSH} disconfirm p .⁴ Similar points apply to a case in which I have no higher-order evidence at all, but where there nevertheless is some evidence I should have had.⁵

But $(++)$ says this cannot happen; rather, it says that one cannot have a probability higher than (have a belief more justified than) that which is supported by one's total evidence *plus* the evidence one should have had. $(++)$ claims that evidence one should have had can defeat your present justification even if you never acquire that evidence, should it be the kind of evidence which would lower that probability (justification). The strangeness of this is revealed by the idea that, on Goldberg's framework, one's justification or evidential probability presumably would not go down once one actually acquires the E_{SSH} : for it already went down just in virtue of there being negative E_{SSH} . But shouldn't one's justification go down once one acquires and updates on the negative evidence one should have had?

It is open to Goldberg, of course, to say that one's justification initially goes down only a *bit* simply due to there being some negative E_{SSH} ; and perhaps he can also say that it goes down even more once one actually acquires it. For example, perhaps the mere existence of negative E_{SSH} initially drops one's justification only *halfway* to the level at which it would be once one does acquire that evidence. But this move would require, I think,

⁴This is the case Goldberg labels: $(E_{HO\emptyset})/(E_{SSH-})$.

⁵What Goldberg labels: $(E_{HO\emptyset})/(E_{SSH-})$.

a more detailed story about this process than that given by the epistemic ceiling principles, for it would need to tell us how much justification goes down in each case. (For example, what would be the difference between the initial downgrade in justification in Goldberg’s unhappy/harmful case (E_{HO^-})/(E_{SSH^-}) and his null/harmful case ($E_{HO\emptyset}$)/(E_{SSH^-})? Does the former result in a larger initial downgrade than the latter?) On the other hand, Goldberg might propose that these difficulties are a reason to eschew talk of justification in terms of probability raising or lowering. It is enough for my present purposes to show that on its probability interpretation, a result of his principles seems counterintuitive.

Note that (**) and (++) *conflict* in certain cases. The reason is that (**) is about what one’s total *current* evidence, including one’s higher-order evidence, permits independently of the effect of E_{SSH} , if acquired, would have; whereas (++) is about what is permitted given one’s total current evidence as it *would be affected* by the E_{SSH} . Now compare two cases in which Sue believes p based on her total evidence E . In the first case, suppose part of Sue’s E includes the higher-order evidence that there is some additional evidence she ought to have which if had would *raise* her justification (probability) for p . Goldberg characterizes this as a “happy” case, which he symbolizes as (E_{HO+}); and let us say that in it, Sue is justified to degree X . (**) says that X is the justificatory ceiling, which is established independent of whether Sue’s E_{SSH} is in fact such that it would raise her justification (probability). Now in the second case, there is E_{SSH} which is such that, if Sue had it, should *lower* her probability for p . Even having said nothing in this second case about Sue’s higher-order evidence—it might be happy, unhappy, neutral, or there might be none—(++) says that Sue is justified to some degree *less than* X from the first case. On (++) in this second case, Sue may only be justified to the degree supported by E & E_{SSH} , for it says that she can be no more justified than she would have been had she both E & E_{SSH} .⁶ In other words, in a case of happy higher-order evidence but negative evidence one should’ve had, namely

⁶In our probability talk, Sue cannot have a higher probability than she would’ve had, had she acquired E_{SSH} and added it to E .

Goldberg's $(E_{HO+})/(E_{SSH-})$, (**) and (++) offer conflicting ceilings.

This is not in itself a serious difficulty, for Goldberg wants these two general principles (**) and (++) to combine to yield the following *Epistemic Ceiling Principle* (ECP):

ECP Given a subject S 's belief that p based at time t on evidence E , if there was evidence S should have had (E_{SSH}) at t , then (i) S 's belief that p at t can be no more justified than it is on E , and (ii) S 's belief that p can be no more justified at t than it would have been if S had acquired E_{SSH} at t (and updated in a rational fashion).

As such, ECP just says that the justificatory ceiling for S must be the lower (if any) of the two ceilings, issued by its clauses (i) and (ii).

Is this principle too strong? We've already considered above a case where the evidence one should have had would, if added to one's evidence, bring down one's justification (probability) for the p in question, but where plausibly one's current justification (probability) for p on one's total evidence seems to be unaffected by this: for example, where one lacks any higher-order evidence at all, or where that higher-order evidence is neutral, or even one in which my higher-order evidence is *happy*. It seems hard to sustain the idea that in these cases the subject loses justification; but Goldberg embraces⁷ the idea that even when one's higher-order evidence suggests that the E_{SSH} will, if acquired, only improve one's justification, the E_{SSH} still reduces one's justification for p if it in fact supports $\neg p$.

⁷"Now it seems to me that anyone who thinks that evidence you should have had can undermine or defeat your justification for believing that p , should regard this as such a case—and this, despite the fact that S had misleading higher-order evidence. If this verdict is correct, it gives us an interesting result: the evidence you should have had *trumps* the higher-order evidence you did have regarding the (existence and bearing of the) evidence you should have had" (p. 11).

Thus much appears to hinge on the mysterious notion of “should have had,” which can generate “normative defeat” (Goldberg *forthcominga*, 11). This “should have had” is a defeating mechanism on which merely being such that you *should have had* some evidence, where that evidence would negatively affect your current evidential support, is enough to defeat the positive status given your belief by that current evidential support. To that notion we now turn.

3 Expectations and “ought to have”

Goldberg thinks the obligation represented by the “ought” of evidence one *ought to have had*, or the “should” of what one *should have known*, is grounded in the reasonable expectations of others where those expectations are made reasonable by the norms internal to a social practice. Where someone reasonably expects a subject to know or have some evidence, and that expectation is grounded in participation in a practice which grants one the entitlement so to expect,⁸ this generates the “should have (had)” status concerning that subject, which is enough to ground that obligation. As Goldberg puts it,

Because our epistemic reliance on others is so profound and systematic, and because this reliance is largely rationalized by the expectations we have (and which we are entitled to have) of other epistemic agents playing various roles in the knowledge community, other subjects are properly held accountable to these expectations when they are legitimate... this “holding accountable” is a species of practical/moral accountability: those who violate legitimate expectations are *prima facie* socially or morally responsible (and are socially or morally blameworthy

⁸See also Goldberg *forthcomingb*, §2.3: “what entitles us to these expectations is nothing deeper than the practical rationality of our information-exchanging engagements with one another... such engagements are practically rational only if these expectations are legitimate; but these engagements *are* practically rational; so the expectations are legitimate.”

unless excused). However, this social or moral failure opens up subjects to epistemic criticism as well. At a minimum, a subject who fails to have all of the evidence she should have had exposes her interlocutors to an additional risk of falsity and unreliability in the information she passes along. (forthcominga, 19)

But there are at least three difficulties with this approach.

The first is that sometimes our practices justify “should have known p ” judgments which are compatible with knowledge of some q which itself depends inferentially upon p which is not known. Cases of inferential knowledge from falsehood are paradigms of such judgments. Sometimes one can believe a falsehood p , yet competently deduce that q from p and thereby come to know that q .⁹ For example:

A teacher counts the students in the class and arrives at a count of 53. The teacher believes that the result of the count is correct. The teacher then infers (and thereby comes to believe) that there will be enough handouts, having made 100 of them. In fact there are 54 students. (Warfield 2005, 407–408)

An employee calls the employer and says “I am sick today.” The employer believes the testimony and infers that the employee won’t be in at work that day. What the employee says is false (she just didn’t fancy coming in that day), but the employer’s inference is to a truth: she does not come in that day. (Cf. Hawthorne and Rabinowitz forthcoming, §2)

Notice, however, that our practice of deduction is one where many, including the inferrer herself, reasonably expect that the one making the inference “should have known” the premise from which it was deduced. And the social expectations associated with the practice of deduction are

⁹See Warfield 2005, Fitelson 2010, and Hawthorne and Rabinowitz forthcoming, among others.

such that others who rely on knowers making such inferences will be entitled to expect those inferring to have known the premises of their deductions. But Goldberg's account would rule out such cases of knowledge from false belief as ones where the subject's belief is less justified (or even unjustified) due to the "should have had" obligation; it is to that extent counterintuitive.

A second problem is that there are cases where another's reasonable expectation of what evidence one ought to have had, or what one should have known, is based on an epistemic responsibility incurred by an *action* the person undertook, even though there may be no relevant evidence or knowledge which the person could have had. For example, many philosophers have argued persuasively that knowledge is the norm of assertion, such that one who flat-out asserts a proposition ought to know what they assert. Thus assertion is a social linguistic practice, governed by the following norm:

(KNA) One must: assert that p only if one knows that p .¹⁰

KNA well-explains a wide range of linguistic data, including the fact that we typically expect asserters to know what they assert. If they do not know what they assert, they are typically criticizable; Austin's datum was: " 'But I took it you *knew*' , said reproachfully" (Austin 1946; repr. 1961, 45). Such a reproach might have instead been "But you should have *known*," because a speaker who opted to assert should have known what she asserted. Similarly, the criticism might be put that she ought to have had (better) evidence for what she asserts. (Thus this point doesn't rest on the KNA.) But it doesn't follow from these correctness of these judgments that the speaker could have known, or that there is some evidence available which she ought to have had. For if the proposition asserted is false, it couldn't

¹⁰E.g. Williamson 2000, Ch. 11, among others; see Benton 2014, §1 for an overview. Goldberg himself sees the attraction of KNA, but opts for a different norm (Goldberg 2015) which itself appeals to the mutual expectations of speakers; see Benton forthcoming for some criticisms.

have been known; and under the right kind of circumstances, there may have been no such evidence for it. Nevertheless, given that the speaker did assert it, the “should have” judgments and their correlated reasonable expectations are correct.

A third problem with Goldberg’s approach is that sometimes our practices make reasonable certain expectations which nevertheless are not grounded in obligations of any kind, but merely in the frequency with which someone achieves their aims. Professor X might be well-known as one who has always submitted her high-quality referee reports within a week of accepting the journal’s request. Her policy is only to accept such requests when she is able to referee the paper within a week. This generates a reasonable expectation from all those who know this (perhaps only herself) that she will referee the next paper within a week. In the unfortunate case where she doesn’t meet this expectation, she has not fulfilled her *aim*; but it does not follow that she has violated any *norm*.¹¹ There is plausibly no “should have done” applicable to her in that case (at least in the robust sense of “should have” to which Goldberg appeals). But if so, this calls into question the foundation on which Goldberg has built his account of “should have had” obligation. Now Goldberg (in his *forthcomingb*) is careful to claim that the kinds of expectations he is invoking are normative ones (as in the assertion case of the previous paragraph), and not epistemic ones of mere expectation; but the question then becomes what exactly it is that grounds these expectations as *normative* rather than merely epistemic. What does it take for a social practice to generate specifically normative expectations?

In connection with this concern, and the question (raised in fn. 2) about whether Goldberg would allow that evidence can be false, consider the interesting case of a spy, Rose, who is a double-agent. If Goldberg allows that evidence can be false, the reasonable expectations of Rose’s

¹¹This point goes overlooked in Goldberg’s discussion (see esp. Goldberg *forthcomingb*, §1.2, including his principle PGE): a practice can have standards, in terms of *aims*, without thereby being norm- or rule-governed. What is needed is an argument for the claim that the practices in question have standards which are set by rules and not just by aims.

colleagues in each of the spy agencies for which she is working will presumably cut against each other in such a way as to rob her of a lot of justified beliefs. Each side may expect Rose to have some (false) evidence, because each side is being misled concerning some truths. Indeed, Rose may be the very one who is providing to each side such misleading evidence (or the misinformation that such evidence exists). In this situation, each side for whom Rose is spying will expect her to have some evidence which she actually knows to be false. Now we might plausibly think it is a constraint on evidence possession that the subject at least believe a proposition in order for it to become part of her evidence; and so Rose, if she knows some propositions to be false, could not have them as part of her evidence.¹² Nevertheless Goldberg's account, if it allows for false evidence, would also allow for false evidence one should have had; and it would thus suggest that Rose is not (as) justified in some of her beliefs simply because there are reasonable expectations which the two sides are entitled to hold about evidence which Rose should have.

To recap the broad lessons of this section, the three difficulties I've raised for Goldberg's proposal are these. First, there is inferential knowledge from falsehood, which seems impossible on the account of normative defeat given by Goldberg. Second, there are correct judgments of "S should have had *E*" where there was in fact no *E* for S to have had. And third, there are cases of reasonable expectations by others derived from a social practice which do not of themselves generate a *norm* of the kind claimed to be supporting the "should have had" deontic claim, and part of the challenge is to delineate those social practices which generate not merely expectations but a norm (and in virtue of what they do so).

¹²This assumption would handle the difficulty, if evidence can be false, concerning whether one who offers to others misleading evidence which she herself knows to be false already "has" those falsehoods as evidence. On this assumption, a knowing liar does not "have" as evidence the false propositions she tells to others. (But on Goldberg's approach, if some of those falsehoods count as evidence the liar ought to have had, it would seem that by lying the speaker thereby fails to know that her lie is false.)

4 Knowledge and normative defeat

There is a larger worry about a project like Goldberg's, which may well extend to any systematic formulation of normative defeat. The worry might be deemed one of circularity insofar as normative defeat is supposed to encroach on justification, which is in turn typically thought to be necessary for knowledge; so on this approach to the project, spelling out the conditions on normative defeat is part of explaining the conditions for knowing. But this is problematic in that whether some information really is defeating evidence for one's belief that p can itself depend on whether one knows that p .

Sometimes one can know a proposition and rationally reject as misleading information which would otherwise be evidence against what one knows (or against that one indeed knows it). And sometimes one can do this even if there is a clear sense in which it is evidence that one should have had. If (case *A*) I know that I am with Blake at the pub, and I get a text on my phone from a friend indicating that Blake is at the park, the text will not be evidence for me that Blake is at the park because I know he's not at the park, but at the pub with me; indeed, if anything, what I have acquired is evidence that my friend is mistaken about Blake's whereabouts. Yet if (case *B*) I had not known Blake's whereabouts, that text would constitute evidence for me that he's at the park; had I merely a note from Blake that he was headed to the pub, that text might defeat my justification for believing that Blake is now at the pub.

Now suppose (in either case *A* or *B*) that I know I have received a text, which I also know concerns Blake's whereabouts, but I haven't yet looked at it. There might be a clear sense in which it is evidence I should have. In case *B*, we can perhaps grant to Goldberg that this evidence I should have had plausibly affects my justification for believing that Blake is at the pub (at least if I wait for too long before reading the text). But in case *A* this evidence I should have doesn't, contrary to Goldberg's ECP and (++), reduce my justification (assuming that such justification is necessary for knowledge): I still know that Blake is at the pub with me.

This shows there to be serious problems with ECP and (++)), but it also shows that the prospects are dim for articulating a non-circular account of normative defeat which does not itself advert to whether one knows. For if normative defeaters (as all such defeaters) must be absent in order for one to know, but the evidence which is a normative defeater in one kind of case does not defeat in a similar case where the subject knows, one may suspect that such defeaters are inadequate to explain judgements of non-knowledge by being responsible for *why* the subject does not know. In other words, defeaters are supposed to explain some cases where you fail to know, rather than some cases where you know explaining why putative defeaters are not operable. The above kind of case, however, provides reason for pessimism that any such account of defeat can be given in other terms than whether one knows.¹³

5 Conclusion

I have argued here that Goldberg’s admirable attempt to give a full account of normative defeat in terms of evidence one should have had suffers from a variety of problems. Even without clarification concerning the exact notion of “evidence” in play, and without an intuitive way to model exactly how evidence generates or defeats justification, we can see that Goldberg’s appeal to evidence or knowledge you “should have had” generates overly strong epistemic ceiling principles. His account of what generates the deontic “should” or “ought,” though based on important insights in social epistemology, does not seem up to the task of capturing all and only the cases of improper belief or reduced justification. This is complicated by the fact that such epistemic obligations are plausibly generated even in cases where one one continues to know, or where one could not have the evidence one “ought” to have had. Finally, it is unclear to me that one can characterize the desired notion of normative defeat in

¹³Indeed, I am pessimistic that defeat is a workable notion in epistemology at all. See Lasonen-Aarnio 2010, 2014, Hawthorne and Srinivasan 2013, and Baker-Hytech and Benton 2015 for such worries.

a way that is separable from, and explanatory of, whether one knows.¹⁴

References

- Austin, J. L. 1946. "Other Minds." *Proceedings of the American Philosophical Society, Supplementary Volumes* Reprinted in Austin (1961).
- Austin, J. L. 1961. *Philosophical Papers*. Oxford: Clarendon Press.
- Baker-Hytch, Max and Benton, Matthew A. 2015. "Defeatism Defeated." *Philosophical Perspectives* 29: 40–66.
- Ballantyne, Nathan. 2015. "The Significance of Unpossessed Evidence." *Philosophical Quarterly* 65: 315–335.
- Benton, Matthew A. 2014. "Knowledge Norms." *Internet Encyclopedia of Philosophy*, ISSN 2161-0002, <http://www.iep.utm.edu/kn-norms/>.
- Benton, Matthew A. forthcoming. "Gricean Quality." *Noûs* .
- Fitelson, Branden. 2010. "Strengthening the Case for Knowledge from Falseness." *Analysis* 70: 666–669.
- Goldberg, Sanford C. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.
- Goldberg, Sanford C. forthcominga. "On the Epistemic Significance of Evidence You Should Have Had." *Episteme* .
- Goldberg, Sanford C. forthcomingb. "Should Have Known." *Synthese* .

¹⁴Thanks to Max Baker-Hytch, Sandy Goldberg, and Dani Rabinowitz for helpful thoughts and discussion. This publication was made possible by support from a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the John Templeton Foundation.

- Hawthorne, John and Rabinowitz, Dani. forthcoming. "Knowledge and False Belief." In Rodrigo Borges, Claudio de Almeida, and Peter Klein (eds.), *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford: Oxford University Press.
- Hawthorne, John and Srinivasan, Amia. 2013. "Disagreement Without Transparency: Some Bleak Thoughts." In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays*, 9–30. Oxford: Oxford University Press.
- Kornblith, Hilary. 1983. "Justified Belief and Epistemically Responsible Action." *Philosophical Review* 92: 33–48.
- Lackey, Jennifer. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lasonen-Aarnio, Maria. 2010. "Unreasonable Knowledge." *Philosophical Perspectives* 24: 1–21.
- Lasonen-Aarnio, Maria. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88: 314–345.
- Warfield, Ted A. 2005. "Knowledge from Falsehood." *Philosophical Perspectives* 19: 405–416.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.