

On Social Defeat

B.J.C. Madison

(Forthcoming in *Canadian Journal of Philosophy*)

Draft Version – Do Not Cite Without Approval

Keywords: Epistemology; Social Epistemology; Defeaters.

Abstract: Influential cases have been provided that seem to suggest that one can fail to have knowledge because of the social environment. If not a distinct kind of social defeater, is there a uniquely social phenomenon that defeats knowledge? My aim in this paper is to explore these questions. I shall argue that despite initial appearances to the contrary, we have no reason to accept a special class of social defeater, nor any essentially social defeat phenomenon. We can explain putative cases of social defeat with our existing epistemological apparatus.

1. *Introduction: Kinds of Defeaters*

Epistemic defeat is a familiar feature of our cognitive lives. Just as acquiring new information can lead to our having better justified beliefs, new information can also take that justification away. For example, suppose that you saw your colleague at work an hour ago, and they tell you that they will be in the office for the rest of the day. On that basis you believe what they say. Such a belief is epistemically justified. But now suppose that a trusted, reliable authority tells you that they just saw that colleague leave the office and board a bus. You now believe the trusted, reliable authority. By believing their testimony, many want to say that you have thereby acquired a *defeater* for your belief that your colleague will be in the office all day, meaning that the belief is no longer justified (or at least it is now less justified than it was before).

It has become common to distinguish between different kinds of defeaters, each individuated in terms of how they defeat: in terms of being believed, in terms of being what the subject *should* believe, and in terms of being true. This corresponds to the traditional three-way distinction between mental state, normative, and factual defeaters. Mental state defeaters are understood to be a subject's mental states, and in Jennifer Lackey's words, are "a doubt or belief that is had by S and that indicates that S's belief that *p* is either false or unreliably formed or sustained." (Lackey, 2008, p. 44) These defeaters function in virtue of being *had* by the subject, and generally are taken to function regardless of their truth-value or epistemic status.¹ A reliable and trusted informant telling you that you are experiencing an optical illusion is an example of such a defeater. If you believe what the informant says, this belief can defeat the justification your perceptual belief might have otherwise enjoyed. The case above about believing that a colleague will be in the office all day, which is then contradicted by the testimony of someone else you believe, is another example of acquiring a mental state defeater.

Sometimes justification is defeated not because of a belief that someone actually holds, but because of a belief that they *should* hold. This is a normative defeater, characterized as "a doubt or belief that S *ought to have* that indicates that S's belief that *p* is either false or unreliably formed or sustained." (Lackey, 2008, p. 44) Like mental state defeaters, these also function regardless of their truth-value or their epistemic status. For instance, suppose that due to wishful thinking and self-deception, one cannot bring oneself to accept overwhelming evidence one has that

¹ Some allow that false justified beliefs can defeat, but they doubt whether *unjustified* beliefs can play a defeating role. For example, to use James Pryor's terminology, unjustified beliefs may rationally "oppose" and / or "obstruct", but they do not defeat. (Pryor 2004; 2018) Such beliefs are taken to be highly epistemically relevant, but not themselves sources of justification, whether positive or negative.

counts against what one believes. This evidence might nonetheless constitute a normative defeater that rebuts or undermines one's justification.

Finally, a factual defeater is commonly understood thus: "a true proposition such that if it were added to S's belief system, then S would no longer be justified in believing that *p*." (Lackey, 2008, p. 45-46, footnote 21) Unlike mental state and normative defeaters, this kind of defeater functions in virtue of being true. Factual defeaters were first introduced as part of an early diagnosis and response to the Gettier problem (e.g. Lehrer and Paxson, 1969).² Here is an updated version of a classic case that illustrates factual defeaters:

LAMB: One of Dr. Lamb's students, Linus, tells her that he owns a Lamborghini. Linus has the title in hand. Dr. Lamb saw Linus arrive on campus in the Lamborghini each day this week. Linus even gave Dr. Lamb the keys and let her take it for a drive. Dr. Lamb believes that Linus owns a Lamborghini, and as a result concludes, 'At least one of my students owns a Lamborghini'. As it turns out, Linus doesn't own a Lamborghini. He is borrowing it from his cousin, who happens to have the same name and [birthdate]. Dr. Lamb has no evidence of this deception, though. And yet it's still true that at least one of her students owns a Lamborghini: a modest young woman who sits in the back row owns one. She doesn't like to boast, though, so she doesn't call attention to the fact that she owns a Lamborghini. (de Almeida & Fett, 2016, pp. 153-154)

Notice that '*Linus does not own a Lamborghini*' is a true proposition such that if Dr. Lamb were to believe it, she would no longer be justified in believing as she does. Crucially, advocates of factual defeaters hold that the reason Dr. Lamb does not know that Linus owns the Lamborghini is *because* her belief's justification is defeated by this true proposition. As such, they contend that this true proposition constitutes a factual defeater which explains why knowledge is absent.

In addition to three traditional kinds of defeaters, are there also social defeaters? Or if not a distinct kind of defeater, is there a kind of uniquely social

² So 'factual defeater' is a technical term that was introduced within the Defeasibility Theory of Knowledge, e.g. in Lehrer and Paxson, 1969. One of the earliest expressions of defeasibility can be found in Chisholm (1964).

defeat phenomenon that falls under one of the three general kinds distinguished above? Some have come to think so, convinced by cases given by Gilbert Harman which purport to show that one can have a justified true belief that fails to amount to knowledge because of the social environment. Here is one such case to illustrate:

ASSASSINATION Version 1: Suppose that Tom enters a room in which many people are talking excitedly although he cannot understand what they are saying. He sees a copy of the morning paper on the table. The headline and main story reveal that a famous civil-rights leader has been assassinated. On reading the story he comes to believe it; it is true...[but] suppose that the assassination has been denied, even by eyewitnesses, the point of denying being to avoid a racial explosion. The assassinated leader is reported in good health; the bullets are said, falsely, to have missed him and hit someone else. The denials occurred too late to prevent the original and true story from appearing in the paper that Tom has seen; but everyone in the room has heard the denials. None of them know what to believe. They all have information that Tom lacks. Would we judge Tom to be the only one who knows that the assassination has actually happened? Could we say that he knows this because he does not yet have the information everyone else has? I do not think so. (Harman, 1968, p. 172)

From this case Harman concludes that: “This reveals the social aspect of knowledge. The evidence that a person has is not always all the evidence relevant to whether he knows. Someone else’s information may also be relevant.” (Ibid., pp. 172-173)

My aim in this paper is to explore the questions of whether there are social defeaters, or else a kind of uniquely social defeat phenomenon that falls under one of the three general kinds of defeaters distinguished above. I shall argue that despite initial appearances to the contrary, we have no reason to accept a special class of social defeater, nor any essentially social defeat phenomenon. We can explain putative cases of social defeat with our existing epistemological apparatus, including appealing only to the three kinds of defeaters commonly postulated. So while justification and knowledge undoubtedly have social dimensions, I shall argue that we have yet to see that there is an inherently social form of defeat.

2. What Would it Take for a Defeater to be Social?

It is not immediately obvious what would make something a social defeater. But we can begin by ruling out the idea that if a defeater has a social source, that would be enough to make it a social defeater in any interesting sense. For instance, suppose I believe that p , but you then tell me something I know to be inconsistent with p , and I believe what you say. Plausibly I now have a defeater for my belief that p , and the source of this defeater, your testimony, is a social one. Nevertheless, this is insufficient to make this a social defeater. Rather, it is just an instance of a regular mental state defeater.

What about Gilbert Harman's widely discussed assassination cases? Might these help us identify what it would take for a defeater to be social? After all, these have been taken to be among the paradigm examples used to motivate the position that knowledge and justification are vulnerable to a kind of inherently social defeat.³ Here is a variation of the case:

ASSASSINATION Version 2: A political leader is assassinated. His associates, fearing a coup, decide to pretend that the bullet hit someone else. On nationwide television they announce that an assassination attempt has failed to kill the leader but has killed a secret service man by mistake. However, before the announcement is made, an enterprising reporter on the scene telephones the real story to his newspaper, which has included the story in its final edition. Jill buys a copy of that paper and reads the story of the assassination. What she reads is true and so are her assumptions about how the story came to be in the paper. The reporter, whose by-line appears, saw the assassination and dictated his report, which is now printed just as he dictated it. Jill has a justified true belief and, it would seem, all her intermediate conclusions are true. But she does not know that the political leader has been assassinated. (Harman, 1973, pp. 143-144)

Like in the first version of the assassination case, Harman suggests that this case also reveals a social aspect of knowledge. Other philosophers have followed suit

³ Other relevant work here includes: Pollock (1986); Meeker (2004); Goldberg (2017), (2018).

in holding that cases like these reveal that social factors can destroy knowledge, although their analyses of how this happens differ (e.g. Pollock, 1986, pp. 190-193; Meeker, 2004). I shall argue that there is nothing inherently social in this kind of case that requires special epistemological treatment. I propose a disjunctive analysis of the assassination cases: either they are cases of knowledge after all, and the luck present is a benign form of *evidential* luck, or else if knowledge is lacking, then this can also be explained by appealing to the notion of epistemic luck, without making essential use of anything social. I shall explain each of these options in turn.

As to the first proposal, I grant that one might be initially tempted to deny that Jill in Version 2 of the case, for example, knows that the political leader has been assassinated (I suggest that the same disjunctive analysis can be applied, *mutatis mutandis*, to ASSASSINATION Version 1). After all, it is (in some sense) a matter of luck that she has the justified true belief that she does, and knowledge is essentially non-lucky (that is surely a principle lesson of the Gettier problem). Jill could very easily have seen the misleading TV news broadcast, and in that case, she would believe that the leader survived the assassination attempt. So her justified true belief that the political leader was assassinated is too lucky to amount to knowledge, one might think.

But upon further reflection, it is clear that there are instances of lucky knowledge. We should not conclude from the Gettier problem that knowledge excludes any and all kinds of luck. Rather, knowledge is compatible with certain kinds of luck, and not others. For example, Duncan Pritchard (2005), following Peter Unger (1968), has argued that knowledge is incompatible with specifically *veritic* luck, whereas what has been called *evidential* luck, by contrast, is no threat to knowledge at all. Very roughly, a belief is veritically lucky if, given how it was

formed, the belief could have very easily been false. A belief is evidentially lucky, on the other hand, if it is a matter of luck that one has the evidence that one does (i.e. while one has the evidence that one does in the actual world, there are nearby worlds where one lacks that evidence).

Applying this distinction between different types of epistemic luck to Harman's ASSASSINATION Version 2 case, for example, I suggest that we should hold that the kind of luck Jill's belief enjoys is epistemically benign evidential luck. She is of course very lucky to have the evidence that she does, namely, the reliably produced accurate newspaper report. Many other people lack that evidence, and we can suppose that she herself might easily not have had it – she could easily have seen the TV news broadcast instead. But *given* the evidence she actually has, it is no accident at all that her belief is both justified and true. In short, I suggest that contrary to Harman's original judgment, and contrary to possible initial appearances otherwise, Jill *does* know that the political leader has been assassinated. She does not lack knowledge, although obviously she almost does. Jill knows luckily, but not problematically so.⁴

Suppose that despite the above considerations, one agrees with Harman that Jill lacks knowledge. Does that motivate the hypothesis that knowledge and justification are inherently social, as Harman contends? I suggest that it does not. Appealing to the same distinction drawn above between veritic and evidential luck, we still have the resources to explain the judgment that Jill lacks knowledge, without having to appeal to social considerations. It is common among those who agree that

⁴ Stephen Hetherington has warned us against confusing lacking knowledge with *almost* lacking it. He has called a bad inference from a counterfactual lack of knowledge to an actual lack of knowledge the "Epistemic Counterfactuals Fallacy" (Hetherington, 1998, pp. 456-9). Given a sort of fallibilism that many epistemologists accept, certain kinds of lucky knowledge are possible. For critical discussion of this so-called Knowing Luckily Proposal as a fully general solution to the Gettier problem, see Madison (2011).

knowledge excludes too much veritic luck to formulate anti-luck conditions that are designed to rule out such luck. Take the safety condition, which has become one of the more popular accounts of the anti-luck condition (e.g. Pritchard 2005).⁵

Safety states that, roughly, it is a necessary condition of S knowing that *p* that, given how the belief was formed, it could not easily have been false. Crucially, safety must be relativized to methods or bases. For instance, safety requires that if the subject were to form their belief in nearby worlds on the same basis as they do in the actual world, the belief would still be true.⁶

Applying this to Version 2 of Harman's assassination case, it is important to consider on what basis Jill forms her belief that the president was assassinated. If the method is construed narrowly as the belief being held on the basis of the newspaper report, as I suggested above, then the belief is safe. If, on the other hand, one takes the method she is using as relying upon trustworthy news sources generally⁷, then her belief is unsafe. This is because that method is unreliable given the nearby worlds where she forms her false belief about the president's fate on the basis of the false TV news reports.

Notice that if this is the correct way of specifying the belief-forming method Jill is relying upon, as some will maintain, her belief that the president was

⁵ A rival, albeit currently less popular, way of attempting to specify the sense in which knowledge excludes luck is in terms of *Sensitivity*. On Nozick's (1981) pioneering version of Sensitivity, it is made up of two distinct conditions, namely, *Variation*: if *p* were false, S would not believe that *p*; and *Adherence*: were *p* true, S would believe that *p*. See Becker (2018) for recent critical discussion of the Sensitivity response to the Gettier Problem. In particular, Becker argues that Nozick wields *Adherence* in inconsistent ways depending upon the point at issue. For recent critical discussion of the Sensitivity condition in epistemology generally, see Becker and Black (2012).

Crucially for our purposes, none of these principles rely upon or involve anything *social*.

⁶ And what is it to form a belief in the *same way* in other worlds? Questions about generality and the individuation of belief forming methods and bases immediately arise. This issue is familiar from discussions of the so-called "generality problem" for reliabilist theories of epistemic justification. See Conee and Feldman (1998) for an early and influential presentation of this problem.

⁷ Sandy Goldberg suggested to me that this is the plausible way of specifying the belief forming method (personal correspondence).

assassinated is infected by too much veritic luck, and so plausibly does not count as knowledge. But there is nothing inherently social here. There is simply too much veritic luck in the environment such that Jill's belief is not safe, and so it does not amount to knowledge.

The upshot is this: we have the resources to explain why Harman's assassination cases are ones of knowledge, but even if they are ones where knowledge is lacking, it is not genuine social defeat. The social sources are causally relevant in explaining why, for example, Jill forms the beliefs she does, but the social dimension is not itself epistemically salient.

In looking for a form of specifically social defeat, I suggest that we follow Sandy Goldberg's recent proposal that social defeat, if there is such a thing, is a kind of normative defeat. Goldberg understands the notion most generally as follows: "doctrine of normative defeat: if S should have known that p , then the proposition that p is a defeater for S." (Goldberg, 2017, p. 2893) While I agree that this is a central kind of normative defeat, this construal is too narrow, as there are surely kinds of normative defeat that need not involve knowledge (in the sense that the subject should-have-known that p). For instance, if you have *justification* to believe that p , but you believe that not- p , then p is a normative defeater for you (and presumably the same is true for other epistemic statuses).⁸ So generalizing somewhat, if S (epistemically) ought to believe that p , then p is a normative defeater for S.

So how, or under what conditions, might there be a kind of normative defeat that is inherently social? After all, clearly not all normative defeat is social defeat, for

⁸ Nathan Ballantyne has argued that awareness of unpossessed evidence sometimes undermines rational belief. For his discussion of the epistemic significance of unpossessed evidence, see his (2015).

example, in cases involving only oneself. For instance, to return to an earlier example, if one cannot bring oneself to accept overwhelming evidence one has that counts against what one believes, this evidence will constitute a normative defeater that rebuts or undermines one's justification. No other agents need be involved. Goldberg's suggestion is that social defeat is a species of normative defeat because it arises from an essentially social should-have-known phenomenon that is grounded in *social expectations*. It is to that proposal I now turn.

3. Goldberg's Analysis of Social Defeat: Should-Have-Known

Goldberg is not primarily interested in cases where someone should have known that *p*, given the evidence they have. That would be a straightforward case of either mental state or normative defeat, and need not in any way be social. Rather, Goldberg's primary focus is, very provocatively, on when one should have known, even when the subject is *not* in a position to know. (Goldberg, 2017, p. 2864) At first that claim might have a paradoxical ring: how could it be true that someone ought to know something, even if there is no way that they could know it? Does 'ought' not imply 'can', after all?

To begin to motivate the most controversial and interesting versions of this claim, Goldberg begins with the least controversial examples. (Goldberg, 2017, pp. 2865-2867). In these examples, the subjects have professional responsibilities and duties to know certain things. If they fail to have the relevant knowledge, it is both true that they should have known, given their roles, and we are entitled to expect that they should have known. For example a daycare worker, given their job, ought to know where the children under their care are. If they do not know where one of the

children they are responsible for is, it is both true that they should have known, and that we are entitled to expect that they have this knowledge. And this is true even if the daycare worker is in no position to know where the child is (for instance, the child might have become lost on an excursion, and given that, the daycare worker now has no way of knowing where they are).

Goldberg also offers Hilary Kornblith's case of epistemically culpable ignorance to motivate what he claims is the social dimension of when subjects should have known:

Jones is a headstrong young physicist, eager to hear the praise of his colleagues. After Jones reads a paper, a senior colleague presents an objection. Expecting praise and unable to tolerate criticism, Jones pays no attention to the objection; while the criticism is devastating, it fails to make any impact on Jones' beliefs because Jones has not even heard it. Jones' conduct is epistemically irresponsible; had Jones' actions been guided by a desire to have true beliefs, he would have listened carefully to the objection. Since his continuing to believe the doctrines presented in his paper is due, in part, to this epistemically irresponsible act, his continued belief is unjustified. (Kornblith, 1983, p. 36)

Here again, Jones should have known that his paper faces damning criticism, and this Goldberg contends is because it is part of the physicist's job to weigh all the relevant evidence for and against hypotheses – they have a professional responsibility to know if the hypotheses they accept are true or well-supported. Goldberg's central contention is that it is in virtue of legitimate social practices that we are entitled to certain epistemic expectations, and that these legitimate expectations ground the should-have-known phenomenon.

From examples of professional and other practice-generated entitlements, Goldberg generalizes as follows:

SHK: *S should have known that p* when (one) there is a subject S* who is entitled to expect that S satisfied a certain epistemic condition, (two) the satisfaction of this expectation requires that S know that *p*, and (three) it is not the case that S knows that *p*. (2017, pp. 2868; 2882)

Goldberg offers a key illustration of this principle, one that is supposed to be a case of should-have-known that is rooted inherently only in social practices, and not in professional responsibilities or other institutions. Call this the case of Rodriguez /

Yang:

Consider the case of Rodriguez, who out of inattentiveness fails to read all of the items on the single-page grocery list he was given by his partner Yang. Here, Yang's expectation that Rodriguez is a competent reader (being able to read and understand all that is on his list) is an expectation of the relevant sort: Yang expects mature, cognitively healthy literate subjects to read with competence (satisfying SHK's condition one), the satisfaction of this expectation in this circumstance requires (as per condition two) that Rodriguez know that they need milk, eggs, sugar, ... (where these are all things on the list), and yet (condition three) Rodriguez fails to know. Hence Yang is warranted in saying that Rodriguez should have known. (Goldberg, 2017, p. 2883)

Let us grant for the sake of argument that our basic epistemic expectations capture our basic epistemic standards. Grant also that each of us is entitled to have these expectations when we engage with each other. Even grant that these expectations warrant should-have-known allegations. Does all that entail that it is *true* that the subject should have known when Goldberg's SHK is met? If so, is that *because* it was socially expected of them? I shall argue for negative answers to both questions.

I agree that Rodriguez should have known what was on the grocery list, but I submit that this has nothing to do with social expectations. It matters *why* he fails to know, and this I suggest, is because of his inattentiveness. A grocery list is normally written to remind one to buy all of the items listed on it. The goal in writing the list is to end up with (at least) everything on it. If one reads only part of the list, given one's

aim, one is very unlikely to meet this goal (but it is not impossible to meet this goal of course: one could randomly buy what just so happened to be on the list; or someone might add items to your shopping cart that unbeknownst to you were on the list, and so forth). Given the point of writing grocery lists, Rodriguez's inattentiveness explains why he should have known what was on it.

To be clear, I am not contending that his inattentiveness is a *defeater*.⁹ It is not that Rodriguez has a belief that Rodriguez and Yang need milk, eggs, and sugar, but that this belief is unjustified, or fails to amount to knowledge due to inattentiveness. Rather, his inattentiveness results in him not forming the target belief at all. This is a case of ignorance, and one where Rodriguez should have known, but it is not one of defeat (normative or otherwise).

To see even more clearly why there is no relevant social dimension in play here, notice that if Rodriguez wrote *himself* a grocery list, the above points equally apply, which suggests it is not the social aspect of the Rodriguez / Yang case that makes a difference. If due to inattentiveness Rodriguez failed to read all the items on the grocery list he wrote himself, it is still true that he should have known all the items on it. Here again, the inattentiveness is not a defeater of justification or knowledge, but it explains his ignorance. Given Rodriguez's goals and purposes in writing himself a grocery list, he should have known everything on it. In short, the Rodriguez / Yang case does not motivate the position that there is a variety of should-have-known, or normative defeat, that is essentially grounded in social practices.

Aside from the question of whether Goldberg's central case motivates the role of social considerations in normative defeat, it is also important to note that there

⁹ Thanks to Sandy Goldberg for impressing upon me the need to clarify the connection between inattentiveness and how this might impact upon when a subject should have known.

seems to be a gap in his overall approach between a principle like SHK and the sort of question we are meant to be asking. What we want to know is when is it *true* that someone should have known that p , so that if they fail to, p amounts to a normative defeater for them? But SHK is couched in terms of *entitlement to expectations* that someone knows. The conclusion of the Yang / Rodriguez argument is also explicitly framed in terms of what Yang is *warranted in saying* about whether Rodriguez should have known.

In the absence of further argument, it seems that when it is true that someone should have known can come apart from whether others are entitled to expect someone to know. Perhaps one is entitled to expect someone to know when one has justification to believe that they know. But evidence, and appearances, can be misleading. It might be perfectly reasonable to expect someone to know something, without this entailing that they ought to know it.

By way of illustration, suppose that you and I are both new in town, but I do not know this about you. Instead, I am told by someone I know to be a local that you were born and raised here, and I have no reason to distrust them. On this basis I might then be perfectly well entitled to expect that you know lots of things about the city: its geography, its best restaurants, and the like. Despite this entitled expectation, surely it does not follow that it is *true* that you ought to know these things. Like me, you too are new in town; since it is not true that I ought to know these things, given that I just arrived, neither should you.

Likewise, granting that Yang might be warranted in saying that Rodriguez should have known is not equivalent to saying that it is true that Rodriguez should have known. Here too one's evidence might overwhelmingly warrant one in

believing that the subject should have known, without it being true that they should have known. Again, it is possible that the evidence is undetectably misleading or incomplete in various ways. In the above example, I might be warranted in saying that you should know the city's landmarks, given that I justifiedly believe that you are a local, without it being true that you ought to have this knowledge.

In response, one might claim that there is no room for the distinction I want to draw between when it is true that someone should have known and when others are entitled to expect them to have known, since being entitled to expect S to X entails that S ought to X.¹⁰ Whether this objection is sound will of course turn on what it is to be entitled to believe or expect something.

Goldberg says that the practice-generated entitlements he is interested in must be distinguished from epistemic entitlements (Goldberg, 2017, p. 2867; Goldberg, 2018, p. 166). And that is all for the good, because on all dominant accounts of epistemic entitlement (e.g. Dretske 2000; Burge 2003; Wright 2004; Peacocke 2004), it is not factive: being entitled to hold a belief (and presumably by extension an expectation) does not entail that the belief is true, and so presumably it would not follow from these accounts of entitlement alone that being entitled to expect S to X entails that S ought to X.¹¹ Entitlements in this regard are more like justification and less like knowledge (while entitlement is nevertheless a notion distinct from justification, according to Dretske, Burge, and Wright; Peacocke, on the other hand, seems to use “entitlement” and “justification” interchangeably).

¹⁰ In fact, Sandy Goldberg himself suggested this response in personal correspondence.

¹¹ See Altschul (n.d.) for a survey of these four accounts of entitlement. See also Graham and Pedersen for relevant discussion (forthcoming).

So if this objection stands, another notion of entitlement is needed. Goldberg does offer a proposal, but it cannot simply be assumed or stipulated that being entitled to expect S to X entails that S ought to X. The onus is on Goldberg to offer such an account that implies this entailment, but this is not something that he has done in this connection. The account of practice-generated entitlement that Goldberg sketches does not have the result that being entitled to expect X to S entails that S ought to X.

Goldberg offers this general analysis of the relevant sense of when a subject has a practice-generated entitlement to an expectation:

PGE: For any subject S, practice ϕ with standards ψ , and expectation E, if (i) S participates in ϕ , (ii) ϕ is a legitimate practice, and (iii) calling into question the legitimacy of E would call into question the legitimacy of ϕ or ψ , then S is entitled to E. (Goldberg, 2017, p. 2869)

PGE is supposed to be wholly general, with the entitled expectations regarding when others should have known only a specific instance of this more general principle. Goldberg motivates and illustrates PGE with the following (non-epistemic) example:

GARBAGE: It is standard practice in Mellissa's [sic] community for the garbage to be picked up on Wednesdays, and standard practice as well that when one has garbage to be picked up one puts one's full bins by the curb (to enable easy pickup by the sanitation trucks). One Tuesday evening Melissa puts her full bins by the curb for pick-up the following day. Unfortunately, Wednesday comes and goes and her garbage is not picked up. Melissa is upset. She expected her bins to be emptied on Wednesday. Given the prevailing standard practice, it seems clear that she was entitled to do so. (Goldberg, 2017, p. 2869)

Here PGE is satisfied: Melissa is entitled to expect her bins to be emptied, since she participates in the practice of garbage collection, the practice is legitimate, and calling into question the legitimacy of her expectation would call into question the legitimacy of the practice of garbage collection or the standard that it is collected on

Wednesdays. I grant that Melissa is entitled to expect her garbage to be collected. What I contest is that this entails that the sanitation workers ought to have collected her garbage.

If the sanitation workers saw her bins but did not empty them, then it is true that they ought to have done so. If they did not see the bins because they were playing on their phones as they passed by Melissa's house, rather than doing their jobs, then I also grant that they should have emptied her bins. But suppose instead that through no fault of their own, no sanitation workers saw the bins.¹² Perhaps the bins were incredibly well camouflaged by local pranksters, such that they were invisible to passersby. In this case, I submit, it is not true that the sanitation workers ought to have collected the garbage.

Or take a more dramatic example: suppose that the garbage truck exploded on the way to Melissa's house, after picking up a box of discarded landmines, tragically killing all the sanitation workers.¹³ Again, I submit it is not true that the sanitation workers ought to have collected the garbage. Assuming that ought implies can, then given that there is no way that they could have collected the garbage (in view of their being dead), it is hard to see how it could be true that they ought to have collected it. None of this detracts from the fact that Melissa would be entitled to expect that her rubbish bins would be emptied.

It is not merely that the sanitation workers cannot be blamed in these cases for failing to collect the rubbish. Blamelessness is not equivalent to justification.¹⁴

¹² Goldberg says in passing that cases of not seeing the bins raise "very interesting questions", although he does not concede that in such cases it is false that the subjects in question ought to do what others are entitled to expect of them. (Goldberg, 2017, p. 2872; footnote 17)

¹³ Thanks to Michael Hannon for suggesting this colourful example.

¹⁴ For an argument that epistemic justification is not blamelessness, see Madison (2014). Neither is epistemic justification excusability; e.g. see Madison (2018).

Rather, the workers cannot be properly criticized in any sense for failing to meet the norm of collecting all the garbage in Melissa's neighborhood on Wednesdays, which is good evidence that it is not true that they ought to have collected it.¹⁵ So even non-epistemic cases of practice-generated entitlements appear to admit of possible gaps between someone being entitled to expect S to X and when S ought to X. Further argument is needed to close the gap. In the absence of such an account, and given the examples I offered above illustrating how it might be perfectly reasonable to expect someone to know something, without this entailing that they ought to know it, my argument against Goldberg's account of the should-have-known phenomena stands.

What is the cost of rejecting Goldberg's should-have-known analysis?

Goldberg says that if we reject his SHK, then we need to reject this plausible principle:

(*) If one is entitled to an expectation whose satisfaction requires another subject S's knowledgeableness, S's failure to know warrants the criticism that S should have known. (Goldberg, 2017, p. 2884)

I grant, however, that this principle is true. But notice that it only says when one is warranted in criticizing S for not knowing. It does not say anything about when it is *true* that S should have known. If you are entitled to expect someone to know something, but they do not, then their ignorance warrants criticism. But I have been arguing that it does not follow from the fact that one is entitled to expect a particular person to know, that they should know (so that if they do not know, then they should have known).

¹⁵ Jennifer Lackey (2007) is one philosopher who is explicit about a common methodology at play in the norms of assertion debate: in order to identify norms of assertion, we need to consider actual and possible cases of assertion, and then look for the conditions under which it is appropriate to *criticize*, and by contrast, *praise* particular assertions. Criticism reveals what we judge to be (epistemically) *wrong* with an assertion; praise, on the other hand, reveals that no norm of assertion is violated.

Goldberg thinks that there is another cost in rejecting his analysis. If we reject the analysis, then what should we say about Kornblith's objection-ignoring scientist, the doctor who failed to keep up with the journals she ought to have kept up with, the daycare worker who failed to read the allergy report of the new child in his care, etc., Do we have to count them as having justified beliefs, if we stipulate that their beliefs conform to the evidence that they currently possess? While the cases differ, I agree that in all of the cases Goldberg discusses the subjects lack justification. But this is for various reasons, none of which are social.

In some cases, the subject mishandles the evidence they have. In others, like the arrogant scientist or inattentive Rodriguez, the subjects do not have the evidence they ought to have, or sometimes they do not even have the beliefs they ought to have, due to intellectual vice (e.g. arrogance, closed-mindedness, inattentiveness, etc.). If the doctor failed to keep up with the journals that they ought to have, it is relevant why they did not. If it was due to arrogance, intellectual laziness, or insouciance, then their failure to know will also have an explanation stemming from their epistemic vices.¹⁶

In some cases, like the doctor or daycare worker, while the subjects are blameworthy and rightly subject to criticism in some sense, it is not obvious that their failing, or their only failings, are purely epistemic. Ignorance is not always in itself an epistemic failing. Doctors and daycare workers have professional and moral obligations to know various things. Accordingly, their failings to have the relevant knowledge in these cases may be professional and moral failings.

¹⁶ According to some forms of virtue epistemology, it is sometimes appropriate to explain questionable beliefs and other epistemic failings by reference to intellectual character vices. For example, in developing what he calls vice epistemology, Quassim Cassam has explored the way that epistemic vices, conceived of as character traits, obstruct effective and responsible inquiry. See Cassam (2016); (2019).

In light of the fact that there appears to be a gap between when it is true that someone should have known and whether others are entitled to expect someone to know, or when one might be warranted in saying that someone should have known, let us explicitly focus on the former issue. We can now ask: is social expectation itself sufficient to ground the truth of when someone should have known?

4. Social Expectations

Let us consider another well-known case Harman used to motivate the thesis that social considerations can make a difference to whether or not a subject has knowledge.

Harman's Postal Case: Donald has gone off to Italy. He told you ahead of time that he was going; and you saw him off at the airport. He said he was to stay for the entire summer. That was in June. It is now July. Then you might know that he is in Italy. It is the sort of thing one often claims to know. However, for reasons of his own Donald wants you to believe that he is not in Italy but in California. He writes several letters saying that he has gone to San Francisco and has decided to stay there for the summer. He wants you to think that these letters were written by him in San Francisco, so he sends them to someone he knows there and has that person mail them to you with a San Francisco postmark, one at a time. You have been out of town for a couple of days and have not read any of the letters. You are now standing before the pile of mail that arrived while you were away. Two of the phony letters are in the pile. You are about to open your mail. I ask you, "Do you know where Donald is?" "Yes", you reply, "I know that he is in Italy." You are right about where Donald is and it would seem that your justification for believing that Donald is in Italy makes no reference to letters from San Francisco. But you do not know that Donald is in Italy. (Harman, 1973, pp. 143-4)

Harman's own analysis of this kind of case is that the reason that you do not know that Donald is in Italy is because of the ease of obtaining the postal evidence.

Harman writes of the same sort of case, but with different character names:

It is important in this case that Mary [the subject] could obtain the misleading evidence. If the evidence is unobtainable, because Norman [or Donald] forgot to mail the letter after he wrote it, or because the letter was burned up in a mail fire in San Francisco, or because the letter was delivered to the wrong building where it will remain unopened, then it does not

keep Mary from knowing that Norman is in Italy. [...] There seems to be two ways in which such misleading evidence can undermine a person's knowledge. The evidence can either be evidence that it would be possible for the person to obtain himself or herself or evidence possessed by others in a relevant social group to which the person in question belongs. (Harman, 1980, p. 164)

Commenting on Harman's postal cases and his analysis of them, John Pollock notes, rightly, that it is not the mere easy possibility of obtaining the postal evidence that makes it a defeater, since if the trick letter from San Francisco had been hidden under the doormat, the subject would still know that Donald is in Italy, despite the letter still being easily available. (Pollock, 1986, p. 192) Looking under one's doormat is easy enough to do, and so is readily available. But that is not enough to make information hidden under one's doormat a defeater.

Instead, Pollock suggests that a defeater arises in Harman's postal case because of social expectations. Pollock writes,

We are 'socially expected' to be aware of various things. We are expected to know what is in our mail. If we fail to know all these things and that makes a difference to whether we are justified in believing some proposition, *p*, then our...justified belief in *p* does not constitute knowledge. (Pollock, 1986, p. 192)

Using the notion of social expectations, Pollock goes on to propose an account of social defeat. Pollock holds that a proposition is a social defeater for *S* if and only if it is of a sort *S* is expected to believe when true. (Ibid.) Similar to Goldberg's account which follows it, it is the violation of a relevant social expectation that is thought to defeat knowledge.

On Pollock's analysis of his variation of the postal case where the trick letter is hidden under the doormat, the hidden letter does not defeat in our social environment, despite being easily available. This is because checking under our

doormats is not something we are socially expected to do. On the other hand, in a society where everyone is expected to check under their doormats when they return home, the hidden letter would defeat your knowledge that Donald is in Italy.

(Pollock, 1986, pp.192)

But do we need a notion of social defeat to explain why knowledge is absent in Harman's postal case? If we cannot explain the case without it, that would motivate postulating the existence of social defeat. But if we can explain the case without introducing a new kind of defeater or defeat phenomenon, then considerations of parsimony dictate that we should not. I suggest that we do not need a notion of social defeat to explain Harman's postal case. I propose the same unified treatment here as the assassination cases above: either you do know that Donald is in Italy, and it is a matter of evidential luck; or you do not know that Donald is in Italy, and that is because there is too much veritic luck present. Either way, nothing inherently social is involved.

For my own part, I think a stronger case can be made for interpreting this example as an instance of knowledge. Like Harman's assassination cases, it is important to consider on what basis you form the belief that Donald is in Italy. On the assumption that the belief is held on the basis of Donald's testimony alone, the belief is safe. Your belief is in a sense lucky, since you could have easily opened the misleading letters. But this appears to be a case of benign evidential luck: you are lucky to have the actual evidence you do, since in very nearby worlds, you have *different* total evidence which includes the misleading letters. But as noted above, evidential luck is compatible with knowledge.

On the other hand, suppose one takes one's cue from Harman. Then the method used could be construed as including all easily available evidence, and so includes not only Donald's testimonial evidence, but also the easily available misleading postal evidence. In that case, your belief is unsafe. This is because the method used is unreliable given the nearby worlds where you form a false belief about Donald's whereabouts on the basis of the trick letters. If this is the correct way of specifying the belief-forming method you are relying upon, your belief that Donald is in Italy is infected by too much veritic luck, and so plausibly does not count as knowledge. But here again, there is nothing inherently social. Just as on one interpretation of the assassination cases, your belief does not amount to knowledge, as there is simply too much veritic luck in the environment such that your belief is not safe. So here too, we have the resources to explain why Harman's post case is one of knowledge, but even if it is one of lacking knowledge, it is not because of any social factors.

Suppose one rejects my disjunctive analysis of the post case, and following Pollock, persists in holding that it is due to social expectations operative in the case that knowledge is lacking. Even then, we should still consider: are mere social expectations sufficient to defeat knowledge? I have already argued that it does not follow from the fact that one is socially expected to be aware of X that one ought to be aware of X. Furthermore, there is no straightforward connection, if any, between a mere social expectation and what one ought to believe. To illustrate that social expectations are insufficient to defeat knowledge, consider this case from Kevin Meeker:

AMANDA: Suppose that Amanda is staying overnight in a hotel of a certain culture that she had previously never heard of because of an emergency landing on the way to another country. Suppose further that this society has certain expectations of all those who are within

its boundaries – but they do not inform those who arrive of these expectations. So Amanda is blissfully unaware of them. [I]magine that this culture expects all those within its borders to check under their doormats for mail and other information. In such a situation, it seems clear that no matter what information is secreted under her doormat, it will not count as a defeater because Amanda *should not be expected* to check there. (Meeker, 2004, pp. 163-4, emphasis in original)

Meeker's AMANDA case is a counterexample to several of the accounts of social defeat we have considered so far. It is a counterexample to Harman's account, since no post under Amanda's doormat would constitute a defeater for her, despite being easily available to her. It is a counterexample to Pollock's account, since no post under Amanda's doormat would constitute a defeater for her, despite the societal expectation that she check under her doormat for mail and other information. This case is also a counterexample to Goldberg's SHK principle: it is not the case that Amanda should have known what was under her doormat, even though the society she finds herself in is entitled to expect that she knows what is under the mat. The case helpfully illustrates the gap between when it is true that Amanda should have known, and when others might be entitled to make the charge that she should have known.

In response to his AMANDA case, Meeker introduces the notion of a *legitimate* social expectation (e.g. Meeker, 2004, pp. 164-165; Goldberg also makes use of the notion of *legitimate* social practices, e.g. Goldberg, 2018, p. 170). Meeker's proposal is that defeaters can arise from legitimate social expectations, in contrast to any old social expectations. A natural question, of course, is what makes a social practice, and its associated expectations, legitimate?

5. *Legitimate Social Expectations*

The case of Amanda illustrates that mere social expectations are insufficient to defeat knowledge. In light of this, the suggestion to consider is that only *legitimate* social expectations can have this effect. In light of the Amanda case, Meeker refines his account of social defeat thus:

A proposition *p* is a socially sensitive defeater for S's belief *b* if and only if *p* is a defeater for *b*, S is *legitimately* socially expected to believe *p*, and S does not believe *p*. (Meeker, 2004, p. 165)

Again, the pressing question is: what makes a social practice, and its associated expectations, legitimate? Somewhat unsatisfyingly, given that the notion of legitimacy is the crux of his proposal, Meeker does not attempt to provide an account.

Instead, Meeker claims that recognizing the legitimacy of social expectations is on par with recognizing instances of knowledge. (Ibid.) To adopt such an approach is to endorse a kind of Chisholmian particularism: just as we can intuitively recognize instances of knowledge when we are presented with them, the suggestion is that we can know legitimate social expectations when we see them too.¹⁷

I am sympathetic to particularism as a methodological approach in exploring the nature of knowledge and epistemic justification. But consistent with this kind of particularism is the thesis that we can attempt to formulate, at least partially, illuminating conditions for the instantiation of certain epistemic properties. I propose the same for elucidating the notion of legitimate social expectations. That is, given an intuitive idea of which social expectations are legitimate, we can seek to give an

¹⁷ This sort of particularism in epistemology is associated with Roderick Chisholm's work on the Problem of the Criterion (e.g. Chisholm, 1966).

account of what all the intuitive particular cases have in common that make them instances of *legitimate* social expectations, rather than something else.

According to Meeker, it is intuitively obvious that the social expectation operative in the Amanda case is not legitimate. I agree. While he does not attempt to offer an account of legitimacy, his remarks on the case are suggestive. Of Amanda's situation, Meeker writes:

The social expectation in this case was, in a word, "illegitimate". Of course, if Amanda *was informed* of this custom, then she could have been blamed for failing to check under the doormat; more precisely, the expectation is illegitimate because Amanda was not properly *informed* of it, not because it is somehow intrinsically or essentially illegitimate to leave information under someone's doormat. (Meeker, 2004, pp.164-165, emphasis added)

Generalizing from Meeker's comments, my suggestion is that the relevant kind of social expectation is legitimate only if you have, or ought to have, evidence for it. If Amanda had been informed of the custom of checking under doormats, she would then have evidence that this was expected of her. As Meeker rightly notes, if having been so informed, and so having evidence of the expectation, she failed to check under her doormat, she could be rightly criticized.

We can see that it is evidence doing the epistemic work, and nothing inherently social, by considering a case which combines Harman's Post Case with Meeker's AMANDA example, but with a further twist. Suppose that Amanda had a justified true belief that Donald is in Italy for the whole summer in the way Harman describes in his original case. Meanwhile, as the result of an emergency landing, Amanda finds herself in a hotel in a certain culture that expects people to check under their doormats for mail and other information, as in Meeker's original case. Now for the twist: learning of Amanda's location, Donald couriers the trick letters saying that he is in San Francisco to the hotel to be placed under Amanda's doormat. Unlike in Meeker's original case where the expectation was illegitimate, however, suppose that

Amanda has been informed of the expectation to check under her doormat for post, but she has not done so.

As Amanda has been informed of the practice, she has evidence that she should check under her doormat for possible information. If she did check under her doormat, she would then have found the letter saying that Donald is in San Francisco for the summer. So the contents of the letter serve as a *normative defeater* for her belief that Donald is in Italy. That is, the contents of the letter amount to a belief that Amanda *ought to have* that indicates that her belief about Donald's whereabouts is false or unreliably formed or sustained. And why ought Amanda have this postal evidence which would amount to a normative defeater? Not because it is merely socially expected of her to have looked under her doormat, as the original AMANDA case shows, but rather, because she has been given *evidence* that she ought to check for post and other information under her doormat.

But what is epistemically salient here is the *evidence*, not the social expectation; the evidence is the thing that makes the difference to what Amanda ought to believe. If this is correct, it is not the social expectation that makes it true that the subject should have known. Rather, it is simply that given the evidence the subject has, they should have known what was under her doormat. In addition, this would not be the kind of case Goldberg is interested in, since it is just a straightforward one of the subject not responding properly to their evidence, not one of someone who should have known, even when the subject is in no position to know.

Suppose instead that Amanda had been informed of the custom of checking under doormats, but she was so incredulous of what she took to be an outlandish practice, that she simply did not believe what she was told, despite being sincerely

told of the practice by trustworthy locals. Still, she would have evidence of what was expected of her, so if she failed to look and find information under her doormat, it would have been true that she should have known. Given the evidence Amanda *should have had*, had she not been so incredulous and mistrusting, she should have known. But then again this is just a regular (non-social) normative defeat case: one ought to have some evidence, but does not.

Crucially, nothing about either of these two cases where Amanda is given evidence of the practice is inherently social. They simply involve evidence about something which happens to be social, namely a practice involving the transmission of information under doormats peculiar to this community. But having evidence about something social does not by itself make the social somehow epistemically important. Compare: if one had evidence of something political or aesthetic that was inconsistent with one's political or aesthetic beliefs, we should not conclude that there is a special class of political or aesthetic defeaters, or kinds of uniquely political or aesthetic defeat phenomena. Any defeat, and the sense in which it is true that Amanda should have known in these cases, can be fully explained by the existence of mental state and normative defeaters. There is no need to appeal to a special class of social defeater, nor any essentially social defeat phenomenon.

Before concluding, it is worth pausing to examine Goldberg's account of what makes a social practice legitimate. While Goldberg does not provide a full-fledged analysis, he does say more than Meeker does on this score. In the course of describing practice-generated entitlements, which I discussed above, Goldberg tends to say things like this about the legitimacy of social practices:

When the practice itself is legitimate – it is an ongoing and *recognized* practice, its standards are widely *acknowledged*, there have been no serious questions as to the propriety of either the practice or its standards [...] (Goldberg, 2018, p. 165, emphasis added)

[...] when it comes to the sorts of practices and practice-generated expectations that emerge in the course of our interpersonal relationships, here it would seem that the legitimacy of the practice itself does depend on the *acknowledgement* of its standards by the people in the relationship. (Goldberg, 2017, p. 2890, emphasis added)

More formally, Goldberg defines the general requirements on the legitimacy of a practice (and so also a social practice) as follows:

LEGIT: A practice ϕ is legitimate only if (i) it is an ongoing and/or recognized practice, (ii) its standards π are widely acknowledged, and (iii) there have been no serious questions as to the propriety of either the practice or its standards. (Goldberg, 2018, p. 170)

What are we to make of requiring that it is a necessary condition of a practice being legitimate that it must be *recognized* and *acknowledged* by those within the practice?¹⁸ My suggestion is that Goldberg's concern is the same as Meeker's when the latter claims that what made the expectation of Amanda to look under the doormat illegitimate was that she was not properly informed of it. Goldberg's remarks on the importance of recognition and acknowledgment fit with my suggestion that the kind of social expectation at issue is legitimate only if you have, or ought to have, evidence for it. If Amanda had been informed of the custom of checking under doormats, and acknowledged and recognized the practice, she would then have evidence that this was expected of her.

But what is epistemically salient here is the recognition and acknowledgment.

Here again it seems that what is important is evidence, and not the social expectation

¹⁸ Note that Goldberg only offers necessary conditions of a practice being legitimate. A complete account will obviously require additional conditions to rule out intuitively bad practices as illegitimate that are nevertheless ongoing, acknowledged, not questioned, and so on. A terrible practice believed good and legitimate by its practitioners is of course insufficient for it to be legitimate. I thank an anonymous referee for this observation.

per se. If this is correct, it remains the case it is not the social expectation that makes it true that the subject should have known. Rather, it is simply that given the evidence the subject has, they should have known. But then again, this is simply a case of a regular mental state defeater.

On the other hand, if the social practice was widely acknowledged and recognized, but Amanda just failed to give it any proper credence, she would still have evidence of what was expected of her. If she then failed to look and find information under her doormat, it would have been true that she should have known. But then this too is just a regular (non-social) normative defeat case: she ought to have some evidence, but does not. Any defeat, and the sense in which it is true that Amanda should have known in these cases, can be fully explained by the existence of mental state and normative defeaters, without the appeal to anything social.

In sum, my aim in this article has been to explore the questions of whether there are social defeaters, or else a kind of uniquely social defeat phenomenon that falls under one of the three general kinds of defeaters that epistemologists tend to admit. I argued that despite initial appearances to the contrary motivated by a number of interesting and influential cases, we have been given no reason to accept a special class of social defeater, nor any essentially social defeat phenomenon. Rather, we can explain putative cases of social defeat with our existing epistemological machinery. So while justification and knowledge undoubtedly have social dimensions, we have yet to see that there is an inherently social form of defeat.¹⁹

¹⁹ Thanks to audiences at the Cardiff University workshop “Social Epistemology and Epistemic Normativity” (in particular, I recall helpful comments from Mona Simion, Cameron Boulton, Alessandra Tanesini, Orestis Palermos, and Sandy Goldberg), and the 2019 Bled Epistemology Conference (thanks to Peter Graham, Jack Lyons, Chris Kelp, Mona Simion, Kelly Becker, Bob Beddor, Michael Hannon, and Hilary Kornblith for their helpful comments). Thanks especially to Sandy Goldberg, Rhiannon

Sources Cited:

- Altschul, Jon. n.d. "Epistemic Entitlement," *Internet Encyclopedia of Philosophy* <http://www.iep.utm.edu/ep-en/>.
- Ballantyne, Nathan. (2015) "The Significance of Unpossessed Evidence". *The Philosophical Quarterly* 65: 315-335.
- Becker, Kelly and Black, Tim (eds.) (2012) *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press.
- Becker, Kelly. (2018) "The Sensitivity Response to the Gettier Problem". In Stephen Hetherington ed. *The Gettier Problem*. Cambridge: Cambridge University Press: 108-124.
- Burge, T. (2003). "Perceptual Entitlement." *Philosophy and Phenomenological Research* 67: 503-548.
- Cassam, Quassim. (2016) "Vice Epistemology". *The Monist* 99: 159-180.
- Cassam, Quassim. (2019) *Vices of the Mind*. Oxford: Oxford University Press.
- Chisholm, Roderick. (1964) "The Ethics of Requirement". *American Philosophical Quarterly* 1: 147-153.
- Chisholm, Roderick. (1966) *Theory of Knowledge*. 1st ed. Englewood, N.J.: Prentice-Hall.
- Conee, Earl and Feldman, Richard. (1998) "The Generality Problem for Reliabilism". *Philosophical Studies* 89: 1-29.
- De Almeida, C and Fett, J.R. (2015) "Defeasibility and Gettierization: A Reminder". *Australasian Journal of Philosophy* 94: 152-169.
- Dretske, F. (2000). "Entitlement: Epistemic Rights without Epistemic Duties?" *Philosophy and Phenomenological Research* 60: 591-606.
- Gettier, Edmund. (1963) 'Is Justified True Belief Knowledge?' in *Analysis* 23: 121-3.
- Goldberg, Sanford C. (2017) "Should Have Known". *Synthese* 194: 2863-2894.
- Goldberg, Sanford C. (2018) *To The Best of Our Knowledge*. Oxford: Oxford University Press.
- Harman, Gilbert. (1968) "Knowledge, Inference, and Explanation". *American Philosophical Quarterly* 5: 164-173.

James, two anonymous reviewers and an Editor at the *Canadian Journal of Philosophy* for detailed written feedback on earlier drafts of this paper.

- Harman, Gilbert. (1973) *Thought*. Princeton: Princeton University Press.
- Harman, Gilbert. (1980) "Reasoning and Evidence One Does Not Possess".
Midwest Studies in Philosophy 5: 163-182.
- Hetherington, Stephen. (1998) "Actually Knowing". *The Philosophical Quarterly*
48: 453-69
- Lackey, Jennifer. (2007) "Norms of Assertion". *Nous* 41: 594-626.
- Lackey, Jennifer. (2008) *Learning From Words*. Oxford: Oxford University Press
- Lehrer, Keith and Paxson, Jr., Thomas. (1969) "Knowledge: Undefeated Justified
True Belief". *The Journal Of Philosophy* 66: 225-37.
- Meeker, Kevin. (2004) "Justification and the Social Nature of Knowledge".
Philosophy and Phenomenological Research 69: 156-172.
- Madison, B.J.C. (2011) "Combating Anti Anti-Luck Epistemology". *Australasian
Journal of Philosophy* 89: 47-58.
- Madison, B.J.C. (2014) "Epistemological Disjunctivism and the New Evil Demon".
Acta Analytica 29: 61-70.
- Madison, B.J.C. (2018) "On Justifications and Excuses". *Synthese* 195: 4551-4562.
- Nozick, Robert. (1981) *Philosophical Explanations*. Oxford: Clarendon Press.
- Peacocke, Christopher. (2004) *The Realm of Reason*. Oxford: Oxford University
Press.
- Pedersen, N.J.L.L., & Graham, P. (2020) "Recent Work on Epistemic Entitlement".
American Philosophical Quarterly 57: 193-214.
- Pollock, John. (1986) *Contemporary Theories of Knowledge*. Lanham, MD:
Rowman & Littlefield.
- Pritchard, Duncan. (2005) *Epistemic Luck*. Oxford: Clarendon Press.
- Pryor, James. (2004) "What's Wrong With Moore's Argument?". *Philosophical
Issues* 14: 349-378.
- Pryor, James. (2018) "The Merits of Incoherence". *Analytic Philosophy* 50: 112-
141.
- Unger, Peter. (1968) "An Analysis of Factual Knowledge". *Journal of Philosophy*
65: 157-170.
- Wright, C. (2004). "Warrant for Nothing (and Foundations for Free)?" *Aristotelian*

Society Supplementary Volume 78: 167–212.