

ON PROVIDING EVIDENCE

CHARITY ANDERSON

charity.r.anderson@gmail.com

ABSTRACT

Obligations to provide evidence to others arise in many contexts. This paper develops a framework within which to understand what it is to provide evidence to someone. I argue that an initially plausible connection between evidence-providing and evidence-possession fails: it is not the case that in order to count as *providing* evidence to someone, the intended recipient must *have* the evidence. I further argue that the following is possible: evidence is provided to an agent, the agent does not *have* the evidence, but it is not the case that the agent is culpably ignorant of the evidence.

INTRODUCTION

Various social relationships place us under obligation to provide evidence to others. When a tenant enters a legal agreement with a landlord, both place themselves under an obligation to divulge certain information. The tenant may be obligated to let the landlord know if damage occurs to the property. The landlord will be expected to give the tenant notice if the property is to become unavailable for rent, or should the property need to be accessed for maintenance. Obligations to provide evidence also arise in less formal settings, such as when a spouse ought to let his partner know if he cannot pick up the children as previously agreed. Similarly, one ought to communicate to one's friend in the event that one is running significantly late, or needs to cancel an appointment. In some cases, serious consequences result if the information is not provided, as when a parent fails to disclose that their child has the measles to the parents of their child's playmate. In these kinds of situations, the ignorant party would be within their rights to complain that the other person ought to have communicated some information, but failed to do so.

Sometimes the means by which the information ought to be provided is clearly specified: an application for a visa might require exactly two forms of acceptable identification, to be submitted via mail. In this case, when one has done what is required on the checklist, one has provided the information. In other cases, the conditions are not nearly so straightforward, as when a significant other wants to communicate during a break-up that they are 'still friends.'

It will be helpful to begin with some cases in mind. The following cases illustrate some of the various kinds of situations in which an obligation to provide evidence may arise:

Childcare: A caregiver arrives at a home to take care of a child for the first time. The parent fails to mention to the caregiver that the child has a severe nut allergy.

Late Friend: You are on your way to meet your friend for coffee at the local café. A call comes in from work that requires immediate attention and you must head to the office instead. You text your friend to let her know you need to cancel.

Seeking Agnostic: Kate is an agnostic but is open to theism. She considers the evidence for the existence of God and finds it lacking. She wonders why, if there is a God, he does not provide her with more evidence that he exists.

In each of these cases, there is a *prima facie* obligation or expectation that an agent will provide evidence or let the other person know some information. *Late Friend* and *Childcare* demonstrate the widespread and ordinary nature of the obligation to provide evidence. We commonly find ourselves in situations where we ought to let someone know some information – even information as simple as the time or place of a meeting. Whether an obligation to provide evidence arises in contexts such as *Seeking Agnostic* is more controversial. Some think that God has obligations to provide evidence to us that he exists. Others think he does not have such an obligation. For the most part, this paper will consider situations of the ordinary kind.

Although the obligation to provide evidence also arises in legal contexts, for the most part, I'll set to one side consideration of legal contexts, as subtleties unique to these contexts may not carry over to a generalized notion.

Given the pervasive nature of these obligations, one question we might be interested in asking is: *Under what conditions does the obligation to provide evidence arise and when has it been met?* Clearly, social norms will play a role. Expectations of one another arise out of both implicit and explicit social norms. Context will be relevant. As will Gricean norms of conversation – in some cases your listener will expect you to say the strongest thing you know. I will not attempt to address here the many (and very interesting) questions regarding specific obligations to provide evidence. Given the variability with which the obligations arise, we can expect the corresponding norms to be similarly unruly. (And the task of uncovering them strikes me as having a more natural home within the domain of ethics. This paper will lay some of the groundwork relevant to such a project.) Thus, I will set aside the issue of identifying specific obligations and norms, and focus instead on what I take to be a more fundamentally *epistemic* question: *What is it, in general, to provide evidence that p to S?* This paper examines that question.

Despite an initial appearance of simplicity, the question proves intractable.¹ My aim, thus, is not to give an analysis of the notion. Rather, my aim is to make progress understanding the notion and some of the surrounding issues. The paper will proceed as follows. Section 1 introduces terminology. In Section 2, I discuss the relationship between *providing* evidence, on the one hand, and *having* evidence, on the other. I show that an initially plausible connection between evidence-providing and evidence-possession fails to be generally true. It is not the case that in order to count as *providing* evidence to someone, the intended recipient must *have* the evidence – or so I shall argue.

1 One might worry that this question does not have an answer. That is, one might adopt a skeptical stance concerning whether there is a unified notion of *providing evidence*. While granting the widely context sensitive nature of obligations to provide evidence, the fact that there is a core epistemic notion of *having evidence* licenses an expectation of a corresponding notion of *providing evidence*. There is clearly some epistemic sense of *having evidence* – even though this sense does not line up with all senses of *having*. I suggest we can expect something similar for *providing evidence*. Thanks to Mark Schroeder here.

Section 3 investigates complications that arise in providing (or attempting to provide) evidence to another person. Factors including the other person's prior evidence, the other person's epistemic character, and additional factors such as expertise, abilities, attention, and cognitive penetration, all problematize the expectation that the agent will 'uptake' the information provided. There are numerous explanations for why agents have the evidence they have and why they fail to have evidence available to them: some explanations are evidential, some non-evidential. One result of these complicating factors is that it is not generally true that when evidence is provided to an agent but the agent does not have the evidence, the agent is *culpably ignorant* of the evidence. Some available explanations involve culpable ignorance, but some involve non-culpable ignorance.

Finally, in Section 4, I make an application. The notion *providing evidence* plays a prominent role in an argument from philosophy of religion: the argument from divine hiddenness. In closing, I'll explore how the points discussed in this paper might be used to shed light on this argument. The results open a new way to think about the argument.

I. PRELIMINARY TERMINOLOGY

To begin, it will be helpful to anchor some key terminology. I'll attempt to give discipline to the terminology of 'providing evidence' by drawing on a Bayesian framework. 'Providing evidence to S that *p*' will be understood along the following lines: 'evidence for *p*' and 'evidence that *p*' will be understood in terms of *probability raising*. Thus, some piece (or set) of evidence *e* is evidence for *p* when *e* raises the probability of *p*. I'll assume that evidence is propositional (though I don't expect anything substantive will hang on this). To simplify matters, I'll think of *e* as evidence for *p* when *e* raises the probability of *p* on the set of initial priors combined with the fixed set of common background knowledge,² and I'll assume that the initial priors are objective. Understanding 'evidence for *p*' as relative to the set of initial priors plus a common set of background knowledge is useful because it permits a framework for understanding how some evidence *e* is evidence for *p* without the complication of taking into account a particular agent's prior knowledge. (Since this set varies agent to agent, taking it into account would not allow a fix for the terminology.)

It's worth noting that on this way of thinking, sometimes *e* is evidence for *p* even if, for some agent, learning *e* lowers the probability of *p* for that agent. (As when, for example, S believes it is 'opposite day,' where people say the opposite of what they mean). Thinking of the *evidence for* relation as relative to a set of objective priors and common background knowledge gives us a stable reference point, and thus a fix on terminology, however idealized.³ As is standard in this framework, when an ideal agent learns new information, she

2 It is an option to drop the addition of the set of common background knowledge and instead think of *e* as evidence for *p* relative only to the initial priors, but this will deliver the result that *e* is evidence for *p* even if it is not the case that *e* raises the probability of *p* for any actual human being with normal background knowledge. There is a choice point here. Of course, it is not a simple matter to specify what belongs in the set of background knowledge. The ideas in this paper do not depend on nor do they provide a particular resolution to this issue. Thanks to Brandon Fitelson for drawing this to my attention.

3 Alternatively, one could use only the objective set of initial priors, without adding a set of common background knowledge. For our purposes, which way one proceeds will not make a significant

updates on her evidence using classical Bayesian conditionalization, resulting in a set of epistemic probabilities that respect the Kolmogorov axioms.⁴ (The fact that we are not ideal agents will matter a bit later on.)

Given this picture, one way to understand *S provides evidence that p to S₂*, is as follows:

Weak provide: S provides evidence that *p* to S₂ when S provides evidence *e* to S₂ and *e* raises the probability of *p*.⁵

One disadvantage of *weak provide* is that it seems to make *providing evidence* too easy. One will count as providing evidence even when the evidence provided only barely raises the probability of *p*. The increase might be negligible.⁶ Our ordinary notion of providing evidence suggests something stronger than mere probability raising – it suggests something more in the ballpark of evidence that indicates *whether p*. In this way, ‘providing evidence’ could be taken as shorthand for ‘providing *enough* evidence,’ where ‘enough evidence’ is something along the lines of evidence that raises the probability of *p* to or above the threshold for outright belief.⁷ I’ll call this *strong support*, and articulate a principle accordingly:

Strong provide: S provides evidence that *p* to S₂ when S provides evidence *e* to S₂ and *e* strongly supports *p*.

where ‘strongly supports’ is understood as follows:

strong support: *e* strongly supports *p* when *e* raises the probability of *p* to/above the threshold for outright belief.⁸

We can identify three initially intuitive roles we might hope that strong evidence for *p* will play: (1) raise the probability of *p*; (2) raise the probability a significant distance in terms of absolute increments – for example, start at a low probability and move to a high probability; and (3) deliver the result that the posterior probability of *p* is at or above the

difference. Both involve idealizations that one might think stretch the limits of plausibility, but these are general issues for the Bayesian framework, rather than problems unique to my project.

4 Whether the agent uses classical conditionalization or Jeffrey conditionalization won’t matter for my purposes here.

5 Here and throughout, I omit time-indexing; complexities may arise when this is added.

6 Of course, if *p* is incredibly improbable to begin with, a ‘weak’ nudge might still be viewed as strong evidence according to certain natural ways of thinking about evidential boosts. What I’m calling ‘weak’ evidence might, for example, double or triple the probability of *p*, and thus be strong evidence using the Bayes factor as a guide to strength of evidence. Since I’m assuming that our ordinary notion of *providing evidence* indicates something more along the lines of placing one in a position to rationally believe that *p*, or to know *p*, it will make sense to call the former kind of boost *weak*, while acknowledging that there are other – and more standard – ways to measure strength of evidence.

7 The question of where exactly the threshold for outright belief is set takes us outside the scope of this paper.

8 One result of thinking about the notion this way is that when one is initially near the threshold, even a small raise that ordinarily would be thought of as weak can count as strong evidence. This is a cost, but there will be costs any way we proceed.

threshold for belief. Each is intuitive but it is not clear the framework can respect them all. For example, (2) is problematic when a proposition has an initial high prior probability. And if the prior probability is initially high, (3) could be satisfied without (1) being met (and moreover, without (2)). There are various choice points here; I won't try to resolve them all. *Strong provide* does not include (2) as a necessary condition, but nevertheless gives us a working gloss suitable to our purposes.

Of course, having some evidence that strongly supports *p* does not guarantee that one's total body of evidence strongly supports *p*. Further difficulties arise when using this gloss to explain what it is to provide evidence *to a subject*. One might think that providing evidence that *p* to *S* requires both that the evidence strongly supports *p* and that it raise *S*'s probability that *p*. This idea is captured in the following principle:

Strong provide-agent: *S* provides evidence that *p* to *S*₂ when *S* provides evidence to *S*₂ that strongly supports *p* and the evidence raises the probability of *p* for *S*₂.

While this is a natural way to extend the picture, as the reader may foresee, difficulties emerge. A central worry concerns the fact that adding *e* to *S*'s body of evidence does not always result in raising the probability of *p* for *S*. (I will address some of these difficulties in Section 3.)

A further issue attending the notion of *providing evidence* is this: there is a natural sense in which I have provided you with evidence that *p* when the evidence is in your possession, but you haven't looked at it yet, and thus your probability that *p* has not changed. Suppose, for example, that I write *p* on a piece of paper, place it in an envelope, and hand it to you. There's a very natural sense in which it is correct to say that I have provided you with evidence that *p*, even though you have no idea what the contents of the envelope are. You *have* the evidence in some sense, but not in the sense the preceding principles describe. This points to the need to distinguish between a loose sense of *providing evidence* and a specifically *epistemic* sense. There will be significant overlap, but not a one-to-one correlation.

(We may place more distance between you and the contents of the envelope if we imagine that I have placed the envelope in the mail, but it hasn't been delivered – so it is not yet in your possession. Intuitions will likely diverge on the question of whether the evidence has been provided to you when the letter is in the mail, but it's not that much of a stretch, in my view at least, to affirm that it has. Note that it is not merely the physical distance that's relevant. It could be that the parcel is in my mailbox, but I'm on holiday in Italy. If we stipulate that the evidence counts as *provided* once it's been delivered, it won't matter that I'm far away.)⁹

Another type of situation that makes trouble for the agent-centric principle is a setting where the recipient already knows *p*. In these cases, providing evidence on the agent-centric conception will be especially demanding – and we might wonder whether it is in every case possible (for example, when the recipient is already certain that *p*).

9 One might be tempted to fix for this by substituting 'would raise *p* if *S* had *e*' for 'raises' in *strong-provide-agent*. But this strikes me as unpromising: not only are counterfactuals unruly, but defeat can interfere.

Thus, although the agent-centric principle is initially promising, I suggest that we not understand what it is to provide evidence that p to S as requiring a raise in S 's probability for p .

A few more preliminary points.¹⁰ I will limit my discussion to cases where the recipient is an individual subject. Much of what I say can be applied to settings where S is a group, though many of the difficulties I raise with respect to providing evidence to individuals will be magnified in a group setting.

An additional point worth noticing is that S can provide evidence that p to someone even if S does not know or even believe p . Suppose Andrew is an arrogant guy. He may not be aware that he has this particular character trait, and thus will not intend to provide evidence of his arrogance when he interacts with others, but he may provide such evidence nonetheless. This draws our attention to two further features of *providing evidence*: first, that it can be done unintentionally; and second, that it doesn't require that the provider *have* the evidence – at least not in the epistemic sense of have (though perhaps in some other sense he has the evidence).¹¹ An interesting consequence of this last point is that *providing evidence* does not involve a condition that some theories of testimony place on the transmission of testimonial knowledge or justification – views whereby the speaker must know p in order for the listener to come to know p on the basis of the speaker's testimony.¹²

A final preliminary observation: there are many structural similarities between *providing evidence* and related notions, such as *letting them know*. Although my discussion focuses on the former, many of the observations I make here can be extended to *letting them know*.

2. PROVIDING AND HAVING EVIDENCE

An initially plausible thought regarding *providing evidence* is as follows: if S provides evidence to S_2 , then S_2 *has* the evidence. Consider an analogy with other instances of *providing*. If Anna provides dinner for Bill, typically we would expect that Bill has the dinner – that is, he eats it, or is about to eat it, or something along these lines. The dinner is in some way in his possession. There is an intuitive connection between *providing* and *having*. We might initially try to capture this connection as follows:

Provide-have: If S provides evidence e that p to S_2 , S_2 has evidence e .

¹⁰ The discussion here suggests the following linguistic observation: the verb 'provide' patterns like a 3-place predicate such as 'put.' It requires a giver, recipient, and something given, as in 'S provides x to y.' It sounds incomplete to say, 'S provides' or even 'S provides x.' Thanks to John Hawthorne for drawing this to my attention.

¹¹ A further issue related to the phenomena of *providing evidence* that I won't be able to discuss is the way in which the act of providing evidence e can itself be the source of additional evidence e^* . The choice of the agent to provide evidence e sometimes reveals something about the agent. Who provides e , and when and how e is provided are facts that may impact how e is interpreted and what additional information is communicated. Thanks to Tom Simpson for pointing this out.

¹² Not all theories of the epistemology of testimony are in agreement on this point. See Lackey (2008) for an example of transmission of knowledge (or justification) via testimony where the speaker does not know, or even believe, the proposition transmitted.

(This principle is neutral with respect to *weak provide* and *strong provide*.) In this section, I argue that this principle is false. It is not a necessary condition on *providing evidence* that the intended recipient *have* the evidence. This is, at first glance, surprising. Perhaps in many cases the recipient has the evidence once it is provided. And typically one aim in providing evidence is that the recipient have it. But that the recipient *have* the evidence is not a constraint on what it is to *provide evidence*.¹³

To elucidate this principle, we'll need a better grasp of what it is to *have* or *possess* evidence. I'll focus on the epistemic sense of *have*, though the conclusion I argue for applies generally. There are various competing options on offer. There are also competing accounts of what *evidence* is. Here I'll limit my discussion to two broad kinds of accounts of evidence in the literature. I'll begin with a family of views, which I'll loosely group together under the classification of 'phenomenal conservatism.' What's distinctive of this group is that its members hold that seemings play a central evidential role – and thus play a role in what it is for S to have evidence.¹⁴ On one side of the spectrum, we find Feldman, who holds that S has *e* as evidence if and only if S is currently thinking of *e*.¹⁵ This is an extremely restrictive picture of what it is to have evidence, and it delivers an equally stringent, and much too demanding, gloss on the notion under investigation. (As soon as I stop thinking about *e*, I no longer have *e* as evidence, on this view.)

A more promising picture of having evidence within this family of views is one where S has evidence *e* for *p* when it seems to S that *e*, and where for S to have evidence for *p*, S's relevant seemings need not be occurrent, but rather where it suffices for S to have the seemings dispositionally. (In cases where *e* = *p* this will mean that S has a seeming that *p*.) The latter obviously provides an account better suited to our purposes here. However, we might still worry that this imposes too strong a requirement on *providing evidence*. It is not clear that it is always within the control of the provider to bring about seemings in another agent, either occurrently or dispositionally. (I'll return to this issue in a moment, as the problem affects all views.)

An alternative theory of evidence, defended by Timothy Williamson (2000), is an account whereby one's total body of evidence consists of all and only the propositions one knows. This picture delivers the result that, for the principle at hand, in order for S to provide evidence that *p* to S₂, S₂ must come to know some proposition(s) that strongly supports *p* (or, alternatively, S must come to know *p*).

It is not straightforward how to capture the relevant phenomenon in a non-evidential framework. I expect that the problems that arise for the above views will in some way arise for non-evidential accounts since, on any plausible non-evidential view, evidence will still be in some sense relevant to one's overall theory. (In Section 3, I discuss an issue more clearly relevant to a non-evidential picture: evidential 'uptake'.)

Each of these options for understanding what it is to have evidence deliver the result that *provide-have* is too demanding. Each gloss requires some kind of cooperation on the part of the recipient – whether it be thinking about the evidence, allowing oneself to experience the seeming, or whether one must believe the relevant propositions. Sometimes recipients are uncooperative. If *having* evidence, in each of these senses,

13 Note that this suggests 'providing' is more similar to 'offering' than to 'giving.'

14 This general view is cashed out in numerous ways, but the nuances won't matter for my purposes. See Tucker (2013) for further discussion.

15 See Feldman (1988). Specifically, he discusses when S has *e* 'available' as evidence.

requires the cooperation of the recipient, then it will not be within the power of the provider to guarantee that the evidence is possessed. Thus, if to count as *providing* evidence, the recipient must *have* the evidence, it will not be in the power of the provider to ensure the evidence is provided. But on the assumption that ought implies can, combined with the observation that we can be under obligation to *provide evidence*, *provide-have* is false.

This conclusion fits with some of the cases already mentioned – both the epistemic and non-epistemic senses. There is a clear sense in which I provide dinner for you even if you refuse to eat it. Whether you eat the food provided is, in some relevant sense, outside the control of the provider. Examples abound: as I tell you some important information, you fall asleep. An instructor announces the exam date to the class, but Sammy isn't listening. The newspaper is dropped on your doorstep but a passerby steals it. (Or, the information is *on the syllabus*, but the students don't read it.) It's very natural to think that in these cases, the evidence has been provided even though the intended recipient doesn't *have* the evidence. The conclusion we should draw, I suggest, is that in general, providing evidence does not require that the recipient have the evidence: *provide-have* is false.¹⁶

It will be helpful to distinguish an *obligation* to provide evidence from the general phenomenon of providing evidence. The two can come apart in various ways. One may *discharge* one's obligation to provide evidence even in situations where the provision is unsuccessful. One may also provide evidence without discharging one's obligation. Consider the following scenario where a subject, Jones, is under obligation to provide evidence in a specific way – namely by testifying in court as a witness to an event. Suppose the day before he is to testify, Jones is speaking loudly about the event and is overheard by the jury. Even if by doing so he provides evidence to the jury (in the sense of probability raising), he still must testify in order to discharge his obligation to provide evidence.¹⁷ His obligation is met only when he provides evidence in a specific way: by testifying in court.¹⁸

Given that something strikes us as initially plausible about *provide-have*, perhaps we can capture the intuitive thought with an alternative principle. One way to revise the principle is to replace 'have' with a less demanding notion. Perhaps what is necessary for evidence to count as *provided* to S is for the evidence to be in some sense *in S's reach*.

Provide-reach: If S provides evidence that *p* to S₂, then evidence that *p* is within S₂'s reach.

This sounds plausible, but to elucidate we need to specify what it is for evidence to be *within the reach* of an agent. One option is that evidence *e* is within S's reach when S is in a position to believe *e* outright:

16 It is open to one to package things differently and include as part of the obligation to provide evidence that the recipient have the evidence. On this way of thinking, fulfillment of one's obligation will require that certain things in the world are in place. Consider a parallel case: I have an obligation to open the door when my friend knocks on the door, but suppose when he arrives the door is jammed and I can't get it open. If we think that *opening the door* is built into the obligation, the result is that sometimes we can't fulfill our obligations – sometimes the world doesn't cooperate. I find this a less natural package, but mention it as an option.

17 Evidence obtained in this way will not be admissible for the purposes of conviction, but the point remains: he has provided evidence to the jury but has not discharged his obligation to provide evidence in court.

18 Thanks to Baron Reed and Sarah Wright here.

Provide-position: If S provides evidence *e* that *p* to S₂, then S₂ is in a position to believe *e* outright.

This improves upon *provide-have* by not requiring the cooperation of the recipient. Still, it introduces the need for an account of what it is to be in a position to outright believe. (Another option is to replace ‘within reach’ with ‘in a position to know.’)

A further qualification we may want to add concerns whether a constraint on *providing evidence* is that the recipient must recognize the evidence as evidence for *p*.¹⁹ (Or, given the considerations just surveyed, whether the recipient must be *in a position to recognize* that *e* is evidence for *p*.) One might provide evidence *e* that *p*, and the recipient come to know *e* but have no idea of the connection *e* bears to *p*. Suppose, for example, that we are at a horse race and I want to provide you with evidence that Horse A will win. I tell you the name of Horse A’s trainer, who is in fact the best in the world, but it is evident that you know nothing about horse trainers. A natural expectation is that I explain the relevance of the trainer to Horse A’s performance in the race. If I fail to do so, do we want to say that I *provided evidence* to you that Horse A will win? There is a choice point here. Quite plausibly, in some contexts, it is an additional part of one’s obligation to make the connection clear.²⁰

Let’s turn our attention for a moment to an alternative principle, one that attempts to specify a sufficient condition for *providing evidence*:

Provide-Have-SUFF: S provides evidence *e* that *p* to S₂ if S₂ has evidence *e*.

Unlike the necessary direction of the principle, this version does not become overly demanding when we supply an epistemic gloss on *have*. Nevertheless, qualification is needed. *Provide-Have-SUFF* does not mention anything about the way in which S₂ comes to have the evidence that *p*. But intuitively, for S to successfully provide evidence that *p* to S₂, it needs to be the case that S₂ has the evidence in some part *due to* S. This points to the need for refinement along the following lines:

*Provide-Have-SUFF**: S provides evidence *e* that *p* to S₂ if S₂ has evidence *e* because of S.

Attempts to give an account of what it is for S₂ to have evidence *because of* S will involve many of the same difficulties that shroud the basing relation. I won’t try to solve those here.

Not surprisingly, a connection between *providing* and *having* is more promising when *having* is a sufficient condition rather than a necessary condition for *providing*. As a sufficient condition *provide-have-SUFF** looks plausible, as a necessary condition it is too demanding.²¹

19 The issue arises because, in fact, we are not ideal agents. If the agent were ideal, this additional consideration wouldn’t be necessary.

20 Relevant here is Sandy Goldberg’s (2017) helpful discussion of the nature of social expectations of this kind.

21 One consequence of what I’ve argued for here is that S can provide *e* to S₂ without S having *e* or S₂ having *e*. Although I agree that it sounds strange to say, ‘She provided him with evidence *e*, but neither she nor he have the evidence,’ I think this can be explained by the fact that typically, we make attributions of providing evidence only in cases where we also assume that the person providing the evidence has the evidence. Thanks to Matthew Benton here.

3. DEFEAT AND UPTAKE

One question we might ask in any given instance where evidence e for p is provided yet the recipient fails to have e or fails to believe p is: *why* does the recipient fail to believe p ? What explains the failure of S_2 to ‘uptake’ or take advantage of the evidence provided? In this section, I discuss various possible answers to this question. I look first at settings where the recipient has a defeater for p , then explore other factors that may interrupt the transmission process – complications concerning the ways agents ‘uptake’ or evaluate evidence. These include epistemic factors that contribute positively to one’s epistemic position – epistemic character, expertise, and skill – but also factors that may contribute in negative ways. I argue that there are cases where evidence is provided, the intended recipient does not have the evidence, but where there is no *culpable ignorance* on the part of the intended recipient.

3.1 Defeat

One type of setting where evidence for p is provided but the recipient fails to believe p is a setting where the recipient has a defeater for p . The recipient may have an undermining defeater: you yell ‘Shark!’ to warn of imminent danger, but the listener believes you to be playing a practical joke (perhaps she knows you are a prankster and often give false warnings at the beach). The recipient may also have an overriding defeater. Matthew Benton (2016) presents a case that nicely demonstrates this phenomenon. In the case, while he is at the pub with Blake, he receives a text from another friend stating that Blake is at the park. Since Matt knows this is false – Blake is sitting across from him – Matt doesn’t believe that Blake is at the park, but continues to believe Blake is at the pub.²² Did the friend provide evidence to Matt that *Blake is at the park*? It seems natural to say he did, even if Benton knows the evidence is misleading. (Note that it’s consistent with *strong provide* that the friend provided the evidence, but according to *strong provide-agent*, he did not.)

Consider another case. Suppose again that we are at a horse race and I tell you all sorts of facts about a particular horse, Horse A – facts which indicate that Horse A will win the race. You are an eager listener and understand the relevance of the facts (you recognize the name of the horse trainer, who is the best in the world, etc.). But you also know something I don’t know: that the horse is sick today, and moreover, has never won a race while sick. Although you don’t believe that *Horse A will win the race*, it’s very natural to think that I have provided you evidence that Horse A will win the race.²³

It’s worth noticing that given the way defeaters are explained in a Bayesian framework, some of what is discussed in this section is already accounted for. With respect to our topic of *providing evidence*, the general phenomenon of *defeat* can be captured as follows: S has a defeater for p when (strong) evidence e for p is provided to S , S has e ,²⁴ but S ’s total body of evidence – including e – does not support p .²⁵ I don’t intend to suggest that all aspects of the

22 Assuming an E=K framework, and that Matt continues to know that Blake is at the pub, Matt will not even change his credence that *Blake is at the park*. The same result holds in a non-E=K framework where Benton is certain that Blake is at the pub with him.

23 Thanks to John Hawthorne here.

24 By ‘has’ here I mean that e is added to S ’s body of evidence.

25 Note that this way of thinking about defeat – where we assume e is added – seems to implicitly take a stand on a controversial issue concerning whether evidence can be false. If it cannot, this will constrain e and it will be less straightforward to determine, in some cases, what e a subject receives. For example,

phenomenon of defeat are uncontroversial in a Bayesian framework – the framework in itself does not settle all disputes: it does not settle the issue of what evidence an agent receives in a given setting, for example. Still, even given this generalized gloss on defeat, there are a few things worth emphasizing with respect to how *defeat* and *providing evidence* interact.

Defeat can interrupt the uptake of evidence even in what we might otherwise consider to be paradigm cases of providing evidence. Consider this example.

Red Cup: I have a red cup in my hand, and clearly set it within your view. But you rationally believe that the room is fixed with trick lighting – lighting that makes objects that are not red look red – so you don't believe the cup is red.

Showing someone a red object is, in general, a paradigmatic case of providing someone with evidence that the object is red. But even given a cooperative recipient, defeat can interfere.

At this point we might wonder whether a requirement on providing evidence to someone is that when the provider is aware that the intended recipient has a defeater, the provider must do more than merely provide evidence that p , she must also defeat the relevant defeater. (Another way to capture the idea is this: the provider must sometimes destroy or remove knowledge or belief in the recipient, not merely add.) Although in some cases one may be obligated to do more than provide evidence, it's implausible that this will be part of what it is, in general, to provide evidence. Principles such as the following might initially seem attractive:

Provide-Defeat: If S knows that S_2 has a defeater for p , for S to provide evidence for p to S_2 , S must defeat S_2 's defeater for p .

But, in general, this kind of principle is implausibly demanding. Whether defeating a defeater is part of the agent's *obligation* to provide evidence will depend on the context. In the Shark case, for example, it's easy to imagine that I am criticizable if I know that you think I am playing a practical joke and I don't do more to convince you that I am not and that the shark is near. Still, if you are not cooperative there will be limits to what I can do to convince you. (Clearly, in some cases there may also be actions I ought to take if you are unconvinced – such as pulling you out of the water if a shark is near – but this goes beyond epistemic considerations.)

The difficulties with the preceding principle may be more obvious when put in Bayesian terms, where the principle will look like this:

Provide-Defeat:* If S knows S_2 has a defeater for p , for S to provide evidence that p to S_2 , S must bring it about that S_2 's total evidence strongly supports p .

One problem worth flagging: if *defeating defeaters* or *making S 's total evidence support p* requires getting S to have some additional evidence or believe some additional propositions other than p , many of the same problems will reappear with respect to the new evidence that must be provided. (So, again, these principles don't look promising.)

in the case above, this issue is relevant to whether Matt gets *Blake is at the park* as evidence, or whether he merely gets *S said 'Blake is at the park'* as evidence. These are disputed matters that make situations involving defeaters less straightforward than is sometimes appreciated. See Lasonen-Aarnio (2010) and Baker-Hytech and Benton (2015) for helpful discussions of defeat.

The preceding observations about the phenomenon of defeat are relevant to our evaluations of agents qua recipients of information. They suggest one way that an agent can *blamelessly* be ignorant of p even when evidence for p has been provided to her. In this way, they suggest the following principle is false:

Provide-Culpable: If evidence e for p is provided to S and S has e but doesn't believe p , then S is culpably ignorant of p .

Settings where the recipient has a defeater are straightforward kinds of cases where strong evidence for p can be provided, the recipient has the evidence, does not believe p , but the explanation does not involve epistemic blameworthiness on the part of the recipient. Together with the observation of the previous section – namely, that the provider can provide evidence (or discharge her obligation to provide evidence) even when the recipient does not have the evidence – we have identified two kinds of situations where evidence is provided, but the intended recipient either does not have e or has e but fails to believe p : situations where the recipient is uncooperative and situations where the recipient has a defeater. Settings involving defeat involve a relatively straightforward explanation for why the evidence provision fails to achieve its aim. But there are less straightforward cases. I'll consider these next.

3.2 *Evidential uptake*

Many factors influence whether an agent uptakes – or fails to uptake – available evidence. Some of these factors relate to what we might think of as an agent's *epistemic character*. In a Bayesian framework, rationality is a matter of conditionalizing on one's evidence. But the Bayesian framework is silent on matters relating to *how an agent gets the evidence she has*. Some epistemologists have argued that features such as an agent's method of inquiry are relevant to our epistemic evaluation of the agent.²⁶ According to these views, even if one's beliefs 'fit' or 'respect' one's evidence, one can be epistemically blameworthy for having the evidence one has – and importantly, for failing to have evidence one *should have had*. Pictures friendly to inclusion of this broad kind of epistemic evaluation are sometimes called *responsibilist* theories. The extent to which *responsibilist* views are compatible with evidentialism is a contested issue.²⁷ I won't take sides on this issue here.²⁸

26 For relevant discussion see: Foley and Fumerton (1982), Heil (1983), Kornblith (1983), DeRose (2000), Feldman (2000), Cloos (2015), Goldberg (2016, 2017), among others.

27 A persistent difficulty concerns how we ought to evaluate other beliefs that would differ in the counterfactual scenario. Suppose I ought to have read a book, but did not. If I had read the book, I would have evidence that I don't have – perhaps I would have learned the contents of the book. It's natural to think that this is evidence I ought to have, or that I should've known. Some go further to argue that such evidence can function as a *normative defeater* for some of my beliefs. (See especially Goldberg 2017. And for discussion of whether evidence we don't possess can have defeating power even when it's not the case that we ought to have the evidence, see Harman 1973 and Lycan 1977). But there are other propositions I don't currently believe but would've believed had I read the book, such as *that I read the book*. Am I currently justified in believing that I didn't read the book? Or is this belief defeated as well – since, if I had read it, as I ought to have, then I would believe that I read it. Perhaps we can hope for a disciplined characterization of *normative defeat*, but until one is available the best way to incorporate the responsibilist idea may be to keep its evaluation distinct from the evidentialist evaluation.

28 A related issue concerns whether a *responsibilist* evaluation is fundamentally *epistemic* or if it falls more squarely in the domain of practical rationality or ethical evaluation. Even if *responsibilist* factors are irrelevant to an agent's rationality with respect to a particular proposition at a time, they might be

Of particular interest are cases where two agents uptake different evidence in the same situation (that is, where the same evidence is *provided* or *available* to each), and where the difference in uptake is not explained by a difference in the agent's respective background evidence. While a Bayesian picture is helpful in explaining why agents don't believe p when they uptake evidence e in cases of defeat, they have little to offer by way of explanation for the kinds of cases discussed in this section – cases where the agent does not even uptake e .

Among the factors relevant to differences in how agents uptake evidence is an agent's expertise, or lack thereof. Agents with expertise of a certain kind will, in some situations, be able to obtain evidence that the non-expert is unable to obtain even when both are placed in the same setting. Sometimes the reason for this will be that the expert can *see* more than the average person or novice: she has a perceptual expertise or skill. For example, someone trained to identify gold will be able to see that a pebble is gold when an untrained person would not. (Susanna Siegel suggests that expert top-down 'seeing' is a kind of epistemically beneficial cognitive penetration.)²⁹ In some cases, expertise may ultimately be reduced to a difference in background information. But, plausibly, not all kinds of expertise can be explained in this way. Consider variations in sensory faculties, such as eyesight. Some people have better eyesight than others. People with perfect eyesight will be able to uptake evidence from objects at a distance that agents with poor eyesight (and without glasses on) are unable to uptake.

The relevance for *providing evidence* is that even in the absence of defeaters, some agents may not be able to uptake evidence provided. An extreme case is that of a blind man who cannot see the red mug you place in front of him. He lacks the requisite ability to acquire evidence via sight. Should we count you as having provided him with evidence that your mug is red if you hold it in front of him? This raises a question about how individualized *providing evidence* is. Usually, holding up your bright red mug suffices to provide someone in front of you with evidence that your mug is red. But this case suggests a more individualized approach. If you know the person is blind (and perhaps even if you do not know) it seems a stretch to say that you provided him with evidence that the mug in your hand is red when you merely hold it out in front of him. (If you do not know the person is blind, this may be a case where you discharge your obligation to provide evidence, despite not providing it.)

Let's consider a less extreme case. Suppose you are giving a philosophy talk to an audience with varied experience and background in the discipline (professors, graduate students, and undergraduates). From your perspective, you provide the same evidence to them all – at least, everyone in the room hears the same words, has the same handout, sees the same slides, and so on. It is reasonable to expect that when each leaves the talk, each will have different evidence. It's not only that they leave with a different set of credences than one another – this is true but could be explained by a difference in background evidence. Rather, the set of propositions each uptakes will be different. Some of this could be explained due to factors such as attention – perhaps one person had coffee right before the talk and so is highly focused, and another was up all night and so falls

relevant to other ways in which we evaluate an agent epistemically, such as the agent's epistemic habits. Insofar as the features are relevant to the agent's aim of believing the truth, the evaluation strikes me as relevantly epistemic. This kind of evaluation will, for instance, be more clearly epistemic than an evaluation of the agent's choice between vanilla or chocolate ice cream.

²⁹ See Siegel (2012). Siegel (2017) is also relevant.

asleep. Others' uptake might be influenced by factors that have their source in poor character: perhaps they are lazy or dislike the speaker and so are predisposed to dismiss anything the speaker argues. But these factors aside, even supposing equally attentive and gracious listeners some may uptake different evidence because of differences in the skill set each possesses: an ability to 'see' connections or to pick up on unsaid implications. Some might need an hour to think through the ideas, while others see the conclusion immediately.³⁰

As the latter case makes evident, not every situation of 'uptake failure' will involve the absence of a relevant ability. Sometimes an agent won't uptake *e* due to practical reasons: there is a bird in the window to my left but I'm focused on my book; I don't turn my head and thus fail to see it. There's a sense in which I'm 'in a position to know' that there is a bird outside my window, but this failure to uptake the evidence is not relevant to the epistemic evaluation of my actual beliefs. It's difficult to tease apart the practical considerations relevant here from the epistemic. We can appeal to clear cases: if I have to leave the lecture for a family emergency, this is a practical consideration. If I like the speaker's shoes and so believe everything they say without scrutiny, this is epistemically irresponsible. I'm skeptical that we can draw clean lines of demarcation.³¹ But even in the absence of strict divisions, we can learn something general from this class of phenomena.

The main point is this: there are cases where evidence *e* for *p* is provided to *S*, *S* lacks defeaters, but *S* fails to uptake *e* and does so non-culpably (both in an epistemic sense of culpable and a non-epistemic sense). Moreover, *S* may non-culpably fail to uptake *e* even when someone else with the same evidence as *S* and in the same situation does uptake *e*. (We can add that it's not due to a past or present failure that *S* lacks the relevant ability or fails to attend to the evidence.)³² Thus, we have reason to reject the following principle:

Culpable Uptake Failure: If evidence *e* for *p* is provided to *S* and *S* does not have *e* and does not believe *p*, then *S* is culpably ignorant of *p*.

In conclusion, it is insightful to distinguish two different types of interruption that can affect the transfer of information when providing evidence. In cases of defeat the problem is that although *e* raises the probability of *p*, when *e* is embedded in *S*'s body of evidence, *p* is not probable enough for outright belief. In such cases the evidence is provided and the recipient has it, but the recipient doesn't thereby believe *p* (and is rational in so doing). A different kind of problem occurs when the evidence is provided, but the recipient does not 'uptake' the evidence. 'Uptake failure' can have many sources. I've argued that these

30 Here, again, we drop the idealization since we are not ideal rational agents.

31 One question we can ask concerns whether each and every listener is equally rational. As long as each updates correctly on the evidence each *in fact* gets, the Bayesian framework will say they all are rational. A Bayesian evaluation stops there. The responsibilist allows for an additional evaluation (we might criticize someone for falling asleep, engaging in wishful thinking, and so on.) This kind of evaluation will have to be taken on a case by case basis.

32 There will be limits to when the lack of ability can be offered as an explanation for uptake failure that excuses both the provider and intended recipient. Suppose I write you a letter in Chinese and you don't know Chinese. You could pay to have it translated, or learn Chinese yourself, but in this case the burden seems to fall on the provider to make the evidence more easily available.

interruptions can occur even when both the provider and the intended recipient are non-culpable. ³³

4. DIVINE HIDDENNESS: AN APPLICATION

The notion of *providing evidence* plays a prominent role in the literature concerning the problem of divine hiddenness. The central line of argument in that literature involves the following thought: a loving God – that is, a God that wants a relationship with us – would provide us with evidence of his existence.³⁴ Or, in alternative terminology, if there is a loving God, he would let us know that he exists. The fact that he hasn't – or, the fact that our total evidence doesn't obviously make God's existence probable, is itself evidence that such a God does not exist. This poverty of our evidential situation, it is argued, itself constitutes evidence against the existence of God. In light of what I've said here, in this section I will make a few applications to this topic.

First, it might initially seem plausible that if God provides evidence of his existence, then everyone who is paying attention or looking for it *has* the evidence. But, as I've argued, we can make no such inference. The lesson is that one cannot argue from the fact that some people who are open to evidence for the existence of God appear not to have it, to the conclusion that the evidence has not been provided.

A second consequence that falls out of my discussion is that we cannot make an inference from 'evidence for *p* was provided and S doesn't *have* the evidence or doesn't believe *p*,' to 'it's S's fault' – that is, we cannot conclude that S is culpably ignorant. The evidence could be provided or available, but S might not *have* the evidence – or S might *have* the evidence but not believe *p*, due to defeaters. So while some theists respond to the hiddenness argument by arguing that all atheists or agnostics must, given the widespread evidence God has provided, be culpable for not believing, we have seen several general reasons why those who do not have evidence provided could be non-culpable in failing to have the evidence or in failing to believe that God exists.

In closing, I'd like to make one further application, drawing on the idea of epistemic expertise. The idea that evidence is *provided* in the sense of an available input that requires expertise to uptake, improves upon an idea already in the literature that the uptake requires some cooperation on the part of the recipient. Some have found plausible an explanation of the problem in terms of an agent's 'resistance' to the evidence (where such resistance is pictured as agents clenching their eyes closed). The idea of expertise – while compatible with the idea that some who lack evidence are resistant – allows for another explanation. Some people might need time – either to evaluate the evidence or to develop the requisite expertise. These agents are not culpable for their ignorance – they 'respect their evidence' in the way a rational Bayesian is required. They simply do not have the skill or requisite expertise to fully appreciate the evidence.

There are non-epistemic features of a person's life that at least partially explain why some people end up with the set of evidence they have, and why they direct their attention in certain ways, and why they engage in the methods of belief formation they engage in.

³³ Relevant here is Kelly's (2008) defense of rational belief polarization.

³⁴ See Anderson (2017) and Howard-Snyder and Moser (2002) for further discussion. For the sake of space, I gloss over many details here.

All of this suggests that there are a wide number of explanations available for why some people's total evidence seems not to point to a God.³⁵

REFERENCES

- Anderson, C. 2017. 'Divine Hiddenness: Defeated Evidence.' *Royal Institute of Philosophy Supplement*, 81: 119–32.
- Baker-Hyitch, M. and Benton, M.A. 2015. 'Defeatism Defeated.' *Philosophical Perspectives*, 29: 40–66.
- Benton, M. 2016. 'Knowledge And Evidence You Should Have Had.' *Episteme*, 13: 471–9.
- Cloos, C. 2015. 'Responsibilist Evidentialism.' *Philosophical Studies*, 172: 2999–3016.
- DeRose, K. 2000. 'Ought We to Follow Our Evidence?' *Philosophy and Phenomenological Research*, 60: 697–706.
- Feldman, R. 1988. 'Having Evidence.' In D. Austin (ed.), *Philosophical Analysis*, pp. 83–104. Dordrecht: Kluwer.
- 2000. 'The Ethics of Belief.' *Philosophy and Phenomenological Research*, 60: 667–95.
- Foley, R. and Fumerton, R. 1982. 'Epistemic Indolence.' *Mind*, 91: 38–56.
- Goldberg, S. 2016. 'On The Epistemic Significance Of Evidence You Should Have Had.' *Episteme*, 13: 449–70.
- 2017. 'Should Have Known.' *Synthese*, 194: 2863–94.
- Harman, G. 1973. *Thought*. Princeton, NJ: Princeton University Press.
- Heil, J. 1983. 'Doxastic Agency.' *Philosophical Studies*, 43: 355–64.
- Howard-Snyder, D. and Moser, P. K. 2002. *Divine Hiddenness: New Essays*. Cambridge: Cambridge University Press.
- Kelly, T. 2008. 'Disagreement, Dogmatism, and Belief Polarization.' *Journal of Philosophy*, 105: 611–33.
- Kornblith, H. 1983. 'Justified Belief and Epistemically Responsible Action.' *Philosophical Review*, 92: 33–48.
- Lackey, J. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. 2010. 'Unreasonable Knowledge.' *Philosophical Perspectives*, 24: 1–21.
- Lycan, W. 1977. 'Evidence One Does Not Possess.' *Australasian Journal of Philosophy*, 55: 114–26.
- Siegel, S. 2012. 'Cognitive Penetrability and Perceptual Justification.' *Noûs*, 46: 201–22.
- 2017. 'How Is Wishful Seeing Like Wishful Thinking?' *Philosophy and Phenomenological Research*, 95: 408–35.
- Tucker C., ed. 2013. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford: Oxford University Press.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

CHARITY ANDERSON is an Assistant Professor of Philosophy at Baylor University. Her research spans Epistemology and Philosophy of Religion.

35 Thanks to Matthew Benton, Brandon Fitelson, Ram Neta, Trevor Nyman, and Sarah Wright for comments on a draft of this paper and also to the participants of the 13th annual *Episteme* conference. I'm particularly grateful to John Hawthorne for discussion of the issues in this paper.