

Belief Dependence: How Do the Numbers Count?

Zach Barnett

Penultimate Draft

Forthcoming, *Philosophical Studies*

Abstract: This paper is about how to aggregate outside opinion. If two experts are on one side of an issue, while three experts are on the other side, what should a non-expert believe? Certainly, the non-expert should take into account more than just the numbers. But which other factors are relevant, and why? According to the view developed here, one important factor is whether the experts should have been expected, in advance, to reach the same conclusion. When the agreement of two (or of twenty) thinkers can be predicted with certainty in advance, their shared belief is worth only as much as one of their beliefs would be worth alone. This expectational model of belief dependence can be applied whether we think in terms of credences or in terms of all-or-nothing beliefs.

1 Introduction

While taking a logic exam, you encounter a problem you don't know how to do. So you decide to cheat. Your friends Anna, Beth, and Chad are seated close by, and you know that they are all quite good at logic – about equally good, in fact. You peek at their answers and find that they did not all agree: Anna answered *not-p*; Beth and Chad answered *p*. What do you do?

Other things the same, it seems clear that you should go with *p*. You have two reliable sources against one. But suppose that you saw Chad copy off of Beth. Given this information, it seems clear that you do not have reason to favor *p*. Even though the case can still be described as 'two against one,' Chad's opinion is *dependent* on Beth's, in some important sense. And, for this reason, it seems to lack epistemic significance (that is, it seems not to provide additional support for *p* beyond that provided by Beth's opinion). To accommodate cases of this kind, we might offer the following general principle: When one opinion is totally dependent on another, the dependent opinion does not provide any additional support for the jointly held proposition.

It is difficult, however, to characterize precisely the type of belief dependence that can play this epistemic role. Lackey (2013) canvasses various attempts to capture the relevant type of dependence, arguing that all are unsuccessful. Ultimately, she goes a step farther, boldly rejecting the received view that completely dependent opinions lack epistemic significance. This paper has two goals. The first is to defend the received view from Lackey's thought-provoking critique. The second is to offer an account of belief dependence – that is, an account of when additional opinions count and when they do not. The account offered is expectational: According to it, one opinion is (completely)

dependent on another when learning about the latter allows one to predict the former (with certainty). Later, we will see that this expectational account can be extended to cover partial dependence and independence and that it can be applied regardless of whether we understand a person's opinions in terms all-or-nothing beliefs or in terms of credences. Before exploring such subtleties, though, we must say a bit more to motivate investigation into this issue.

2 The Importance of Belief Dependence

One reason to care about this issue is its connection to disagreements involving shared evidence – a topic that has received substantial attention. Much of that controversy turns on just how we should see the interaction between ordinary evidence (e.g. the information given in the logic problem) and the evidence provided by the opinions of others. On the one hand, 'conciliatory' views demand that a person revise her opinion when she encounters disagreement from a reliable source who shares her evidence – even if she had in fact responded perfectly to the shared evidence in the first place. On the other hand, 'steadfast' views hold that those who actually respond correctly to the shared evidence should not revise their beliefs in light of disagreement.

Given this distinction between views, it is tempting to think that the interest of the belief dependence issue is limited to conciliationists. But this is a mistake. Virtually everyone, whether conciliationist or not, should allow that outside opinions can have some evidential force, when such opinions are the only evidence a person has about some proposition. (Everyone should allow, for example, that a reliable meteorologist's predictions, other things equal, provides some evidence about whether it will rain.) As long as one admits that outside opinions can provide evidence sometimes, one will need to be able to handle questions involving belief dependence. We can ask, for example: How much more evidence is provided by the agreement of independent experts than is provided by that of dependent experts? So questions surrounding dependence are not just of interest to conciliationists; they should be taken seriously more generally.

In order to focus on the issue directly, we will concentrate on cases that allow us to stay neutral on how to adjudicate the disagreement debate. So, for instance, in the logic exam case, we will not ask: What should Anna think, in light of the information given in the logic problem and the disagreement from Beth and Chad? Answering this question would require us first to decide what Anna should think, in light of the logic problem information and disagreement from Beth alone, which is a controversial matter. Instead, we will ask: What should an outsider think, given Anna's, Beth's, and Chad's opinions, and importantly setting aside any evidence provided by the logic problem itself?¹

¹ In the original story, we imagined that you had seen the logic problem but were unsure how to solve it. In the cleanest version of the case, we'd stipulate that you had not even seen the logic

3 Lackey's Argument

In the logic exam case, it seems clear that Chad's opinion should not count for much, if it really is true that he simply copied off of Beth. On the basis of this judgment, it seems reasonable to posit the following general principle:

Belief Dependence: When one opinion is totally dependent on another, the dependent opinion does not confer any additional support for the jointly held proposition.

Precedent for such a principle is easy to find. Here is Elga:

[A]n additional outside opinion should move one only to the extent that one counts it as independent from opinions one has already taken into account. (2010, p. 177)

Elga regards this claim as “completely uncontroversial” and suggests that “every sensible view on disagreement should accommodate it” (2010, p. 178). Kelly, writing from the other side of the disagreement debate, shares Elga's outlook:

[E]ven in cases in which opinion is sharply divided among a large number of generally reliable individuals, it would be a mistake to be impressed by the sheer numbers on both sides of the issue. For numbers mean little in the absence of independence. (2010, p. 148)

But despite the widespread appeal of Belief Dependence, Lackey (2013) argues, persuasively, that it is tricky to characterize the type of dependence operative in this principle. Indeed, she goes a step further, suggesting that there is simply no good way to do it. Let's take a look at Lackey's argument.

Lackey restricts her attention to cases involving *epistemic peers* (who, for Lackey, are “evidential and cognitive equals” with respect to the issue at hand²). On the face of it, it may seem strange to invoke peerhood here. After all, Belief Dependence states simply that if one person's belief is dependent on another person's, then the dependent belief does not confer additional support for the opinion shared. Peerhood seems irrelevant to the issue.

Though Lackey does not explain why this restriction is imposed, there is a compelling rationale for it. In assessing the import of incoming opinions, it is important to distinguish two questions: (1) How strong are the respective epistemic credentials of the sources of these opinions? (2) To what extent do these sources depend on each other in their thinking? Since, presumably, the relevant sort of dependence can occur when the involved people are on equal epistemic footing, it seems better, methodologically, to focus on cases of this type. Framing the question in terms of epistemic peers allows us to

problem but were still trying to determine its solution through the efforts of your classmates.

² See Lackey (pp. 243-245). We'll discuss the role this assumption plays in the following section.

control for a confounding variable. So, following Lackey, let us focus on a more restricted version of Belief Dependence:

Belief Dependence for Peers: When a person's opinion is totally dependent on the opinion of an epistemic peer, the dependent opinion does not confer any additional support for the jointly held proposition.

Lackey considers several ways one might try to understand this notion of dependence so as to render the principle true.³ But she argues that each is no good, for reasons we will see shortly. Ultimately, Lackey argues that this widely held principle cannot be sustained, suggesting instead that dependent beliefs do tend to confer additional support:

I shall show that where one disagrees with two (or more) epistemic peers, the beliefs of those peers can be dependent in the relevant sense and yet one cannot rationally regard this as a single instance of disagreement when engaging in doxastic revision. (2013, p. 245)

It would seem that Lackey is flat out rejecting the intuition elicited by the logic exam case – provided that peerhood between Beth and Chad is stipulated. With peerhood in place, Lackey's view seems to entail that, contra appearances, Chad's opinion, together with Beth's, somehow counts for more than Beth's opinion does alone. And if we add that Anna, too, is a peer of Beth and of Chad, then it seems to follow that you, the cheater, would have reason to favor Beth and Chad's joint answer (*not-p*) over Anna's answer (*p*) – even if you were certain that Chad got his answer directly from Beth.

Lackey offers an intriguing diagnosis of this result. She points out that the case is underdescribed. Though we know that Chad's opinion was, in some sense, grounded in Beth's, we are not told whether Chad was at all critical in his decision to endorse Beth's opinion. Here, Lackey distinguishes what she calls *autonomous* and *non-autonomous* dependence:

The autonomous version of this dependence involves a subject exercising agency in her reliance on a source of information, critically assessing its reliability, monitoring for defeaters,

³ It is worth noting that Lackey's version of the principle is slightly different:

When A disagrees with peers B, C, and so on, with respect to a given question and A has already rationally taken into account the disagreement with B, A's disagreement with C, and so on, requires doxastic revision for A only if the beliefs of C, and so on, are independent of B's belief. (2013, p. 244)

This version is more closely intertwined with the issue of disagreements involving shared evidence. When A disagrees with peers B, C, and so on, with respect to a given question and A has already rationally taken into account the disagreement with B, A's disagreement with C, and so on, requires doxastic revision for A only if the beliefs of C, and so on, are independent of B's belief. (2013, p. 244)

This version is more closely intertwined with the issue of disagreements involving shared evidence, for A must weigh the import of incoming opinions while maintaining her own point of view. In this paper, we set this complication to the side. Lackey's arguments apply equally to both versions of the principle.

and comparing the content of the belief that she forms with her background beliefs. This, I take it, is the minimum required for rational belief formation. (2013, p. 249)

Applying this distinction to the case at hand presents two possibilities: Either Chad was autonomous in his reliance on Beth or he was not. Whichever way we go, Lackey thinks, we will not need to endorse anything like Belief Dependence for Peers to deliver the correct verdict.

First, suppose Chad was autonomous in his decision to copy Beth. So we can presume either that Chad engaged in some double-checking of Beth's answer or, at the very least, that he thought about whether Beth was a reliable source, prior to copying her. Consider each option in turn.

Double checking: If Chad engaged in a bit of double-checking before endorsing Beth's answer of *not-p*, then it seems plausible that his agreement does confer at least some additional support upon the answer they both favor. After all, Anna's answer was not double-checked, and it seems clear that a double-checked answer is a better bet than an un-double-checked one, from an outside point of view. So Chad's opinion must be providing some support of its own.

Copying a vetted source: What if Chad did not double-check Beth's answer, but did at least confirm Beth's reliability before resolving to trust her? Here, too, we can make it plausible that Chad's agreement should carry some epistemic weight. To see this, imagine, realistically, that your own reliability assessments of Anna and of Beth are less than certain: You have good reason to regard each as reliable, but you recognize that these assessments may be off base. Under these conditions, learning that Chad agreed with Beth is evidence that Chad assessed Beth's reliability favorably – which does seem to render their shared opinion at least slightly more credible than Anna's opinion. After all, we now have more evidence for Beth's reliability than we do for Anna's.

In either case, we find that – so long as Chad's reliance on Beth was autonomous – Chad's apparently dependent opinion seems still to have some epistemic significance.⁴ But what if Chad's reliance on Beth was not autonomous? What if, to use Lackey's term, Chad simply *parroted* Beth? Here, Lackey agrees that Chad's opinion does not provide additional support for the position he and Beth share. But Lackey notes that we do not need to appeal to Belief Dependence for Peers to explain this. Since Chad is non-autonomous in his reliance on Beth, he would defer to her even if she were thoroughly unreliable; he would adopt her beliefs even if they were patently false. On this issue,

⁴ One could object that, by making Chad's reliance autonomous, we have rendered his opinion at least partially *independent* of Beth's. I am sympathetic to this point of view; section 6 discusses an account of belief dependence that can deliver this result. However, I still see an intuitive sense in which Chad's opinions still are dependent on Beth's (e.g. causally), and, in this sense, Lackey's verdict seems to be exactly right.

Chad's belief-forming process is manifestly irrational. And, for Lackey, this makes it the case that Chad is not a peer of Anna or Beth – in which case, the situation is irrelevant to the principle at issue (2013, p. 253). Lackey also considers what would happen if we stipulate that Anna and Beth are just as irrational as Chad. In such a case, she points out that *none* of their opinions would have much epistemic significance at all, for reasons that have little to do with belief dependence. The key point is that – no matter how we describe the case – we do not need to invoke anything like Belief Dependence for Peers to explain why Chad's parroted opinion lacks epistemic significance.

The logic exam case seemed to illustrate the need for some kind of Belief Dependence principle. Taking into account Lackey's insights, it is not clear that such a principle is needed to accommodate this case. More generally, it is tempting to think that we can get by without ever appealing to Belief Dependence – at least in cases involving epistemic peers.⁵ The next section discusses some reasons for thinking that we do still need a principle closely resembling Belief Dependence while respecting Lackey's observations.

4 Rationality, Accuracy, and Dependence

There is something attractive about the argument discussed in the previous section. In many cases that seem, at first, to exhibit a problematic sort of belief dependence, the distinction between autonomous and non-autonomous dependence is an important one to draw. But the argument deserves a bit more scrutiny. In particular, it is important to note a way in which the argument depends on a particular conception of epistemic peerhood.

Christensen (2014) distinguishes two ways to think about epistemic peerhood: Two people are said to be **rationality peers**, with respect to some issue, if and only if they are equally likely to form *rational* beliefs about that issue; In contrast, two people are said to be **accuracy peers**, with respect to some issue, if and only if they are equally likely to form *accurate* beliefs about that issue (2014, p. 586).

Of course, these two notions will quite often coincide, since, presumably, rational beliefs tend to be accurate. Rationality peers tend to be accuracy peers, and vice versa. But the two notions can come apart. Compare two math students – one who solves certain problems properly and correctly, and another who regularly commits a pair of errors that cancel each other out. The two students might well be accuracy peers (scoring

⁵ One way to react to Lackey's argument is to retreat to a weaker thesis: To the extent that one opinion is dependent on another, the dependent opinion provides *less* additional support for the jointly held proposition. Lackey rejects even this more cautious formulation of the position, suggesting that "the amount of doxastic revision required in the face of disagreement does not track the amount of independence possessed by the target belief" (p. 265). We set this issue aside, as, in the end, the paper will aim to salvage a position resembling the original, unweakened, Belief Dependence.

equally well on exams, perhaps) without being rationality peers (since one of them makes so many mistakes).

Let us apply this distinction to one step of Lackey's argument. Return to the case in which Chad copies Beth non-autonomously. That is, he copies her opinion blindly – without any regard to her reliability and without any regard to the plausibility of the belief copied. Further, let's suppose that Beth, unlike Chad, is quite rational and reliable in her thinking about the relevant sort of issues. Given these details, the intuition elicited is clear: Chad's "parroted" opinion adds nothing to Beth's.

Lackey's answer to this case was to point out that Chad and Beth are not epistemic peers (rendering the case irrelevant to the principle under consideration). If we are thinking about rationality peerhood, then Lackey's answer seems just right: Chad and Beth are not equally likely to form rational beliefs about the relevant issue. But, if instead we are thinking about accuracy peerhood, then we will have to say that Chad and Beth are epistemic peers, with respect to the relevant issue, since the two are bound to be equally accurate, given the setup. So the argument against Belief Dependence is unsuccessful, if an accuracy-based notion of peerhood is assumed.

Importantly, this observation isn't a major objection to Lackey's argument: Lackey goes in for a rationality-based notion of peerhood explicitly (2013, p. 243). But since some of the targets of her critique – such as Elga (2010) – do favor accuracy-based notions of peerhood, it is worth observing how both varieties of peerhood interact with belief dependence.

The first lesson, then, is that whether a Belief Dependence principle is required to explain 'parrot cases' will depend on which type of epistemic peerhood we employ. If we think in terms of rationality, then we may be able to explain why parroted opinions lack epistemic significance without appealing to any Belief Dependence principle. But if we think terms of accuracy, then such a principle does seem necessary to deliver the correct verdicts in such cases.

At this point, one might wonder: Are parrot cases the only cases that call for (or at least allow for) some kind of appeal to Belief Dependence in order to derive the intuitive verdict? If this is so, we may be able to get by without ever invoking this principle, simply by insisting on rationality-based measures of epistemic credibility. But as it turns out, there are cases that demand some kind of appeal to Belief Dependence which do not require us to think in terms of accuracy to appreciate their force.

5 The Indispensability of Belief Dependence

Consider the following (highly unrealistic) case.

Chicken-Sexing: A *chicken-sexing heuristic* is a reliable, but fallible method that can be used to discern the sex of a chicken by examining a certain superficial fact about how it looks or

moves.

Dawn knows a heuristic – method A – that uses the chicken’s head movements as a guide. Millions of other people know a different heuristic – method B – that uses the chicken’s strut as a guide. Everyone has equal evidence for the efficacy of their respective method. As it happens, both method A and method B are 90% reliable at determining a given chicken’s sex.

A chicken walks by. Dawn, using her method, judges it to be female. Everyone else, using the other method, looks at the same chicken and judges it to be male. You’re staring at the chicken and have no idea which sex it is. You do not know how to apply either method, but you do know all of the above information. How confident are you that the chicken is male?

Not 99.9999% confident, presumably. Despite that you have millions on one side and only one person on the other, it seems clear that the chicken could quite easily be male or female. Indeed, if we idealize the case so that you are certain both heuristics can never be misapplied, then we can make it reasonable for you to afford equal confidence to the female and male hypotheses.

Here is an argument for this result. Given the assumptions made, it is certain that each heuristic was correctly applied. In this case, they produced divergent judgments. It follows that one of the heuristics gets this chicken wrong. Presumably, there are some chickens that are misclassified by method A but not by method B while there are other chickens that are misclassified by method B but not by method A. There must be about as many chickens in each of these groups – otherwise one of the two methods would be more reliable than the other. Absent any reason to suspect that the mystery chicken was pulled from one of these groups, it is reasonable to split one’s confidence equally between both options.

If we trust this result, it is worth exploring how it bears on the issue of belief dependence. For it is clear that all of the chicken-sexers in the story are epistemic peers with respect to chicken-sexing – both accuracy-wise and rationality-wise. If it were really true that numbers mattered, even in the absence of independence, then we would be forced to conclude that the male hypothesis was somehow more credible in the imagined situation. Since this is not a very plausible result, it seems that we will want to hold onto some version of Belief Dependence to explain what is going on in a case like this.

An opponent of Belief Dependence for Peers can push back against this case, though. The setup suggests that Dawn has some evidence for method A, while everyone else has evidence for method B. Doesn’t this imply that they have different evidence? And if so, wouldn’t this undermine the suggestion that the case is relevant to the principle in question (since the involved parties are not all epistemic peers)?

In response, there are two points worth making. First, it should be noted that one may be able to avoid this objection by revising the case. Suppose that all of the chicken-sexers have access to both heuristics, but, for whatever reason, Dawn uses method A,

while the others all use method B. So imagined, Dawn and her counterparts may well be peers, despite having used different methods on this occasion.

But leave this point aside. Even if we leave Dawn out of the story altogether, there still seems to be a clear need for some kind of Belief Dependence principle. For, compare two situations: in the first, we learn that millions of chicken-sexers (using method B) all judged the chicken to be male; in the second, we learn that a single chicken-sexer (using method B) judged the chicken to be male. If we again idealize so that method B can never be misapplied, is there any more support for the male hypothesis in the first situation? Clearly not. In both situations, it seems reasonable to have a confidence of .9 that the chicken is male. But this can be true only if the additional agreeing opinions confer no additional support.

No matter how you slice it, we need to be able to make sense of a certain sort of belief dependence – one that can render additional dependent opinions epistemically inert. The next section investigates the nature of this dependence.

6 Belief Dependence⁶

6.1 An Expectational Account of Complete Dependence

Dependent beliefs do sometimes lack epistemic weight. In the chicken-sexing example, Dawn's opinion is worth just as much as is the shared judgment of the many method B users. But an important question remains: *In what sense* are the opinions of the method B users really dependent? After all, their opinions are not necessarily *causally* dependent: These chicken-sexers may well have been causally isolated from one another, perhaps all discovering method B separately.⁷ Even if this condition is stipulated, the evidential import of their shared opinion does not change. So long as we know, in advance, that they are using the same method (and that the method cannot be misapplied), it seems to follow that their shared opinion should 'count as one.'

If causal dependence is not what matters in these cases, where should we look instead? Here is one angle. In determining whether two sources are dependent in the relevant sense, what matters is not whether one causes the other, but rather, *whether they*

⁶ The expectational account of dependence and independence discussed in the following sections are in broad agreement with remarks made by Jeffrey (1987, p. 392), Goldman (2001, pp. 99-101), and Elga (2010, p. 177). The task at hand is to develop a complete account of dependence, independence, and everything in between – along the way exposing a number of subtleties not discussed by these other authors.

⁷ In light of this observation, one might worry that the discussion in the previous section is unfair to Lackey's position. But Lackey does not only want to reject causally-based dependence principles – for example, she examines and rejects Goldman's account of dependence, which does not cast dependence in causal terms (pp. 257-260). And apart from questions about the details of Lackey's view, it is important to emphasize the non-causal character of the type of belief dependence that can serve to render agreeing opinions evidentially inert.

should be expected, in advance, to reach the same conclusion.⁸ When Chad copies Beth's answer uncritically, we can see in advance that the two students will come away with the same opinion. In the chicken-sexing example, too, we can see in advance that all of the method B users will issue the same judgment about the sex of the mystery chicken. The best way to capture the relevant sort of dependence should, I think, appeal to this observation. With this in mind, consider the following account.

Complete Dependence: C's belief (about some proposition *p*) is *completely dependent* on B's belief (about *p*) just in case, in advance,⁹ it is rational to be certain that C would arrive at his belief (about *p*) given that B arrived at her belief (about *p*). When this condition is met, the two beliefs will together provide just as much evidence for or against *p* as would be provided by B's belief, on its own.

It is worth pausing to clarify two key aspects of this view before applying it to the familiar cases we have examined. First, this account of complete dependence is *expectational*. It makes essential reference to what it would be rational for someone to predict, given certain information. So, on this view, whether a given belief is completely dependent on another will turn on what information an evaluator possesses.

Second, it is worth pointing out that there are different ways in which one belief can be dependent on another. Most commonly, complete dependence can occur when it can be seen in advance that two sources will inevitably *agree* about some proposition (e.g. when Chad blindly copies Beth). We can call this *ordinary* complete dependence. But, notably, complete dependence can also occur when it can be seen in advance that two sources will inevitably *disagree* about some proposition (e.g. when Chad the contrarian blindly negates all of Beth's answers, perhaps – after all, knowing Beth's opinion would still enable us to predict Chad's opinion with certainty here). We can call this *perverse* complete dependence, since it is impossible for two sources to exhibit this kind of dependence and still both be reliable (that is, to answer correctly more than half the time). Though certain forms of perverse dependence will arise later on, ordinary complete dependence proves more relevant to the cases we have so far examined.

Chicken-Sexing: Recall the chicken-sexing example. In assessing whether the joint opinion of the method B users should count for more than one of their opinions alone, we must ask: In advance, how likely was it that the method B users would all agree?

⁸ Vavova (2014) makes a similar observation. There, Vavova argues that, other things equal, a dissenting opinion is significant to the extent that it is surprising.

⁹ It is worth clarifying what is meant by the qualifier "in advance". The purpose of that provision is to block the following kind of reasoning: "Yes, I can predict Chad's opinion with certainty, because I already know for a fact that he ultimately came to believe *p*. So his opinion must be dependent on Beth's." The relevant assessment is supposed to be made in a way that sets aside (or "brackets") information about what Chad actually concluded.

Given the setup – in particular, given that their shared heuristic cannot be misapplied – it was certain that they would all arrive at the same verdict. So, this is a case of complete dependence – and for this reason, their shared opinion counts only as heavily as any one of their opinions would.¹⁰

Logic Exam – Blind copying: Recall the version of the logic exam case in which Chad blindly copies Beth. Chad adopts Beth’s opinion uncritically – without any regard to Beth’s reliability or to the plausibility of the opinion adopted. In assessing whether their jointly held opinion should count for more than Beth’s opinion alone, we must ask: In advance, how likely was it that Beth and Chad would agree? Given the setup, it was certain that Chad’s opinion would match Beth’s. This is another instance of (ordinary) complete dependence – for this reason, their shared opinion counts only as heavily as Beth’s would alone.

Each of the previous two examples clearly exhibits ordinary complete dependence. Considering Lackey’s variations on the simple copying example brings out important subtleties.

Logic Exam – Copying with double-checking: Recall the version of the logic exam case in which Chad copies Beth, but only after reflecting at least somewhat critically on the solution he steals from her. In assessing the significance of their shared opinion, we ask: In advance, how likely was it that they would agree? Here, there are two cases to consider.

On the one hand, we might know that when Chad double-checks a stolen answer, he never actually changes it. If we are aware of this fact, then this case is not importantly different from the blind copying case, for we will be able to see, in advance, that Chad and Beth will surely come away agreeing. Chad’s agreement would not confer any additional support.

On the other hand, we might know, somewhat more plausibly, that Chad does sometimes revise stolen answers during the double-checking process. Specifically, let us suppose that he has a 50% chance of discovering and correcting a mistake – when there is a mistake. Given this setup, we cannot be certain, in advance, that Beth and Chad will end up agreeing – since Beth may make a mistake, and Chad may find it. For this reason, Chad’s opinion is not completely dependent upon Beth’s, according to the expectational account. This explains why Chad’s agreement with Beth – if indeed they do end up agreeing – would have its own epistemic significance, as Lackey rightly suggests.

Logic Exam – Copying from a vetted source: Finally, recall what might seem to be a problematic version of the logic exam case. In this version, Chad copies Beth without

¹⁰ Additionally, it is worth noting that even if we had learned about the details of the setup after learning about the distribution of opinions, we would still have equal reason to regard this as a case of complete dependence. Once we are in possession of all relevant information, we can see that the method B users were, *in advance*, sure to arrive at the same conclusion.

double-checking Beth's answer at all. However, Chad's deference is not totally blind, as he does assess Beth's reliability in general before resolving to copy her answer. At first, it seems that this case is quite problematic for the account being considered. Given the setup, we can see in advance that Beth and Chad will come away agreeing. Nonetheless, as Lackey points out, it is intuitive that we would gain additional reason to trust Beth's answer after learning that Chad agreed with her. Isn't this a problem?

Despite appearances, this case actually confirms the expectational account of complete dependence. There are two versions of this case. In both, we know, going in, that Chad will assess Beth's general reliability, resolving to copy her answer if his assessment is a favorable one. In one version of the case, though, we do not know, in advance, how Chad's reliability assessment turned out. In the other version, we know, going in, that Chad did deem Beth reliable. Let us discuss each version in turn.

In the first version, Chad's agreement with Beth tells us something important. It gives us some evidence that Chad did assess Beth favorably. This is epistemically significant, and it may well provide additional support for Beth's answer, as Lackey suggests. But notice that the dependence condition is not satisfied in this case: Given the information we had, in advance, we could not be sure that Beth and Chad would end up with the same opinion. Chad might have deemed Beth unreliable, in which case the two could have disagreed. Since there was some probability of disagreement between Beth and Chad, their opinions are not completely dependent. Chad's agreement – should he agree – can and does confer additional support.

In the second version, we know in advance that Chad vetted Beth and found her to be trustworthy. This version does meet the condition provided: Given the setup, we can see in advance that Beth and Chad will inevitably agree. According to the view under consideration, it would seem to follow that Chad's agreement should not confer additional support upon Beth's answer. But this turns out to be the appropriate result. To be sure, Chad *is* providing us with a reason to be more confident of Beth's answer. But what is doing the work is Chad's favorable assessment of Beth – not Chad's agreement with her. And, importantly, these two bits of information can be pulled apart: Chad might be sick and miss the exam (and so he might never actually come to agree with Beth). Nonetheless, so long as we know that Chad assessed Beth favorably, we would have some reason to place more confidence in Beth's answer. Whether Chad happens to be healthy on the day of the exam (and is therefore able to copy Beth's answer) is evidentially inert. So the expectational account handles this case effectively.

6.2 Tricky Cases – Group Dependence and Noisy Copying¹¹

¹¹ Thanks to an anonymous referee for posing questions that illustrated the need for this section's inclusion.

In thinking about complete dependence, we are investigating a certain way in which opinions can be evidentially worthless. Of course, there is another way an opinion can fail to “count”: It may come from an unreliable source (that is, a source who is no more likely to answer correctly than incorrectly, given what is known). Are these the only two ways? It is tempting to suppose so. But we must pause to deal with two interesting complications to this picture. First, consider the following case, which involves dependence upon a group.

Group Dependence: Before answering a logic problem, Dawn surveys the answers of her reliable friends – Anna, Beth, and Chad – who all work independently. In the end, Dawn comes to believe whichever answer is most popular among her three friends.

Given the setup, Dawn’s opinion is not completely dependent on Anna’s, nor on Beth’s, nor on Chad’s. After all, her opinion cannot be predicted (with certainty) purely on the basis of any single friend’s opinion. And, given the setup, it is clear that Dawn is, in general, reliable. Indeed, she may be the most reliable member of her group.¹² So Dawn’s opinion is not completely dependent on any other opinion, and she is reliable in general. Still, her opinion is worthless, so long as the other three opinions are already known. Provided we know, in advance, how Dawn operates, learning her opinion tells us nothing new.

What this shows is that an opinion can be completely dependent on the opinions of a group without being completely dependent on the individual opinion of any one group member. This type of dependence is expectational, just as individual dependence is: One’s opinion is completely dependent upon the opinions of some group if it can be predicted with certainty, given full knowledge of the opinions of all group members. So Dawn’s opinion is completely dependent upon the group opinion, and this is why it does not provide additional evidence.

The second complication can be illustrated by way of an example involving a type of dependence that is mixed with random noise.

Noisy Copying: Before answering a logic problem, Chad rolls a fair, six-sided die. If the result is anything other than a six, Chad copies Beth’s answer. If the result is a six, Chad “anti-copies” Beth (i.e. goes with the opposite of whatever Beth put down).¹³

Given the setup, Chad’s opinion is not completely dependent on Beth’s. After all, his opinion cannot be predicted (with certainty) solely on the basis of Beth’s opinion. And, given the setup, it is clear that Chad is reliable: $\frac{5}{6}$ of the time, he copies Beth; $\frac{1}{6}$ of the

¹² This fact relates closely to the Condorcet’s (1785) jury theorem.

¹³ Estlund (1994) also discusses a case of what might be called ‘Noisy copying,’ though he puts the example to different use.

time, he anti-copies her. On balance, Chad will be right more than he is wrong. So Chad's opinion is not completely dependent on Beth's, and he is reliable in general. Still, his opinion is worthless, so long as Beth's opinion is already known. Provided we know, in advance, how Chad operates, learning his opinion tells us only whether he rolled a six with his die, which is not relevant to the issue at hand.¹⁴

What the preceding example illustrates is that an additional opinion can be evidentially worthless even when it is not completely dependent and comes from a generally reliable source. This should not trouble us. As noted above, an outside opinion can be evidentially worthless for either of two reasons. It may come from an unreliable source, or it may be dependent on opinions that are already known. The interesting upshot of the Noisy Copying case is that an opinion can be worthless due to a mixture of these defects. In such cases, the opinion in question will not provide any additional evidence. So, in effect, Chad's opinion contains information about two sources – Beth and the die. Beth's opinion is already known, and the die's "opinion" isn't relevant.

To conclude this section, it will be instructive to consider an example that combines group dependence and noisy copying.

Random Copying: Before answering a logic problem, Chad flips a coin. If the result is heads, he copies Anna (whose opinions are right 90% of the time); if the result is tails, he copies Beth (whose opinions are also right 90% of the time, and whose opinions sometimes differ from Anna's).¹⁵

Given the setup, Chad is reliable. Like Anna and Beth, Chad's opinions will be true 90% of the time. And, given the setup, Chad's opinion is not completely dependent on Anna's or Beth's, nor is his opinion completely dependent on the group. After all, his opinion cannot be predicted (with certainty) solely on the basis of information about either or both of their opinions. However, Chad's opinion is still evidentially worthless, so long as Anna's and Beth's opinions are already known. Why? Chad's opinion carries information about three things – Anna's opinion, Beth's opinion, and the result of the coin flip. Anna's opinion and Beth's opinion are known, and the coin flip isn't relevant.

In this section, we have seen that, by analyzing dependence expectationally, we can accommodate a wide range of cases – parrot cases, cases of autonomous copying, cases involving people that employ identical methods, cases involving randomness, and cases involving dependence on groups. But, a wider project remains. The sorts of cases we have so far examined comprise only a small portion of the cases we might encounter in

¹⁴ While Chad's opinion is worthless, given Beth's, the reverse is not true. Just as it can be helpful to obtain an original document, even after you have in hand a photocopy, it could be helpful to learn of Beth's opinion even after you are aware of Chad's.

¹⁵ Thanks to an anonymous referee for proposing a version of this example.

real life. Sometimes, the evidential value of an incoming opinion might be diminished, but not entirely eliminated, in light of other opinions we already know. A full treatment of belief dependence should do more than describe when an opinion is worthless; it should tell us just how much additional opinions are worth. The remaining sections of this paper take preliminary steps toward completing this wider project.

7 Belief Independence

7.1 An Expectational Account of Complete Independence

The previous section observes that two beliefs are completely dependent when, in effect, learning one belief tells you what the other is. In turn, we might be tempted to conclude that two beliefs are completely independent just in case information about the content of one of them tells you nothing at all about the likely content of the other. (Note that under ordinary circumstances, separate coin flips will have this very same property.)

But while this condition may capture *an* important kind of independence, it cannot be the kind of independence we are seeking. Why? Because the beliefs of sufficiently reliable thinkers will never ever meet this condition.¹⁶ To see this, suppose that Beth and Chad are known to be highly reliable: Each tends to answer correctly $\frac{5}{6}$ of the time. From this alone, it follows that their judgments will be well correlated, in a certain sense: They will reach the same conclusion at least $\frac{2}{3}$ of the time. And, therefore, information about the content of, say, Chad's beliefs will tell us a lot about the likely content of Beth's beliefs – solely because they are both so reliable.

Still, this incomplete approach serves as a helpful step toward the goal of characterizing complete independence of the salient kind. There are two kinds of cases we would like to be able to distinguish: On the one hand, there are cases in which Beth and Chad are correlated because one is copying the other, or because they are applying relevantly similar heuristics; on the other, there are cases in which Beth and Chad are correlated *only* because they are both reliable. The latter case is one of complete independence. With this observation in mind, we can offer the following.

Complete Independence: *C*'s belief (about some proposition *p*) is completely independent of *B*'s belief (about *p*) just in case:

- (1) In advance, on the supposition that *p* is *true*, learning that *B* arrived at his belief (about *p*) does not raise or lower the probability that *C* will arrive at her belief (about *p*).
- (2) In advance, on the supposition that *p* is *false*, learning that *B* arrived at his belief (about *p*) does not raise or lower the probability that *C* will arrive at her belief (about *p*).

When these conditions are met, the two beliefs – as long as *B* and *C* are reliable – provide more evidence than is provided by either belief, on its own.

¹⁶ This point is due to Jeffrey (1987, p. 392). Goldman (2001, p. 101, fn. 18) also emphasizes this point.

Let's take a look at a few variations on the usual example to see how this condition can be applied and to see just how much agreeing independent beliefs are worth.

7.2 Variations on the Two-Person Case

Suppose that Beth and Chad each answer correctly $\frac{5}{100}$ of the time. And suppose that we learn that Beth arrived at p . At this point, our confidence in the truth of p is $\frac{5}{100}$. Suppose that we are about to learn of Chad's answer, and we are wondering: If Chad ends up agreeing with Beth, how much should our confidence in p increase?

To answer this, we must assess whether Chad's belief is completely independent of Beth's. For simplicity, let's suppose that p is true.¹⁷ In effect, we want to know how (if at all) the fact that Beth answered correctly bears on the probability that Chad will answer correctly. Given the setup, the probability that Chad gets it right would seem to be $\frac{5}{100}$. When we take into account the fact that Beth answered correctly, does the probability of Chad's success change?

Of course, here it will depend on what else we know about Chad and Beth. Obviously, if we know that Chad blindly copies Beth always, then Chad's probability of success increases from $\frac{5}{100}$ to 1. In other cases, the path forward may be less obvious. Suppose we know that neither Beth nor Chad copies off the other. However, we also know that their teacher tends to teach certain topics extremely clearly and effectively, while teaching other topics somewhat confusingly. As a result, all the students tend to get the same problems right/wrong. In this case, too, learning that Beth got a given problem right would raise the probability that Chad also got it right (though it would not guarantee Chad's success). After all, if Beth answered correctly, the question is more likely to have been about the well-taught material, in which case, Chad is more likely to have answered correctly. So here, Chad's opinion would not be completely independent of Beth's.

One might wonder, then: Under what conditions *would* Chad's opinion be independent of Beth's? Let's suppose that we have elaborate track records for both students. Suppose we survey all of the cases in which Beth answered correctly and find that, in those cases, Chad answered correctly $\frac{5}{100}$ of the time. Here, we would have good reason to think that the information about Beth's having answered correctly has no bearing on Chad's probability of success. So in this case, Chad's opinion would be completely independent of Beth's.

But we aren't done. We began this section trying to determine how much the agreement of completely independent thinkers is worth. In this case, we want to know

¹⁷ Strictly, to assess independence, we would have to consider the case where p is false, in addition. But in most cases (including the one at hand), it suffices to consider just one of the two cases.

how much support for p Chad's completely independent agreement would provide. Recall that our all-things-considered confidence in p started out at $\frac{5}{6}$, given only the fact that Beth arrived at p . Supposing that Chad independently comes to agree with Beth, how much should our confidence in p increase? There are two possibilities: Either they both are right (in which case p is true), or they both are wrong (in which case p is false). How likely are each of these options? Because their opinions are independent, the probability that both are right is $\frac{5}{6} \times \frac{5}{6}$ (which is equal to $\frac{25}{36}$).¹⁸ The probability that both are wrong is simply $\frac{1}{6} \times \frac{1}{6}$ (which is equal to $\frac{1}{36}$). Since these are the only two possibilities, the final probability that both are right is $\frac{25}{26}$ (which is approximately .96). And because p is true if and only if they're both right, this result should be our new confidence in p . So our confidence in p increases from about .83 to about .96 after we learn of Chad's independent agreeing belief. More generally, the agreement of reliable, completely independent sources will always be worth more than that of completely dependent sources.

7.3 *Returning to the Random Copying Case*

Here, it will be instructive to consider and respond to an objection, as the discussion may help to illustrate how the independence constraint outlined above is to be applied. The objection appeals to a certain version of the Random Copying case. In that example, Chad copies Anna when his coin lands heads and copies Beth when his coin lands tails. (Anna and Beth answer correctly 90% of the time, and answer independently from one another.) To make things somewhat more difficult, let's suppose that Anna and Beth in fact disagree – Anna believes p , and Beth believes *not-p*. Will Chad's opinion carry additional evidential weight? As we observed earlier, it is clear that Chad's opinion shouldn't carry any weight at all in this case. But it is not obvious that the proposed picture can deliver this result.¹⁹

The tension is as follows. First, it seems clear that Chad is reliable, in a certain sense. He arrives at his views by randomly copying either of two very reliable sources. Second, it seems clear that Chad's opinions are completely independent of the other opinions, in the expectational sense put forward previously. For, given that Anna and Beth disagree, Chad's probability of answering correctly seems to be 50%, regardless of whatever new information is introduced. If indeed it is true that Chad is reliable and that his opinions are completely independent, the expectational account would say that his opinions are worth something, even after Anna's and Beth's opinions are known. Which they're not.

¹⁸ Let 'B' stand for the proposition that Beth is right; let 'C' stand for the proposition that Chad is right. We obtain: $P(B\&C) = P(B) \times P(C|B)$. Independence tells us that $P(C|B) = P(C)$, so: $P(B\&C) = P(B) \times P(C)$.

¹⁹ Thanks to an anonymous referee for offering a version of this objection.

The problem with the foregoing reasoning is an incongruity between the assessment of Chad as generally reliable and the judgment that his probability of answering correctly is 50%. There are two ways of thinking about the case. Both are acceptable, and both have the consequence that Chad's opinion is worthless. But we mustn't mix and match.

If it is part of the setup of the case that Anna and Beth disagree, then, given what we know, Chad's opinion should not be seen as reliable. For Chad will answer correctly only when his coin leads him to defer to the accurate source, which will be half the time. On this way of looking at things, Chad's opinion is indeed completely independent of Anna's and Beth's. But his opinion still lacks significance simply because it is unreliable.

On the other hand, suppose we do not take Anna's disagreement with Beth to be part of the setup of the case. With this more minimal setup in place, Chad's opinion seems clearly to be reliable: He will answer correctly 90% of the time. Using this version of the case, however, Chad's opinions will turn out not to be completely independent of the opinions of his peers. Let's start by examining whether Chad's opinion is completely independent of Anna's opinion (leaving Beth out of it altogether). We just said that, from the setup alone, Chad's probability of answering correctly is 90%. Now, suppose we learn that Anna answered correctly. Does this information raise or lower the probability that Chad answered correctly? Yes. Chad's probability of success increases: If he copied Anna, then he answered correctly for certain; if he copied Beth, then the chance he answered correctly is the usual 90% (since Beth's answers are completely independent of Anna's). Because Chad is equally likely to have copied either of them, the updated estimate of Chad's probability of success is the average of these values – 95%. So Chad's probability rises from 90% to 95% when the new information about Anna's opinion is accounted for. For this reason, his opinion is not completely independent of Anna's – according to the expectational account. A totally parallel argument can show that Chad's opinion is not completely dependent of Beth's opinion either.

But might Chad's opinion be completely independent of Anna's and Beth's opinions, considered together?²⁰ To see why not, recall that the baseline probability that Chad answers correctly, given the setup of the case, is 90%. Now suppose we learn that Anna and Beth reached different conclusions. Does Chad's probability of success change? Yes. Suddenly, the chance that Chad answers correctly falls drastically – to 50%. After all, Chad chose at random which of them to copy, and only one of them answered correctly. Because the new information that we gained about Anna's and Beth's opinions does

²⁰ Strictly speaking, the expectational constraint outlined earlier only describes what it takes for one opinion to be completely independent of another. To assess whether an opinion is completely independent of multiple opinions, we simply examine whether knowledge of the multiple opinions has any bearing on the content of the first opinion – both on the supposition that the proposition in question is true, and on the supposition that the proposition in question is false.

change Chad's probability of success, Chad's opinion is not completely independent of the opinions of Anna and Beth, considered together. So, in sum, the expectational account does not force us to say, implausibly, that Chad's randomly copied opinion has epistemic significance.

8 Partial Dependence

So far, we have seen how to understand complete dependence and complete independence. But this still leaves out a lot of cases. And, it turns out that considering these cases – cases of partial dependence – introduces some interesting complications.

We have seen that the agreement of reliable, completely independent sources provides more evidence than is provided by the agreement of comparably reliable, completely dependent sources. For this reason, it is tempting to conclude that the value of partially dependent agreement falls somewhere in between these extremes. And, often, this will be true. But not always. Certain forms of partial dependence can actually render the agreement of reliable sources even more valuable than that of comparably reliable completely independent sources. To see this, we will start with a simple case, where the evidential value of partially dependent agreement does land in between that of completely dependent and completely independent agreement. Next, we will consider the more unusual sort of case.

The following case can be described as 'two against two' – with two completely independent sources believing one thing, and two partially dependent sources believing the opposite. Suppose that Beth and Chad – who, as usual, each answer correctly $\frac{5}{6}$ of the time – occasionally copy each other (and are therefore partially dependent). In addition, suppose that Dawn and Eric also happen to be in the same logic class. Dawn and Eric are undercover logicians, so, for any question, they know what the right answer is. But in order to keep their identities secret, they answer incorrectly $\frac{1}{6}$ of the time by using a six-sided die. Each keeps a personal die and answers incorrectly whenever '6' is rolled – so Dawn's and Eric's answers are completely independent.

You are sitting in the center of all four of them, and you know all of the above information. On this particular problem, you can see that Beth and Chad arrived at p , while Dawn and Eric both arrived at $\sim p$. You do not know anything about this logic problem, and you did not see the dice that were discreetly rolled. Which answer should you go with? As it turns out, the kind of partial dependence exhibited by Beth and Chad renders their shared belief less significant than the shared opinion of Dawn and Eric, which makes it reasonable for you to be more confident of $\sim p$. A simple argument can show this.

We know that something somewhat unlikely happened. Either Beth and Chad both arrived at the wrong answer, or else two fair dice both came up '6.' Which of these is

more likely? The second possibility involves the coincidence of two unlikely events, which have no bearing on one another – Dawn’s rolling ‘6,’ and Eric’s rolling ‘6’. The first possibility, in contrast, involves the coincidence of two unlikely events that frequently tend to coincide – Beth’s erring and Chad’s erring. Because one sometimes copies off of the other, it is – relatively speaking – not that uncommon for Beth and Chad to get the same problem wrong. This is the better explanation of what we have observed, so Dawn and Eric’s answer of $\sim p$ is more likely to be correct. So, in this case, the agreement of partially dependent sources was not worth as much as that of comparably reliable, completely independent sources.

If we were to take this example as a guide, we might well conclude that the agreement of completely independent sources is *always* worth more than that of comparably reliable partially dependent sources. But, as noted earlier, this turns out not to be the case. Partial dependence, like complete dependence, comes in two main varieties: *Ordinary* partial dependence obtains whenever two sources agree more often than comparably reliable independent sources would. *Perverse* partial dependence obtains when two sources agree less often than comparably reliable independent sources would. In the example above, the dependence between Beth’s and Chad’s beliefs took the ordinary form: They agreed with each other more often Dawn and Eric do. But, one might wonder, what would happen if two reliable sources exhibited partial perverse dependence? How much would the agreement of two reliable contrarians be worth? Intriguingly, it turns out that the agreement of such thinkers is worth more than that of even completely independent sources.

Suppose that Beth and Chad are contrarians, always looking for opportunities to disagree with one another. But they do not always disagree: On the whole, each is right $\frac{5}{6}$ of the time – so they still must agree quite often. However, let’s suppose that they disagree as much as they possibly could, given their reliability: $\frac{2}{3}$ of the time, they agree and are both correct; the other $\frac{1}{3}$ of the time, they end up disagreeing, with Beth’s being correct in exactly half of those remaining cases (so each is still wrong $\frac{1}{6}$ of the time overall). Given this setup, it follows that Chad and Beth never agree on a wrong answer: At least one of them is always right. Immediately, it follows that their agreement is quite valuable. Given the setup, it is clear that, when these two do agree, we can be totally certain that their answer is correct.

One may wonder about what sort of lesson should be taken from this observation. For one thing, the example is unrealistic: After all, how often can we be totally sure, in advance, that two thinkers could not both settle on the wrong conclusion? This seems to me to be a fair complaint. In response, it is worth noting that even if an extreme case like the one described is unlikely to crop up in real life (that is, a case where agreement guarantees truth), milder instances of the same phenomenon (reliable sources exhibiting per-

verse partial dependence) almost certainly do. Imagine, for example, that Chad and Beth are not contrarians, but instead have complementary skill sets – such that they almost never make mistakes in the same area. This does not seem unrealistic, and (depending on the details) it could be that the two rarely, if ever, agree upon wrong answers – because every question falls in an area where at least one of them is quite strong. In such a case, it would be reasonable for us to be quite confident of any answers the two do agree about. While perverse dependence (of reliable sources) may be somewhat rarer than ordinary dependence, it still seems worth taking seriously.

The following table summarizes the observations made in these last few sections. The table assumes, where applicable, that each person answers correctly $\frac{1}{3}$ of the time.

Example	Type of dependence	Confidence after learning one opinion	Confidence after learning of both opinions when they agree
Chad always copies Beth	ordinary complete dependence	~ .83	~ .83
Chad sometimes copies Beth	ordinary partial dependence	~ .83	between ~.83 and ~.96
dice-rolling logicians Dawn/Eric	complete independence	~ .83	~ .96
Chad/Beth have complementary skill sets	perverse partial dependence	~ .83	between ~.96 and 1
Chad always anti-copies Beth	perverse complete dependence	~ .83	n/a ²¹

To conclude this section, it will be useful to look at another toy example to see how these observations can be applied to realistic situations. Though we cannot say anything hard and fast about many real-world cases, the understanding of dependence and independence that has been developed here may help to highlight the kinds of questions that are relevant, in thinking about the degree of support outside opinions provide.

Moved by Argument: Suppose that Beth and Chad agree because Beth came up with an argument and Chad found it convincing. And suppose that they disagree with Anna, who has her own argument, but has not yet tried it on anyone. How much of a boost does Chad's agreement provide?

To simplify matters, a first assumption to make is that Anna, Beth, and Chad are reliable – and equally so – with respect to the topic under dispute. Obviously, if the topic is, say, politics, or philosophy, or religion, then questions of a person's reliability will be

²¹ If Chad always anti-copies Beth, then the two will never agree. Moreover, if Beth's reliability is $\frac{1}{3}$ in such a case, then Chad's reliability will necessarily be only $\frac{1}{3}$. So the anti-copying case is incompatible with the setup.

quite difficult, if not impossible, to settle. But such assumptions are necessary if we are to say anything substantive about how the issue of dependence interacts with cases like these.

The next questions to ask concern dependence, particularly that between Beth and Chad: First, if Beth had offered a good argument, how likely is it that Chad would have been persuaded? Second, if Beth had offered a bad argument, how likely is it that Chad would have been persuaded? In answering these questions, we can't use the fact that Chad ended up agreeing with Beth – since these questions are to be evaluated *in advance* of their agreement. However, we can use certain facts about the people involved. If, for example, Beth is known to be especially charismatic, or if Chad tends to be a contrarian around Beth, then our answers to the above questions should accommodate these data. To the extent that we expected, in advance, that Beth would surely convince Chad (no matter how good or bad her argument was), Chad's agreement will add little credibility to the proposition they jointly hold. So Anna's position and Beth's/Chad's position will be on a par, epistemically. To the extent that we expected, in advance, that Chad could well have disagreed with Beth (despite being just as reliable as she) Chad's agreement will make their position more credible – and it will be reasonable to place more confidence in it than in Anna's position. The size of the advantage that Beth's and Chad's shared view enjoys over Anna's view will depend (in part) upon how likely it was that Chad would have disagreed.

9 Application to Credences

So far, we have only talked about dependence and independence exhibited by *beliefs*. But some epistemologists prefer to think and talk in terms of credences (degrees of confidence), rather than in terms of all-or-nothing beliefs. Within this framework, one should consider the question of how the credences of others should affect one's own credences.²² A complete answer to this question would include some treatment of dependence and independence. So it is well worth pointing out that the expectational account of dependence and independence discussed in the previous sections can be applied to a framework that uses credences in place of beliefs – with minimal modifications. Here are suitably modified versions of the principles earlier presented.

Complete Dependence for Credences: *C*'s credence of *c* (in some proposition *p*) is completely dependent on *B*'s credence of *b* (in *p*) just in case it is rational to be certain, in advance, that *C* will have a credence of *c* (in *p*) if *B* has a credence of *b* (in *p*). When this condition is met, the

²² For a general theory of how to aggregate credences, see Easwaran et al (2016). They do not discuss the kind of dependence we have focused on. For this reason, the expectational account of dependence and independence developed here may serve to complement to their well developed view.

two credences provide just as much evidence for or against p as is provided by B 's credence, on its own.

Complete Independence for Credences: C 's credence of c (in some proposition p) is completely independent of B 's credence of b (in p) just in case:

- (1) On the supposition that p is true, the fact that B has credence b in p does not raise or lower the probability that C will have credence c in p .
- (2) On the supposition that p is false, the fact that B has credence b in p does not raise or lower the probability that C will have credence c in p .

When these conditions are met, the credences of reliable sources provide more evidence than is provided by either credence, on its own.

Applying the first of these principles to an example will help to bring out an important feature of the expectational account – a feature we largely ignored in our earlier discussion.

To the usual logic exam setup, we add the following details. Beth's confidence in her answers is sometimes .9, and never higher. The answers she has .9 confidence in are the answers she is most sure about. And, Beth always makes a written note of her exact credence next to whichever answer she selects. Chad can see Beth's notes and uses the following procedure: When he sees that Beth is at her highest level of .9, he simply copies her answer, coming to have a .9 credence in that answer as well; but when he sees that Beth is somewhere below .9, he works through the problem himself, and comes to have a credence of .8 in whichever answer he arrives at.

With these details in place, one might ask: Do Chad's opinions provide evidence beyond that provided by Beth's opinions? Answer: It depends. On the one hand, if Beth and Chad both end up with a credence of .9 in some proposition p , then, according to the expectational account, Chad's opinion is completely dependent on Beth's – and, therefore, it provides no additional evidence for p . On the other hand, if Beth and Chad both end up with a credence of only .8 in p , then Chad's opinion would provide some additional support for p beyond that provided by Beth's credence. What this package of observations illustrates is that opinions, rather than sources, are what should be assessed for dependence or independence, in a given case. One specific opinion can be completely dependent on another even when the sources of those opinions sometimes arrive at their views independently. In a similar vein, one specific opinion can be completely independent of another, even when the source of that opinion sometimes copies the other source. Though it is sometimes useful to speak of "the agreement of dependent/independent sources" to summarize general observations, it is important to remember that, strictly speaking, opinions, not sources, are the proper unit of analysis.

10 Conclusion

In some sense, it is clear that the numbers count. That is, it is clear the number of think-

ers on a given side of a disputed issue is typically relevant to the degree of support their opinions provide. It has often been suggested that the extent to which these opinions are dependent or independent also has substantial epistemic import in such situations. In this paper, we have seen how to characterize the type of dependence that can play this role. Though many cases that illustrate the operative phenomenon tend to involve dependence of a causal variety, we have seen that dependence of the relevant kind should be understood expectationally, not causally – though causal relations are a particularly salient way of *producing* expectational dependence. Additionally, we have seen that we can make sense of expectational dependence regardless of whether we think in terms of credences or in terms of all-or-nothing beliefs. Though in any real-life circumstance, applying the account discussed will inevitably be an inexact science, a greater understanding of the nature of the relevant kind of dependence and independence puts us in a position to understand better how and when the numbers count.²³

References

- Christensen, D. (2014) “Conciliation, Uniqueness, and Rational Toxicity,” *Noûs* 49 (4): 1–20.
- Condorcet, M. (1785) “Essay on the Application of Mathematics to the Theory of Decision-Making,” in Baker, K. (transl. and ed.) *Condorcet: Selected Writings*. Bobbs-Merrill, 1976: 33–70.
- Easwaran, K., Fenton-Glynn, L., Hitchcock, C., and Velasco, J. (2016) “Updating on the Credences of Others,” *Philosopher’s Imprint* 16 (11): 1–39.
- Elga, A. (2010) “How to Disagree About How to Disagree,” in *Disagreement* (eds. Feldman and Warfield), Oxford University Press: .
- Estlund, D. (1994) “Opinion Leaders, Independence, and Condorcet’s Jury Theorem,” *Theory and Decision* 36 (2): 131–162.
- Goldman, A. (2001) “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63 (1), 85–110.
- Jeffrey, R. (1987) “Alias Smith and Jones: The Testimony of the Senses,” *Erkenntnis* 26 (3): 391–399.
- Lackey, J. (2013) “Disagreement and Belief Dependence: Why Numbers Matter,” in *The Epistemology of Disagreement: New Essays* (eds. D. Christensen & J. Lackey), Oxford University Press: 243–268.
- Vavova, K. (2014) “Confidence, Evidence, and Disagreement,” *Erkenntnis* 79 (1): 173–183.

²³ For helpful advice, I’m grateful to an anonymous reviewer at *Philosophical Studies*, Nomy Arpaly, Anna Brinkerhoff, Jamie Dreier, Nina Emery, Dave Estlund, Arianna Falbo, Tobias Fuchs, Kelly Gaus, Yongming Han, Rachel Leadon, Nick Leonard, Han Li, Chad Marxen, Adam Pautz, Justin Pombrio, Brett Topey, Leo Yan, and members of the Dissertation Workshop at Brown University. Special thanks to Josh Schechter and David Christensen.