

---

## Disagreement and Belief Dependence: Showing When and How the Numbers Count

Word Count: 2999

---

While taking a logic exam, you encounter a problem you do not know how to do. So you decide to cheat. Your friends Alice, Bob, and Carol 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: Alice obtained  $\sim p$ ; Bob and Carol obtained  $p$ . What do you do?

Other things the same, you should go with  $p$ . You have two reliable sources against one. But suppose that you saw Carol copy off of Bob. With 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,’ Carol’s opinion is *dependent* on Bob’s, in some important sense. And, for this reason, it seems not to “count” (i.e. it seems not to provide additional support for  $p$  beyond that provided by Bob’s opinion). To accommodate cases like this, we might offer 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 Adam Elga:<sup>1</sup>

[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.

Elga regards this claim as “completely uncontroversial” and suggests that “every sensible view on disagreement should accommodate it.” Tom Kelly, writing from the other side of the disagreement debate, shares Elga’s outlook.<sup>2</sup> But despite its widespread appeal, Jennifer Lackey (2013) argues, persuasively, that Belief Dependence cannot be generally true:<sup>3</sup>

---

<sup>1</sup> See Elga (2010, p. 177).

<sup>2</sup> See Kelly (2010, p. 148).

<sup>3</sup> See Lackey (p. 245).

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.

This paper investigates the issue. The first section summarizes Lackey's argument. The second section defends Belief Dependence from Lackey's attack. The third and final section offers a positive theory of this kind of dependence.

## 1 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 hand<sup>4</sup>).<sup>5</sup> 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 a peer's opinion, 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.<sup>6</sup> But she suggests that each is no good. Ultimately, Lackey argues that this widely held principle cannot be sustained.<sup>7</sup>

---

<sup>4</sup> See Lackey (p. 243 and p. 245).

<sup>5</sup> 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 engage this concern, I think that it is clear how it can be addressed. 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 control for a confounding variable.

<sup>6</sup> Strictly speaking, Lackey's version of the principle is slightly different (p. 244):

**Belief Independence:** 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.

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

<sup>7</sup> See Lackey (p. 245).

It would seem that Lackey is rejecting flat out the intuition elicited by the logic exam case – provided that peerhood between Bob and Carol is stipulated. With peerhood in place, Lackey’s view seems to entail that, contra appearances, Carol’s opinion, together with Bob’s, somehow counts for more than Bob’s opinion does alone. And if we add that Alice, too, is a peer of Bob/Carol, then it seems to follow that you, the cheater, would have reason to favor Bob+Carol’s joint answer over Alice’s answer – even if you were certain that Carol copied Bob.

Lackey offers an intriguing diagnosis of this result. She points out that the case is underdescribed. Though we know that Carol’s opinion was, in some sense, grounded in Bob’s, we are not told whether Carol was at all critical in her decision to endorse Bob’s opinion. Here, Lackey distinguishes what she calls *autonomous* and *non-autonomous* dependence:<sup>8</sup>

The autonomous version of this dependence involves a subject exercising agency in her reliance on a source... critically assessing its reliability, [and] monitoring for defeaters... This, I take it, is the minimum required for rational belief formation.

Applying this distinction to the case at hand, we can observe that either Carol was autonomous in her reliance on Bob or she was not. Whichever way we go, Lackey thinks, we will not need to appeal to anything like Belief Dependence for Peers to deliver the correct verdict.

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

First option: If Carol simply engaged in a bit of double-checking before endorsing Bob’s answer of *not-p*, then it seems plausible that her agreement does confer at least some additional support upon the answer they both favor. After all, Alice’s answer was not double-checked, and

---

<sup>8</sup> See Lackey (p. 249).

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 Carol's opinion must be providing some support of its own.

Second option: What if Carol did not double-check Bob's answer, but did at least confirm Bob's reliability before resolving to trust him? Here, too, we can make it plausible that Carol's agreement should carry at least some epistemic weight. To make this point as clearly as possible, imagine, realistically, that your own reliability assessments of Alice and of Bob 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 Carol agreed with Bob is evidence that Carol assessed Bob's reliability favorably – which does seem to render their shared opinion at least slightly more credible than Alice's opinion. After all, we now have more evidence for Bob's reliability than we do for Alice's.

In either case, we find that – so long as Carol's reliance on Bob was autonomous – Carol's apparently dependent opinion seems still to have some epistemic significance.<sup>9</sup> But what if Carol's reliance on Bob was not autonomous? What if, to use Lackey's term, Carol simply *parroted* Bob? Here, Lackey agrees that Carol's opinion does not provide additional support for the position she and Bob share. But she notes that we do not need to appeal to Belief Dependence for Peers to explain this. Since Carol is non-autonomous in her reliance on Bob, she would defer to him even if he were thoroughly unreliable; she would adopt his beliefs even if they were patently false. On this issue, Carol's belief-forming process is manifestly irrational. Here, Lackey might well ask: Are Bob and Alice just as irrational as Carol? If they are, then it

---

<sup>9</sup> One could object that, by making Carol's reliance autonomous, we have rendered his opinion at least partially *independent* of Bob's. I am sympathetic to this point of view; section 3 discusses an expectational account of belief dependence that can deliver this result. However, I still see an intuitive sense in which Carol's opinions still are dependent on Bob's (e.g. causally), and in this sense, Lackey's verdict seems to be exactly right.

seems clear that *none* of their opinions has much epistemic significance at all, for reasons that have little to do with belief dependence. But if Bob and Alice are more rational than Carol, then the imagined scenario is irrelevant to the principle at hand, for the principle applies only to cases involving epistemic peers. In neither case do we need to invoke anything like Belief Dependence to explain why Carol's parroted opinion lacks epistemic significance.

The logic exam case seemed to illustrate the need for some kind of Belief Dependence principle. But upon closer inspection, it is not at all clear that such a principle is needed to accommodate this case.

## 2 The Indispensability of Belief Dependence

As it turns out, we cannot abandon Belief Dependence. Consider the following 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.

Dia 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 both 90% reliable at determining a given chicken's sex.

A chicken walks by. Dia, 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. Given all of the above information, how confident are you that the chicken is male?

Not 99.9999% confident, I take it. 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 both heuristics can never be misapplied, then, plausibly, we can make it reasonable for you to afford equal credence to the female and male hypotheses.

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

A proponent of Lackey’s view can push back against the problem posed by this case. The setup suggests that Dia 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, it suffices to revise the case. Suppose that all of the chicken-sexers have access to *both* heuristics, but, for whatever reason, Dia uses method A, while the others all use method B. So imagined, Dia and her counterparts may well be peers, despite having used different methods on this occasion.<sup>10</sup> So it is clear that 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 final section investigates the nature of this dependence.

### **3 Belief Dependence as Expected Correlation**

We have seen that dependent beliefs do sometimes lack epistemic weight: the shared judgment of the many method B users counts no more heavily than does Dia’s rival judgment. But an

---

<sup>10</sup> Leave this point aside. Even if we omit Dia from the story altogether, there still seems to be a clear difficulty for Lackey’s view. 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. Intuitively, is there 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.

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.<sup>11</sup> Even if this condition is stipulated, the verdict does not seem to 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 opinions 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 or more thinkers are dependent in the relevant sense, what matters is not whether one causes the other(s), but rather, *whether they should be expected to reach the same conclusion*. When Carol copies Bob'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:** Multiple opinions are completely dependent just in case it is rational to be certain, in advance, that the opinions will match.

When this condition is met, the agreeing opinions confer no more support for the proposition believed than would be provided by any one of these opinions, on its own.<sup>12</sup>

Let us apply this expectational account to some of the cases we have discussed thus far.

---

<sup>11</sup> 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).

<sup>12</sup> This principle is consistent with that found in Goldman (2001, pp. 99-100), though it is simpler. For Goldman, two opinion-holders X and Y are dependent with respect to some hypothesis H, just in case:

$P(X \text{ believes } H \mid Y \text{ believes } H \ \& \ H \text{ is true}) = 1$  and  $P(X \text{ believes } H \mid Y \text{ believes } H \ \& \ H \text{ is false}) = 1$ .

If suspension of judgment is ignored, this condition is equivalent to that advanced above. In a framework that allows for suspension of judgment, Goldman's condition is necessary but not sufficient for total dependence.

*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 they would all agree? Given the setup (especially: given that the heuristic which they are all applying cannot be misapplied), it was certain that they would all arrive at the same verdict. This is a case of complete dependence – their shared opinion counts only as heavily as any one of their opinions would.

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

*Logic Exam – Copying with double-checking:* Now recall the version of the logic exam case in which Carol copies Bob, but only after reflecting at least somewhat critically on the solution she steals from him. In assessing the significance of their shared opinion, we ask: How likely was it that they would agree? Here, there are two cases to consider.

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

On the other hand, we might know, somewhat more plausibly, that Carol does sometimes revise stolen answers during the double-checking process. Specifically, let us suppose that Carol 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 Bob and Carol will end up agreeing – since Bob may make a mistake, and Carol may find it. Carol's opinion is not fully dependent on Bob's, according to the expectational account. This explains why Carol's agreement with Bob – 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 somewhat problematic version of the logic exam case. In this version, Carol copies Bob without double-checking Bob's answer at all. However, Carol's deference is not totally blind, as she does assess Bob's reliability in general before resolving to copy his answer. At first, it seems that this case is quite problematic for the expectational account. For, given the setup, we can see in advance that Bob and Carol will come away agreeing. Nonetheless, as Lackey points out, it is intuitive that we would gain additional reason to trust Bob's answer after learning that Carol agreed with him. Isn't this a problem?

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

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

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

## References

- Elga, A. (2010) “How to Disagree About How to Disagree” in *Disagreement* (eds. Feldman and Warfield).
- Goldman, A. (2001) “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63 (1), pp. 85-110.
- Kelly, T. (2010) “Peer Disagreement and Higher-Order Evidence” in *Disagreement* (eds. Feldman and Warfield).
- Lackey, J. (2013) “Disagreement and Belief Dependence: Why Numbers Matter” in *Disagreement: New Essays* (eds. Christensen and Lackey), pp. 243–268.