Confirmation, Coincidence, and Contradiction

> It so rarely happens that witnesses of the same transaction perfectly and entirely agree in all points connected with it, that an entire and complete coincidence in every particular, so far from strengthening their credit, not unfrequently engenders a suspicion of practice and concert.
>
> Thomas Starkie, *A Practical Treatise of the Law of Evidence*

1. Introduction

Suppose that two witnesses agree in asserting that some event has occurred--say, a robbery or a car accident at some specified place and time. What other aspects of their individual testimonies, as they compare with one another, are helpful or harmful to the confirmation of the proposition that this event took place?

It is natural enough to assume that the more the witnesses agree on the specifics of their testimony, the more strongly they confirm both one another and the event. If we stipulate that the testimonies are conditionally independent on the proposition that the event did *not* take place (which I will call ~H in this article) and that they are not negatively relevant to one another conditional on the proposition that the event did take place (which I will call H in this article), then this natural assumption holds (see Huemer 2007).

Matters are far more complicated, however, if we do not make such stipulations. I propose to describe a non-trivial, though somewhat surprising, type of situation in which exact agreement between two testimonies on some minor detail actually disconfirms the proposition that the broader event took place, while outright contradiction on that detail strengthens the case. Statements from the witnesses that do not constitute direct agreement on the detail but do agree on the detail in an indirect fashion (to be explained in section 5) do not generate this oddity and are “better” for confirmation of H in these situations than either outright contradiction or direct and explicit agreement.

This somewhat paradoxical situation arises when there is significant background reason to think that, at least if the event did not occur (and perhaps even if it did), one of the alleged witnesses or sources is copying his testimony from the other in a fashion that would permit him to be careful about avoiding contradictions. Such a situation would be especially likely to arise when the best-supported dependence hypothesis is a literary theory: The author of one document had access to another and therefore had the opportunity to copy it or at least to avoid contradicting it in details as well as in general outline.

Another important aspect of the situation that gives rise to such a probability distribution is that the detail is sufficiently specific or difficult that it would be fairly easy for truthful, well-informed witnesses with access to the main event to make an honest mistake about the detail in their testimony. An honest mistake could give rise to contradiction in independent or partially independent witness testimony to a real event. In the situations I have in mind, it would be more probable that an honest second witness to the event got the detail wrong, resulting in a contradiction, than that a colluding or copying source contradicted another source concerning that detail.

2. Sloppy collusion: The flat tire example

I begin with a standard probability theory example. Four students miss a Monday exam. They all come to their professor with the same appeal: According to the story, they went away together for the weekend, meaning to be back in time for the exam, but they were delayed in their return by a flat tire and hence missed the Monday exam. Now they would like the chance to take it. The professor, suspecting that they have made up a false story to cover for the fact that they were partying on the weekend and needed extra time to recover and study, grants their request to take a late exam. He puts them in separate rooms and asks just one question for the exam: "Which tire?"

Assuming that they each randomly select a tire for their answer, one can easily calculate the probability that they will all name the same tire (1/64). The assumption of the exercise is that, if they all were to agree on the tire (which the professor does not expect), this would be significant evidence that their story is true. If their story is false, there was of course no flat tire. The professor assumes that they did not collude on the question, "Which tire?" though they did (he suspects) collude on the main outline of the story. Thus, if their story is false, each student must name a tire at random on the exam.

I will call the scenario that the professor suspects sloppy collusion. In sloppy collusion, alleged witnesses or sources are copying one another concerning the main story, but they either don't have the opportunity or don't take the trouble to agree about some detail. In this case the professor is guessing that the students didn't think ahead well enough to collude on the tire. Once he separates them for the "exam," the opportunity has passed. Thus they are forced to guess at random, unless their story happens to be true.

When we take it that disagreement on the detail would be significant evidence *against* the flat tire story, we are also assuming that, if the story were true, it would be quite easy for the students to get the detail right. The weekend in question was in the recent past, and there are only four tires to choose from. The assumption is that they would have had the opportunity to know which tire went flat and to remember it. One can, of course, finesse the example and make it more complicated by asking what the effect would be if three out of four of the students agreed on the tire. This might be considered enough to confirm the story, since one student might have been asleep during the incident or incompetent to help change the tire. Maybe he stayed in the car. The agreement of three out of four on the specific tire is still quite improbable on chance (1/16).

The set-up tacitly assumes that agreeing testimonies on the detail are positively relevant given H. Why is this the case on the assumed background evidence? Suppose that H is true. Then the students were present and had some opportunity to know which tire went flat. Hence if one of these witnesses, who is truthful *ex hypothesi* on the main event, says that the rear passenger side tire was flat, one has some additional reason to expect the next alleged witness (who is, given H, truthful on the main event) to say the same. The probability that Witness 2 will say that it was the rear passenger tire, given that the event took place and that Witness 1 said so, is greater than ¼--better than chance. By the same token, contradictory testimonies are negatively relevant given H, though the degree of negative relevance is not necessarily the same as the degree of positive relevance (see next section).

Of course, if the students' testimony could be taken as *entirely* independent given ~H, their agreement even on the main outlines of the story would be *very* strong evidence all by itself in favor of its truth. Agreement on the question, "Which tire?" would only increase this effect. After all, if the flat tire story is entirely false and they didn't collude *at all*, it would be remarkable for all four of them independently to decide to tell the teacher that they were together that weekend and had a flat tire. It's obvious in the example that the teacher is not granting *that* degree of independence and is, in fact, taking quite seriously both ~H (there was no flat tire) and the sloppy collusion hypothesis. While it is certainly true that entirely independent testimony to an event will mount up quite quickly into a strong case for the event (Earman 2000, pp. 55ff, Holder, p. 53), the assumption that we can absolutely rule out all positive dependence given that the event did not occur is generally something of a probabilistic fiction.

A strong argument can be made that the desire to reduce dependence given the negation of the hypothesis in question is the reason why we take varied evidence to be so valuable (L. McGrew, 2016a). If two different kinds of cancer tests both come out positive, this is better evidence that the patient has cancer than if the very same test were to come out positive twice within a short period of time. Why? Because whatever factors would produce a false positive result on that test the first time seem likely to produce it again the second time. (We know that cancer is not like a viral illness, where the causal factors that affect a test may vary greatly within a short time period.) In contrast, the factors that might produce a false positive on, say, an x-ray would be (on background evidence) unlikely to produce a false positive on a biopsy.

Analogously, witnesses who tell the same story in varied words and giving varied specifics do not look as much like they are colluding or copying from one another as witnesses who use the same terms and give the same specific details. Alleged witnesses who tell their stories in varied ways do not appear to have gotten the story from a common source or from one another, meaning that, if the main story is *false*, the two witnesses would have had to attest to this false story *via* different causal routes (e.g., both independently mistakenly believing it to be true or both deciding to lie in that way) rather than by the simple route of colluding or copying.

It is provable that, all else being equal, the reduction of positive dependence given ~H between multiple items of evidence that are individually positively relevant to H is helpful to H. (The relatively straightforward proof is found in L. McGrew 2016a, pp. 270-272, 290-291.) So assuming that we cannot absolutely rule out causal dependence given ~H, which would affect probabilistic dependence, the next most informative thing we can look for is *evidence* for or against dependence given ~H. We can seek this evidence both in the background and in the stories themselves. It is also, of course, relevant to look for evidence concerning dependence given H. There are cases, including the one under consideration in this section, where items of evidence are positively dependent given H, which is also helpful to confirmation.

The assumption of complete independence given ~H is too simple for most real-life situations, but so, interestingly, is the assumption of sloppy collusion. In the flat-tire case, can we really be so confident that the students, if they colluded on the story, would *not* have thought to collude on the tire? Perhaps not, but it does not seem all that far-fetched to picture them doing so. The assumption of independence on the specific tire, given ~H, is a simplifying assumption made to produce a probability problem for educational purposes of our own. The hypothetical professor, we gather, does not have a very high regard for his students' intelligence or, perhaps, for their sobriety at the time when they were making up their false story. He therefore pictures them colluding hastily and not thinking of such detailed matters. One can well imagine his having background evidence to support this assumption. The fact that they originally told their story to him orally and that they had no access to each others' answers when they were pressed to say which tire went flat is also relevant.

In real life, if they *did* all say on their separate tests that it was the rear passenger-side tire that went flat, one would be confronted with a choice between thinking that their story was true after all and thinking that they were smarter at jointly making up a false story than one had previously suspected. One can well imagine scenarios where the hypothesis of sloppy collusion on a particular detail is actually quite improbable, and the cases we will discuss below are of that sort. The more sophisticated way to test the students' truthfulness would be to question them separately, letting them tell their stories at more length, noting detailed points of agreement, contradiction or apparent contradiction, wording, indirect confirmation or disconfirmation of each others' stories, etc. Asking, "Which tire?" would be just be one part of a larger process of investigation.

3. Agreement equals good, contradiction equals bad?

It is easy to assume that, when two or more sources agree on a main story, direct agreement on details is uniformly "good" for the confirmation of the main story and contradiction uniformly "bad." It would be tempting to jump to the conclusion that, if agreement is confirmatory of some H, contradiction must be *just as* disconfirmatory. For example, suppose that we take the ratio P(E|H)/P(E|~H) (the simple ratio form of the Bayes factor) as a measure of confirmation. (For the reason for using this measure in this paper, see footnote 3.) Let H be the proposition that some main event (say, a bank robbery), described in specific enough terms to pick it out, took place. Then, if agreement between two sources on some further detail had, let us say, a Bayes factor of 2/1 in favor of H, it might be tempting to think that a direct contradiction on that same detail would have a Bayes factor of 2/1 in favor of ~H.

It is necessarily true that, if agreement on some detail confirms H, disagreement on that same detail disconfirms H, all else being equal. Necessarily, if E confirms H, then ~E disconfirms H, and if "agreement on this detail" is the E in question, contradiction is certainly one variety of "non-agreement on this detail." But it does not at all follow that the confirmation and disconfirmation are of the same order. We can see something similar in evaluating arguments from silence. The fact that an author or witness attests to a particular event is evidence for the event if he has any credibility at all. By the same token, his silence is *some* evidence against the event, but the force of sheer silence may be negligible. The argument from silence often has very little weight even if the positive testimony of the same source to that same fact would be strong (see T. McGrew 2014).

By analogy, the first adjustment we should make to the idea that direct agreement on detail is good for confirmation and that contradiction is bad is to question the assumption of symmetry in evidential force. If we stipulate that, conditional on ~H, the witnesses testify independently concerning the detail in question, how much their agreement confirms H and how much their disagreement confirms ~H will depend entirely upon the degree of dependence that their agreeing testimonies have conditional on H. In other words, how much would we expect agreement on that detail if the main event really happened?[1] The probability that the witnesses

[1] This corresponds to asking, in the evaluation of historical arguments from silence, how probable it is that a speaker or author would mention an event if it really took place. If that probability is low in absolute terms, that will make the argument from silence weak (T. McGrew 2014), *ceteris paribus*. The argument from silence is only an analogue to an argument from disagreement on detail, but it is a useful comparison.

will agree on the detail if the main event took place depends in turn on how accurate or fallible we expect them to be concerning that detail.

Let us assume that the witnesses are forthcoming in the scenario; it is not an option for them not to address this detail. This models the case of the test, in which the students must answer, “Which tire?” Let us simplify further by considering only two students rather than four and evaluating their agreement or disagreement. Since we have no other way of knowing which tire went flat if there really was a flat tire, nor any reason for preferring one tire over the other on chance, the probability that the first alleged witness says any particular tire is .25 on either H or ~H, a factor which can simply be canceled out in considering the overall Bayes factor for agreement. So, once we have the first testimony in hand, agreement or disagreement depends entirely on what the second student says. Suppose that we think that, if the flat tire really occurred, the witnesses are infallible about which tire went flat. Then of course a contradiction *entails* that H is false. For if the witnesses are infallible on this detail given that H happened, the probability, given H, that Witness 2 says the same thing as Witness 1 is 1 and the probability of contradiction is 0. Hence, the power of agreement to confirm H is obviously much less than the power of disagreement to disconfirm H. Disagreement is absolutely fatal to the truth of the main story, while agreement merely makes H more probable.

But why should we consider the alleged witnesses to be infallible even if there really was a flat tire? Matters are much different if we permit even a moderate amount of fallibility on the part of the witnesses. Even if the flat tire story is true and both witnesses have testified truly that it happened, suppose that they are not infallible concerning which tire was flat. Suppose, for example, that there is only a .7 probability that the second witness will name the same tire as the first, conditional on H (and given that both agree already on H). The agreement on the detail remains, let us say, random conditional on ~H. Then agreement confirms H by a factor of 7/2.5 or about 2.8/1. Conditional on H, the probability that the second witness disagrees with the first, naming some tire other than the tire named by the first, is .3. The probability that the second witness disagrees with the first conditional on ~H is .75. So disagreement confirms ~H by odds of 75/30 or about 2.5/1, which is somewhat less than the confirmation of H by agreement.

As the fallibility of otherwise truthful witnesses, reflected in lesser dependence between the two testimonies conditional on H, rises (up to a limit, see below), this disparity continues. For example, if the second witness has only a .5 probability of saying the same tire as the first, conditional on H, then the Bayes factor of agreement in favor of H is 2/1 and of contradiction in favor of ~H is 75/50 or about 1.5/1. Even if a second truthful witness is just slightly more likely than chance (say, .3) to agree with the first truthful witness, this favors H by .3/.25 or about 1.2/1, whereas contradiction in that case favors ~H by 75/70 or only about 1.07/1. By this time this is not a very great disparity, but it is still present. The general point is this: It is only at high levels of expected agreement on the detail, conditional on H, that agreement favors H less than or even equal to the extent to which contradiction favors ~H. When fallibility is permitted for otherwise truthful witnesses, even to a modest extent, an asymmetry arises whereby agreement favors H more than disagreement favors ~H.

Of course, it would be far too simple, and obviously false, to say that making individual witnesses more fallible is a good thing for the confirmation of H, full stop. When the probability of hitting upon a particular number purely by chance is held constant, the limit of the effect discussed here is reached when an individual witness becomes no better than a chance indicator of the truth about the specific tire, even if the flat tire story in general is correct. This is, quite simply, because for purposes of this example we are treating the probability of hitting upon a given tire by chance as P(E|~H) = .25. Therefore, if, due to some background information, P(E|H) is also merely .25, the agreement of the second witness with the first about what tire went flat is of no value at all.[2]

Suppose that we do not hold constant the probability of a particular item of testimony by chance. Suppose, instead, that the probability of a certain report on a detail, selected by chance, goes down. Instead of asking about a tire, say that we are asking about a number from 1 to 100. The probability then that a speaker will pick any particular number at random is .01. Now suppose that we envisage a situation where the dominant ~H hypothesis leaves this number to vary at random. For example, suppose that we suspect that two witnesses are colluding to invent a robbery story while both stating that the alleged perpetrator was wearing a jersey of a certain kind. We know independently that such jerseys have numbers on them ranging from 1-100, inclusive. As in the tire story, we note that the witnesses agree on the main outlines of the story but then we separate them in order to ask, "What number was on the jersey?" assuming that, even if they are colluding on the story as a whole, they did not collude on this particular detail (sloppy collusion).

In this case we might want to allow quite a bit of fallibility for genuine, truthful witnesses, and hence a probability of more than .5 that they will disagree when forced to state a number. It would be plausibly more difficult for them to agree exactly on a number between 1 and 100 than for them to agree exactly on one of four tires, even if they are both real witnesses. Once again, we envisage a situation where the witnesses must be forthcoming. But even supposing that we made the probability only .4 that a second otherwise truthful witness gives the same exact jersey number as the first, this greatly swamps the .01 probability of agreement given ~H. The odds are 40/1 of agreement given H over agreement given ~H. In contrast, the probability that the second witness disagrees with the first given H is .6. The probability that they disagree given ~H is .99. A contradiction on this detail therefore favors ~H by odds of 99/60 or about 1.65/1, which is obviously much less impressive than the 40/1 odds by which agreement favors H.[3]

---

[2] A reviewer for this journal has pointed out the generalization that, when P(E|H) and P(~E|~H) are equidistant from 0.5, E favors H exactly as much as ~E favors ~H. This is true. It is, in effect, the same probabilistic fact as the one just noted: Namely, by the Bayes factor analysis used in this paper, when P(E|H) = P(E|~H), E favors H just as much as ~E favors ~H. Which is to say, neither E nor ~E favors either hypothesis over the other. The connection between the two points is as follows: For any hypothesis H and any univocal evidence E, P(E|H) and P(~E|H) must sum to 1. Therefore, P(E|H) and P(~E|H) are always equidistant from .5, in opposite directions. E.g. if P(E|H) is .7, P(~E|H) is .3, and so forth. Moreover, if P(E|H) = P(E|~H), then P(~E|H) = P(~E|~H). Therefore, P(E|H) and P(~E|~H) are equidistant from .5, in opposite directions, just in case P(E|H) = P(E|~H).

[3] This example illustrates an important part of the reason for my use of the likelihood ratio as the measure of confirmation in this paper. My goal is to focus on the force of the evidence in itself, independent of the prior probabilities of H and ~H. The ratio P(H|E)/P(H), a strong competitor in the popularity contest among measures of confirmation (see Fitelson 1998), is sensitive to the prior probabilities, which obscures the point in question. Suppose that in the case just discussed the prior probability of H were .99 in the distribution in question. Then, when E is the agreement of the alleged witnesses on the detail and the likelihoods are as stated—P(E|H) =.4, P(E|~H) = .01—P(H|E) ≈ .9997 and P(H|E)/P(H) ≈ 1.0098. When E is the witnesses' contradicting each other on the detail and we have P(H) = .99 and the stated likelihoods—P(E|H) = .6, P(E|~H) = .99—P(~H|E) ≈ .01639, and P(~H|E)/P(~H) ≈ 1.6393. In other words, if we use this measure with these priors, contradiction confirms ~H more than agreement confirms H. But this is a function of the priors, as can be seen if we keep the likelihoods the same and simply reverse the priors, so that the prior of H is .01 instead. Then, for E = agreement, P(H|E)/P(H) ≈ 28.776, while for E = contradiction, P(~H|E)/P(~H) ≈ 1.003. In other words, in that case, by that measure, agreement confirms H far more than contradiction confirms ~H, but this is because H rather than ~H has the low prior probability. In this paper I will sometimes compare the impact of differing evidence on the same hypothesis, for which the prior probability can

Again, this effect only holds so long as, and to the extent that, if H is true, the alleged witness has a better-than-chance probability of getting the number right. But when the probability of getting the same number on chance is quite low, yielding a low P(E|~H), this condition is fairly easy to achieve, even when the witnesses are fallible. Given the many factors that can bring about witness fallibility on a highly specific number, while the fallible witness is still a better-than-random indicator of the number, given H, the probability of a contradiction as a result of witness fallibility will often not be as epistemically significant as the sheer improbability that the second witness will say the same thing as the first purely on chance. This means that agreement on a difficult detail has a strong effect in favor of H; the alleged witnesses appear to be especially accurate. But direct contradiction on that same detail does not disfavor H nearly as strongly.

In the same context as the epigraph for this paper, Starkie remarks,

> It is here to be observed, that partial variances in the testimony of different witnesses, on minute and collateral points, although they frequently afford the adverse advocate a topic for copious observation, are of little importance, unless they be of too prominent and striking a nature to be ascribed to mere inadvertence, inattention, or defect of memory. (Starkie 1876, p. 830)

The argument thus far provides a probabilistic rationale for Starkie's point. For a difficult detail, fallible witnesses who agree are showing a striking degree of accuracy. Such agreement may not be very probable given H, but it is highly improbable given ~H. By the same token, contradiction on a highly specific detail is not all that unexpected and may even be more probable than not given truthful testimony to H, which means that it does not provide all that much evidence *against* H if it occurs.

This is already an epistemically interesting result, though a relatively simple one. Of course the specifics will vary greatly with empirical evidence concerning the probability that the witnesses agree on the detail if they know about the larger event described in H.

Thus far, in keeping with the original flat tire example, we have assumed only "sloppy collusion" given ~H—that is, that there is no dependence between the alleged witnesses on the specific detail (which tire or which shirt number) if the main event didn't happen. But matters become much more complicated when greater overall dependence, including dependence on the detail itself, is antecedently probable. There, as we shall see, situations can arise where exact agreement on a difficult detail is positively *detrimental* to the confirmation of H.

4. Contradiction confirms H? How it works when the witness is forthcoming

For this example I will define an hypothesis that I will dub the Echo Chamber hypothesis—EC. When EC holds, the second witness or document attesting to a main event described in H is guaranteed to echo whatever the first witness or document says. Given EC, contradiction is impossible. EC does not necessarily rule out the possibility that the second source will add to what the first source says (though I will make no specific use of this possibility), but he must say

be held equal. At other times I will compare the impact of different evidence upon hypotheses with different prior probabilities. So it would be impossible to illustrate all of the interesting effects I wish to discuss while using a measure that is sensitive to prior probability; by choice I am leaving prior probabilities out of consideration as much as possible.

at least all that the first says and not contradict it nor disconfirm it in any way. I introduce EC to show that contradiction on a matter of detail can be helpful to the confirmation of H and that exact agreement on some matter of detail can (surprisingly enough) disconfirm H when there is significant reason ahead of time to suspect a great deal of dependence between two sources.

To make things even more difficult, I will suppose that EC is possible even if H is true. This is a significant departure from the case of the students and their alleged flat tire, where the only point of collusion was to bolster a lie. One of the points of introducing the dependence hypotheses in this section is to attempt to model a literary case where we have two putatively historical documents that tell about some event described in H, but we do not know for sure how dependent they are. The author of the second document could be innocently copying the first document even if H is true. Perhaps he deems the first document to be a legitimate historical source. So copying (whether echoing or partial) would not need to be tied to an attempt to relate a known falsehood.

For this example I will also allow for the possibility of partial copying. Under this hypothesis, the second witness agrees with the first as to the main event, but he will not simply echo or copy him on the specific detail in question (which will be a number between 1 and 100). He may still happen upon that same number by chance, but not by copying or echoing. Once again, to make the scenario reflect a possible literary hypothesis, I will allow for this possibility both if the event actually happened and if it did not.[4] The most important point concerning partial copying is that, whether or not the main event in H happened, the second witness or document has *no independent access* to what actually happened. So, if the first document or witness is no longer available to the second person, he must guess at the number just as he would have to do if H were false.

Since I am allowing both Echo Chamber and Partial Copying even if the event in H occurred, it is important to be clear about the third possible scenario if H occurred: Independent access. Under this scenario, H occurred, and the second witness or author has *some* access to what occurred that is not causally or epistemically "routed" either through the first or through a common source other than reality. Either the second witness or author was present at the event himself or he has access to some testimony or document going back to a *separate source* (and hence back to reality) from that used by the first witness or document. This point was assumed in the flat tire story, since the students all claimed to have been personally present. By definition they were all witnesses to the event if it happened at all. Here that may not be literally the case (that both sources are eyewitnesses), but it is a possibility. Independent access is also possible even if neither one was a personal witness. They might, for example, have spoken with different witnesses or in some other way be the recipients of independent information about what really happened.

Note, too, that even though "partial copying" is mutually exclusive with "independent access," this does not mean that a witness or source who has independent access must be *utterly* independent of the first witness or source concerning all matters connected with the event. This is a tricky point that must not be missed, especially in literary contexts. "Partial copying" is defined so as to be distinct from independent access: Even if the event occurred, the partial copyist has no independent knowledge of what took place. However, the converse is not

[4] This does not have to mean that partial copying has the same probability given H that it has given ~H, especially not in the distribution representing a situation where both witnesses have attested to H. In fact, I will assume throughout that, once we have conditionalized on the fact that both documents or witnesses describe the main event in H, the absolute probability of both dependence hypotheses given H is lower than it is given ~H. This is because, if H is true, we could have independent attestation to a real event, whereas that is impossible if H is false.

necessarily true. Independent access need not mean that the second witness has *no* knowledge of what the first witness said or has not been *in any way* using the work of the first source or witness. I will not attempt to model that possibility explicitly, but it is compatible with the model, and I note it for the sake of completeness. It could, for example, influence the extent to which the sources' agreement on H itself has already confirmed H. My modeling concerns the agreement or disagreement on a detail *beyond* the fact that both sources agree on H. The major point for independent access is that the second witness does possess at least some separate knowledge and that, when it comes to the detail in question (the number), he writes or speaks as he does *because of* that separate knowledge, not out of a desire to copy, follow, or confirm the first document. He may even be wholly independent in the sense of not knowing what the first witness said at all, either as to the main event or as to the detail. Whether he is wholly independent or not, he does have some independent access which can influence him in what he attests on the further detail (the number).

Let us now suppose that we are considering a distribution after both witnesses or documents have attested to a fairly specific proposition H, asserting the occurrence of that "main event." We also have the explicit testimony of the first witness to a further specific detail involving a number between 1-100. As suggested before, suppose that this is a number on a shirt supposedly worn by someone involved in the incident. We have not yet conditionalized on the testimony of the second document or source concerning this number. Suppose that we keep the forthcomingness requirement in place: The second document must or will speak specifically to that number. Silence and indirect confirmation or disconfirmation alone are not options in this still somewhat simplified scenario.

Let us suppose that the dominant ~H hypothesis is that the second source is an echo chamber of the first source instead of a sloppy colluder (partial copyist). This, we will say, is a result in part of some background information that produced a high prior probability for complete dependence between the two sources. That probability has been further enhanced by the fact that both sources testify to H itself. In fact, we can imagine that the H itself is already specific enough that we can rule out the possibility that there is *no* dependence between the two sources if H is false. There is no way that the two sources both attested to H *entirely* independently, given that H did not occur. Hence, ~H is completely covered at this point by the two dependence hypotheses (Echo Chamber and Partial Copying), with Echo Chamber predominating at .8 conditional probability.

Given H, we are allowing for the possibility of both Echo Chamber and Partial Copying, but the possibility of Independent Access is a very live option and could fully or partially explain the fact that both sources attest to H. This is why Echo Chamber is not nearly as probable given H as given ~H.[5]

Suppose that we have the following conditional probabilities. EC stands for Echo Chamber. PC stands for Partial Copying. IA stands for Independent Access.

---

[5] I am not assuming that the evidence thus far about the main event has equal probability given Echo Chamber and Partial Copying, either given H or given ~H. This is important, since they would have to have the same proportions to one another given H and ~H respectively in the current distribution if they both had equal likelihood *vis a vis* the evidence thus far and if they both had the same prior probabilities before any of the current evidence was taken into account. Our illustrative numbers in this distribution are P(EC|~H) = .8 and P(PC|~H) = .2, but those same respective conditional probabilities given H are .5 and .1. 8/2 is not the same proportion as 5/1. But, since EC and PC did not constitute a partition of either H or ~H in the initial distribution, before *any* testimony was taken, then it is possible for them to stand in different proportions to one another in the current distribution given H and ~H.

$P(EC|{\sim}H) = .8$
$P(PC|{\sim}H) = .2$

$P(EC|H) = .5$
$P(PC|H) = .1$
$P(IA|H) = .4$

Notice that I am saying nothing at all about the prior probabilities of H and ~H in this distribution. My goal here is to develop a thesis about the *evidential force* of agreement and disagreement on a matter of detail between two sources or witnesses who might be attesting truly or falsely and who might be dependent or independent. For that purpose it does not matter whether H itself is antecedently probable or improbable. (This, again, is why I am using the Bayes factor as a measure of confirmation throughout this paper. See footnote 3.)

Given the forthcomingness requirement, agreement and contradiction are the only two possibilities. Contradiction, of course, is impossible on Echo Chamber. Given that, on Partial Copying and H, there is no independent access to anything that really happened, and given that, if H is false, the partial copyist is not going to copy the first document or witness and can agree (A) only by chance, $P(A|PC) = .01$. This probability is the same for both $P(A|PC \,\&\, H)$ and $P(A|PC \,\&\, {\sim}H)$.

As in the previous section, let us allow a serious possibility that one witness or other (or both) will err about this specific detail (a number between 1 and 100), even if the main event occurred, so that (given the forthcomingness requirement) contradiction is somewhat more probable than agreement. Suppose, then, that $P(A|IA \,\&\, H) = .4$ and $P(C|IA \,\&\, H) = .6$.

With all of these stipulations in place, here is the probability of contradiction given ~H:

$P(\text{contradiction}|{\sim}H) = (0)(.8) + (.99)(.2) = .198$

And here it is given H:

$P(C|H) = (0)(.5) + (.99)(.1) + (.6)(.4) = .099 + .24 = .339$

This yields the surprising conclusion that, in this scenario, contradiction on this rather difficult detail actually *confirms the event*, using the simple likelihood ratio as a measure, by a factor of about 1.7 to 1. The confirmation is not very high but is not nothing, either, and of course it is somewhat counterintuitive that contradiction confirms the main event at all.[6]

By the same token, exact agreement on the number in this scenario just slightly disconfirms the main event, by a factor of about 1.2/1.

$P(A|{\sim}H) = (1)(.8) + (.01)(.2) = .802$
$P(A|H) = (1)(.5) + (.01)\,(.1) + (.4)(.4) = .5 + .001 + .16 = .661$

This effect would be even more dramatic if we were insisting that the witnesses are completely independent given H—that is, that the dependence hypotheses are relevant only if the sources

[6] A contradiction in this complex scenario confirms H about as much as a contradiction confirms ~H in the simpler scenario in section 3 where we had a forthcomingness requirement and some degree of fallibility given H (namely, a probability that the second source agrees with the first on the specific number, given H, of .4).

attest falsely. But again, because I am trying to include possible literary scenarios of innocent documentary dependence, I want to allow for the possibility of echoing or partial copying given H as well as ~H.

How do these odd results come about? Are they just an artificial construct in an example, or are they significant? There are three main causes of the results in this section: First, there is the fact that agreement is guaranteed and contradiction is impossible given Echo Chamber. Second, there is the fact that Echo Chamber rather than Partial Copying is the dominant dependence hypothesis, in contrast to the flat tire example in which the professor assumes exactly the opposite—that colluding students have definitely not colluded concerning which tire went flat. The case modeled in this section is like a case in which the first student writes out a statement that is available to the others when they are asked, "Which tire?" Then, obviously, even if there was no flat tire, they do not have to pick a tire by chance. They can just agree with the first statement. Third, there is the fallibility of otherwise truthful witnesses given H. Even if the main event really occurred and the second source has independent information about it, there is the real possibility that the two sources will disagree about this specific detail. In various ways there is more "flex" built into H for disagreement on this detail.

While the modeling is of necessity somewhat artificial in order to make the example mathematically tractable, it is not out of touch with potential real-life examples. If we have strong prior reason to think that putative sources are dependent if the event didn't happen, and if that dependence would be of such a kind that they could avoid contradiction, we will usually expect them not to contradict one another if the main event did not occur.

If an event did occur, a second witness or source with independent access to the event may have a different honest opinion from the first one concerning the specifics of a detail. In fact, even if he actually knows what the first one has said, he may decide to say something different based on some other information; in other words, he may attempt to correct the first statement. Or suppose that the second honest source does not have the first one available for that detail. Given the requirement to be forthcoming, even an otherwise truthful witness may have to hazard a guess about the number based upon hazy memory. Independent access to the event combined with human fallibility could yield contradiction while slavish copying would not. If, then, the witnesses contradict each other on a difficult matter of detail, this has some value for H in virtue of the fact that it allays prior suspicions of problematic dependence.

Starkie, as quoted in the epigraph to this paper, is pertinent, for he noted this same practical result in court testimony:

> It has been well remarked by a great observer, that "the usual character of human testimony is substantial truth under circumstantial variety." It so rarely happens that witnesses of the same transaction perfectly and entirely agree in all points connected with it, that an entire and complete coincidence in every particular, so far from strengthening their credit, not unfrequently engenders a suspicion of practice and concert. (Starkie 1876, p. 831)

5. Contradiction confirms H? More complex examples

We can try to make the results more widely applicable (though the analysis is necessarily more complex) if we loosen the forthcomingness requirement and permit the second witness to remain silent or (which I will treat as the same thing) to say nothing relevant concerning the detail in question. We can also allow for the possibility that, without saying anything relevant *directly* to the detail, the second source says something else that is *indirectly* pertinent to that detail, perhaps without even realizing that he is doing so.

Suppose, for example, that we know by independent evidence that the type of shirt in the example already sketched (picked out by some wording on the shirt on which the witnesses agree as part of H) have a strikingly different color for different number ranges. Shirts with numbers from 1 to 10 are green, those from 11 to 20 are orange, and so forth. The second witness could indirectly confirm the first if the first stated that the number on the shirt was 85 while the second stated that the shirt was bright purple. That is, if our background evidence shows that a shirt of that type and number is indeed bright purple, the two sources indirectly confirm one another. By the same token, the second witness might state that the shirt was orange, which would indirectly contradict the testimony of the first to the number 85. Note that a color incompatible with the number *entails* the falsehood of the specific number, while a color compatible with the number does not *entail* the more specific number, though of course it makes it more probable than it is if all we know is that it is a number from 1 to 100.

When we allow for both indirect confirmation and silence on the part of the second witness concerning the specific detail, while continuing to allow for fallibility concerning the detail on the part of two otherwise truthful witnesses, we have the same effect that we observed in the previous section: Since Echo Chamber is the dominant dependence hypothesis, it becomes the subhypothesis of ~H that most needs to be defeated in order to confirm H. In this case, direct agreement by the two witnesses on a specific detail can actually *disconfirm* H (because it confirms Echo Chamber) while direct contradiction confirms H (because it disconfirms Echo Chamber). Silence and indirect confirmation have their own effects, interesting in their own right.

For this example, I will retain some of the conditional probabilities from the last section. EC stands for Echo Chamber. PC stands for Partial Copying. IA stands for Independent Access. Definitions of these subhypotheses are the same as in the last section.

$P(EC|{\sim}H) = .8$
$P(PC|{\sim}H) = .2$

$P(EC|H) = .5$
$P(PC|H) = .1$
$P(IA|H) = .4$

As before, I am modeling a scenario in which both witnesses have attested to H and in which H is sufficiently specific that, if the main event in H did not occur, some sort of copying or collusion must be taking place. Therefore, PC (partial copying) and EC (echo chamber) are now the only possibilities given ~H.

Now let us consider the following possibilities given each of these: Direct agreement (Witness or Source 2 says exactly the same thing as Witness or Source 1 about the detail of the number between 1-100), contradiction (Witness 2 says something incompatible with what witness 1 has said about the number), silence (Witness 2 says nothing either positively or negatively relevant to the number), indirect confirmation of the number (Witness 2 states a shirt color that is the only color compatible with the number in question, though it does not entail it). These possibilities are treated as mutually exclusive for simplicity. Witness 2 will either directly or indirectly confirm the number, but not both. This means that, since he must directly confirm it given Echo Chamber, the probability of his indirectly confirming given Echo Chamber is 0. These are all of the possibilities.

Suppose that we have the following conditional probabilities under ~H:

P(direct agreement|~H & EC) = 1
P(indirect confirmation|~H & EC) = 0
P(contradiction|~H & EC) = 0
P(silence|~H & EC) = 0

P(direct agreement|~H & PC) = .0045
P(indirect confirmation|~H & PC) = .04
P(contradiction|~H & PC) = .45
P(silence|~H & PC) = .5055[7]

Since these four are all the options, under ~H conjoined with each further supposition they sum to 1.

The conditional probabilities for silence, contradiction, indirect confirmation, and contradiction given (H & Echo Chamber) and (H & Partial Copying) are the same as they are for ~H. It is possible that a partial copyist would have either more or less tendency to venture to make statements about a number if H were false than if H were true, especially if he *knew* that he was colluding in a falsehood. But I am not trying to model any difference in boldness or forthcomingness given Partial Copying between H and ~H. This is partly because, in the case of literary dependence, the author of the second source may sincerely believe that H is true, based upon reading the first source, even if H is in fact false.

Suppose that we have the following conditional probabilities for H:

P(direct agreement|H & EC) = 1
P(indirect confirmation|H & EC) = 0
P(contradiction|H & EC) = 0
P(silence|H & EC) = 0

[7] The reader may wonder about the origin of these highly specific numbers. When one is trying to allow for four different possibilities which must all sum to one in a given portion of the distribution, and when one is trying to model particular epistemic intuitions, some numbers will appear over-precise so as to result in a coherent distribution. As explained in more detail below, for Partial Copying the goal is to model what we might picture happening if H is true but the second document has no independent access to the relevant events and is not going to copy the detail from the first document. (In this sense Partial Copying is similar to what was called "sloppy collusion," above.) The author might well be cautious on such matters of detail, leading to silence on this one. He might try to guess, leading to contradiction. I discuss below why I have made silence more probable than contradiction. He would be very unlikely to produce a delicate indirect confirmation of the detail. And he would be even more unlikely to give the exact same detail without copying it. Trying to model all of this for four possibilities is what produces such artificially specific numbers.

P(direct agreement|H & PC) = .0045
P(indirect confirmation|H & PC) = .04
P(contradiction|H & PC) = .45
P(silence|H & PC) = .5055

P(direct agreement|H & IA) = .2
P(indirect confirmation|H & IA) = .3
P(contradiction|H & IA) = .25
P(silence|H & IA) = .25

If H is true, we at least have an option for Witness 2 to have his own access to the events (Independent Access). Obviously, if the main event didn't happen, there is nothing for the second source to have independent access to. Independent Access is impossible given ~H but a live option given H.

Due to the loosening of the forthcomingness requirement, a second witness of the events is not required to state a particular number from 1 to 100. He is free to speak entirely about other matters, to state only the basic events of H and nothing more, or to focus on another issue such as the color of the shirt, which would be easier to notice and remember than a particular number.

If a second witness does state a particular number, he will have some fallibility if he is a genuine witness or has a genuine independent source of information. At the same time, there will be a better-than-random chance that a second source with independent access gets the detail right. If both witnesses state a particular number and are correct, their numbers will agree, and truthful witnesses do sometimes say exactly the same things. But with Echo Chamber dominating the ~H probability space, the odds are very high that, if H is *false*, they will agree exactly, stating the very same number on the shirt, even if the shirt and the entire scenario are unreal. As in the previous section, there is more looseness built into H than into ~H, since real witnesses can contradict one another on a particular point. And here there is in a sense even more looseness given H, since real witnesses can indirectly confirm one another instead or remain silent.

So what is the effect of these numbers? Consider the probability of direct agreement between the witnesses (Witness 2 says exactly what Witness 1 says about the specific number) conditional on ~H and H, respectively:

P(direct agreement|~H) = (1)(.8) + (.0045)(.2) =.8009

Due to the high probability of Echo Chamber given ~H, the probability of direct agreement is quite high given ~H.

Given partial copying and either H or ~H, the second witness may decide not to try to get too specific about the number on the jersey. There is no forthcomingness requirement, and he knows that if he guesses, he may say something that contradicts something that someone else has said. Since *ex hypothesi* in Partial Copying he is not going to copy the specific number from the first source (perhaps because he no longer has access to it), he may "go vague" on details, preferring silence. The probability that he will explicitly state the same number as the first witness given Partial Copying is therefore even worse than a chance guess, since he may not even try to guess. He has other options. So does a truthful witness with independent access; he, too, may choose to speak about something else. The overall probability of direct agreement given H is therefore even lower than it is when the second witness has to give a number.

P(direct agreement|H) = (1)(.5) + (.0045)(.1) + (.2)(.4) = .58045

The result here is rather similar to the one in the previous section: Direct agreement actually confirms ~H, though by a somewhat greater margin of about 1.4/1, due to the fact that the second witness has even more options given H than he did in the last section.

As before, contradiction confirms H by lowering the dependence between the hypotheses—disconfirming Echo Chamber.

P(contradiction|~H) = (0)(.8) + (.45)(.2) = .09
P(contradiction|H) = (0)(.5) + (.45)(.1) + (.25)(.4) = .045 + .1 = .145

Contradiction confirms H by odds of about 1.6 to 1. The confirmation is slightly less than in the example in the previous section with the forthcomingness requirement. Given H, a second, otherwise truthful source may simply not address the matter, or he may indirectly confirm the first via a different truthful detail (a point I will return to momentarily). The probability of contradiction given partial copying (on either scenario) is significantly higher than the probability of direct agreement, since of course it is much easier for the second witness who has no access to what happened to contradict the first than for him to say exactly the same thing by chance. Even if he speaks out on a matter like color, which seems safer, he may accidentally say something that entails the falsehood of what the first witness said, given the color coding of the shirts.

Silence on the matter of the specific number slightly favors H, though not quite as much (with these numbers) as contradiction.[8]

P(silence|~H) = (0)(.8) + (.5055)(.2) = .1011
P(silence|H) = (0)(.5) + (.5055)(.1) + (.25)(.4) = .15055

Why have I made silence somewhat more probable than contradiction given partial copying? Suppose that Source 2 has access to the statements of Source 1 but has no independent access to the events if they happened, and (obviously) no independent access to the events if they did not happen. But suppose that Source 2 is not an echo chamber—he is not going to echo everything that Source 1 says, and in particular he is not going to echo what Source 1 says about the specific number. This would be one variety of partial copying, in which Source 2 knows what is in Source 1 but is not going to echo it on this point. But in that case, would he contradict it? That would be enormously unlikely. If he were not going to echo the number in the first source, but he knew what that number was, he would be more likely to avoid the topic as much as possible.

[8] Because there are more than two exclusive and exhaustive possible outcomes, it is possible for more than one of them to favor H. When we have given such a detailed analysis of multiple possible outcomes, we cannot make epistemically enlightening generalizations about the effect of the mere *negation* of one outcome—e.g., "*not* silence." The negation of silence must favor ~H if silence favors H. "Not silence" does *very* slightly favor ~H, overall, due to the impact of exact agreement in further confirming copying. But "not silence" *includes* a possible outcome that significantly favors H on a Bayes factor analysis (indirect confirmation) and one that slightly favors it (contradiction).

Contradiction thus disconfirms the idea that the second document even has access to what the first document says about the number.

This point is particularly relevant when we are considering literary dependence hypotheses. Even if we think for some other reason that there is some dependence between them, if one literally contradicts the other, it is most probable that the dependence *does not extend to that particular point*. Perhaps the second author had set aside the first document or no longer had access to it when he made the contradictory statement. Otherwise, since he has no separate knowledge allowing him to correct the first source, one would expect that he would *at least* take the trouble not to contradict it.

The most interesting new point to emerge in this more complex scenario concerns indirect confirmation. The sort of indirect confirmation in view here is what has been called elsewhere an undesigned coincidence.[9] In an undesigned coincidence, one alleged witness or putatively historical document states one piece of information, which fits together in an apparently casual fashion with a piece of information from another source. In this case, suppose that Witness 1 (or Document 1) says that the shirt was number 85, Witness 2 (or Document 2) says that the shirt was purple. Our background knowledge tells us that shirts of the type that they are both talking about are purple for the numbers 81-90 but not for any of the other numbers between 1 and 100. Witness or Document 2, let us suppose, does not mention the specific number on the shirt, and Witness or Document 1 does not mention the color. There is therefore an appearance of casualness in the connection between them, and the connection is accurate, based on our other information.

If one of them does have access to the statements of the other and is deliberately attempting to fit his information together with those statements indirectly, he is doing it in so subtle a way that the connection might be overlooked by his target audience, especially if the audience does not know about the color connection. Moreover, the witness or author who mentions the color may not know about the color connection himself. These considerations are part of what causes the low setting of the probability for this indirect confirmation (undesigned coincidence) given Partial Copying. If the second person has access to the information in the first source, it is a highly complex hypothesis that he would deliberately try to fit his statements with it in so indirect a way. If, on the other hand, he does not know or does not remember what the first source says about the number, it is unlikely that he would just happen upon the color that fits with the specific number named by the first, since he has no independent access to the truth and is not copying this particular detail.

If H is true and the second witness or document really does have independent access to real events, then he may not mention the specific number. But he may mention the color, which would be (plausibly) both easier to notice and easier to remember. Then, because he is speaking truth about real events, his true detail fits with and confirms the detail stated by the other source.

I do not claim that such an indirect confirmation is absolutely more probable than not (>.5) given H and Independent Access. After all, there are plenty of things that a truthful, independent source could mention that would not independently confirm the number. I have set this probability at a modest .3. Even so, the effect on the confirmation of H is quite notable.

$P(\text{indirect confirmation}|{\sim}H) = (0)(.8) + (.04)(.2) = .008$
$P(\text{indirect confirmation}|H) = (0)(.5) + (.04)(.1) + (.3)(.4) = .004 + .12 = .124$

---

[9] The term was originally coined in the context of Pauline studies. See Paley (1850), pp. 1-8. For a recent Bayesian treatment of undesigned coincidences in detail, see L. McGrew 2020.

Using these numbers, indirect confirmation is more likely given H than given ~H by odds of about 15.5 to 1. Even though I did not assume that the two sources *must* have independent access if H is true, the plausibility, given H, of independent access and indirect confirmation means that the big loser here is ~H.[10]

6. Epistemological implications

In the complex scenario just considered, indirect confirmation (undesigned coincidence) gives us the best of both worlds as far as confirming H goes. On the one hand, it is impossible given Echo Chamber, and it is unlikely given Partial Copying. On the other hand, it is more likely than contradiction given H and Independent Access. This is partly a function of the fact that it is easier for a second truthful witness or document with independent access to the event to confirm the number indirectly than to get it right to the exact digit. It is the kind of thing that might well happen if we had two independent (even partially independent) witnesses to a real event.

There is nothing counterintuitive about the fact that indirect confirmation of a detail by a second source confirms H. What is notable is that it can confirm H even when, due to heavy suspicion of collusion or some other form of dependence, direct agreement *disconfirms* H. An undesigned coincidence disconfirms those suspicions of dependence while confirming the truthfulness of both witnesses in a way that doesn't have the drawbacks of exact agreement.

This effect would be even more striking if we were not allowing for the possibility of Echo Chamber and Partial Copying if H is true—that is, if we were treating them solely as subhypotheses of ~H. But even when it is possible for a second source to be innocently dependent (partially or wholly) upon a first source in reporting a true main event (as in a case of historical literary dependence), the statement of details that indirectly fit together confirms both that the second source has some independent access to the events and that the main events did occur as both sources tell us.

One way to see that indirect confirmation is especially valuable in these cases is to consider the effect of iteration. Even without spelling out another conditionalization of the same kind in detail, we can see that conditionalizing on another contradiction with these conditional probabilities in place would not confirm H a second time. Why not? Because after the first contradiction, Partial Copying is the only hypothesis left given ~H. Echo Chamber has been completely ruled out by the first contradiction. But Independent Access would of course remain a possibility given H. Contradiction is less probable given H and Independent Access than given ~H and Partial Copying. After one contradiction, Partial Copying is the *only* possibility given ~H. So the probability of a second contradiction *must* be higher given ~H than given H. This mirrors the intuitive epistemic point that we would not actually try to confirm H by showing repeatedly that the alleged sources or witnesses contradict one another on matters of detail!

---

[10] Scenarios involving all four options (exact agreement, indirect confirmation, silence, and contradiction) are so complex that a distribution is even possible in which both exact agreement *and* contradiction disconfirm H, counterintuitive as that may seem. This is a result of the fact that there are several possible outcomes given H that take portions of the probability space. Suppose that Partial Copying rather than Echo Chamber is dominant, though both are possible. Suppose that P(Partial Copying|~H) = .8, P(Echo Chamber|~H) = .2, P(Partial Copying|H) = .2, P(Echo Chamber|H) = .05, and P(Independent Access|H) = .75, and all other conditional probabilities remain as stated above in section 5. Then both contradiction and exact agreement confirm ~H somewhat, though the confirmation from exact agreement is negligible. Indirect confirmation remains the strongest evidence for H. Proof omitted.

In fact, once strong dependence (Echo Chamber) has been ruled out, we revert to what seems intuitively to be a more "normal" epistemic situation in which contradictions between witnesses disconfirm H. Even if Echo Chamber is not *entirely* ruled out, and even when indirect confirmation and silence are permitted, contradiction disconfirms H and exact agreement confirms H when the probability of Echo Chamber falls low enough. For example, if after both witnesses have attested to H and the first witness has attested to the detail, P(Partial Copying|~H) = .9 and P(Echo Chamber|~H) = .1 and P(Partial Copying|H) = .1 and P(Echo Chamber|H) = .011 (with the rest of the probability given H going to independent access), and all other conditional probabilities remain as stated in section 5, exact agreement does confirm H somewhat while contradiction does disconfirm H somewhat.[11] As discussed in earlier sections, *how much* they disconfirm H will depend on other factors such as how difficult the detail in question would be for real witnesses to notice or remember and how improbable it is for two witnesses to agree upon that detail by chance, which affects the probability of both testimonies given ~H.[12]

There is, however, no reason why a second indirect confirmation should not confirm H again, since indirect confirmation of a detail is more probable if both sources have independent access to real events than if the second source has no independent information at all. The fact that contradiction can sometimes confirm the main event is an oddity that arises from major concerns about exact copying. By contrast, indirect confirmation of details *continues* to allay concerns about copying or collusion while simultaneously continuing to confirm H in a "normal" fashion—by showing that both witnesses appear to get things right. In this sense, indirect confirmation is often a win-win, having epistemic advantages for H over exact agreement, since it helps to disconfirm dependence, and certainly over contradiction, which confirms H only in highly specific circumstances that seem rather unnatural.

Formal epistemologists realized some time ago that greater coherence of the contents of testimony is not *per se* truth conducive for the conjunction of those contents. This is the case even when we stipulate that the testimonies are "screened off" from one another by the truth or falsity of their own contents—in other words, that they are independent given their own contents. (See Olsson 2005, pp. 134ff, pp. 211ff, Bovens and Hartman 2003, pp. 19ff, Shogenji 2013, L. McGrew 2016b, pp. 331-341.) When we do not stipulate independence, the problems with arguing that coherence of contents is truth conducive become especially acute, since in that case exact agreement of contents could indicate "bad" dependence (Huemer, 2007).

I have argued elsewhere that the intuition that coherence of contents is truth conducive arises from the recognition that, when two sources say very similar things, they are normally both "pointing toward" the same hypothesis, i.e., individually confirming it. If we carefully define the hypothesis to which multiple lines of evidence point and stipulate that the lines of evidence are independent given its negation and not negatively dependent given its truth, then the fact that they all point in the direction of that hypothesis means that they form a cumulative case for it (L. McGrew 2016b, pp. 337, 342-349). This is, all else being equal, a good thing for the confirmation of that hypothesis. But it still does not follow that we can develop a measure of their similarity (coherence of contents) that makes them *in general* stronger evidence for some H

[11] Proof omitted.

[12] With these conditional numbers in place, something similar is true of silence. Silence absolutely rules out Echo Chamber but is more probable given Partial Copying than given Independent Access. Hence, once we have one case in hand where the second witness does not address the detail at all, a further instance of the same kind will necessarily disconfirm H at least somewhat.

as they become more and more similar in content. Evidential strength is a far more complicated matter than that.

The results given here push back yet again against the idea that similarity of contents is truth conducive by showing that contradiction itself can be confirmatory in situations when we already strongly suspect dependence given ~H. Moreover, even in cases where exact agreement confirms H and contradiction disconfirms H (the more intuitive result), we should not assume that the confirmation and disconfirmation are proportional. The disconfirmation from a contradiction may be relatively trivial while the confirmation provided by agreement on that same specific detail is very high. This is especially likely to happen when the probability that multiple sources agree on the detail given ~H is low.

By expanding the complexity of the examples to allow for both silence and indirect confirmation, we can see further that what confirms H best of all is just the right *blend* of similarity and difference. Here I have modeled that blend by envisaging an indirect connection, an undesigned coincidence, that simultaneously confirms details of the two testimonies while disconfirming dependence. In that way, we see what Starkie calls "substantial truth under circumstantial variety." We have reason to believe that both testimonies are true, though they are not identical. They fit together because they are both taken from truth.

Though some of these results may seem counterintuitive, the explanations above show that they all have an intuitive rationale. What they challenge is not so much epistemic intuition generally as a hasty assumption about the value of identical testimony and the negative force of contradictions, even real contradictions when the testimonies cannot be reconciled. When contradictions in details of the narrative are a result of such normal causes as "inadvertence, inattention, or defect of memory," they confirm what we already know even about truthful, reliable, human witnesses who are not trying to change the facts—namely, that they are not infallible. This sort of fallibility may have little tendency to undermine the main story on which witnesses agree; in special circumstances, evidence of honest error can even confirm the main story by disconfirming some type of dependence that we have not otherwise ruled out. But better still are those undesigned coincidences that provide variety while leaving open the possibility that *everything* the sources relate is true.[13]

[13] My thanks to Timothy McGrew for useful conversations in which my ideas on these issues were tested and clarified.

Works Cited

Bovens, Luke and Hartmann, Stephan (2003), *Bayesian Epistemology*. Oxford: Oxford University Press.

Earman, John (2000), *Hume's Abject Failure.* Oxford: Oxford University Press.

Fitelson, Branden (1998), "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity," *Philosophy of Science* 66, Supplement, "Proceedings of the 1998 Biennial Meetings of the Philosophy of Science Association," Part I, pp. S362-S378.

Holder, Rodney (1998), "Hume on Miracles: Bayesian Interpretation, Multiple Testimony, and the Existence of God," *BJPS* 49:1, 49-65.

Huemer, Michael (2007), "Weak Bayesian Coherentism." *Synthese* 157:337-346.

McGrew, Lydia (2016a), "Evidential Diversity and the Negation of H: A Probabilistic Account of the Value of Varied Evidence," *Ergo* 3 (10):263-292.

McGrew, Lydia (2016b), "Bayes Factors All the Way: Toward a New View of Coherence and Truth," *Theoria* 82:329-350.

McGrew, Lydia (2020), "Undesigned Coincidences and Coherence for an Hypothesis," *Erkenntnis*, 85 (4):801-828. (First published in on-line first format, 2018.)

McGrew, Timothy (2014), "The Argument From Silence." *Acta Analytica*, 29 (2):215-228.

Olsson, Eric (2005), *Against Coherence*. Clarendon Press, Oxford.

Paley, William (1850). *Horae Paulinae, or the truth of the scripture history of St. Paul evinced by a comparison of the epistles which bear his name with the Acts of the Apostles and with one another*. Ed. T. R. Birks. London: The Religious Tract Society.

Shogenji, Tomoji (2013), "Coherence of the Contents and the Transmission of Probabilistic Support." *Synthese* 190:2525-2545.

Starkie, T. (1876), *A Practical Treatise of the Law of Evidence*. Philadelphia: T. & J. W. Johnson & Co.