

[Penultimate Version; the published version is in 2018 *Proceedings and Addresses of the APA*. If you don't have access to the APA website, please feel free to email me for a pdf of the published version: david.christensen@brown.edu]

On Acting as Judge in One's Own (Epistemic) Case **Marc Sanders Lecture**

David Christensen, Brown University

Introduction

“Nemo iudex in causa sua,” we are told: no one should be a judge in their own case. But while this may be a good rule to follow in legal proceedings, its epistemic analogue would be harder to uphold. In fact, we're often put in a position where we have no choice but to judge our own epistemic performance. We're put in this sort of position when, for example, we form the opinion that a female candidate's qualifications are slightly less good than her male competitor's—while aware of strong evidence that we're likely to undervalue women's CVs relative to men's. Or when we form an opinion about the results of some economic policy that's tightly connected to our passionate political views—while aware of strong evidence that political passions frequently distort people's reasoning on this type of matter. A small-plane pilot is put in this position when she's deciding whether she has enough fuel to make it to an airport a bit further away than her original destination, while aware that her altitude makes it likely that she's affected by hypoxia, which notoriously affects this sort of judgment while leaving its victims feeling totally clear-headed. A medical resident is put in this position when he forms an opinion about the appropriate drug dosage for a patient, while aware of strong evidence that he's been awake so long that his thinking about appropriate dosages is likely to be degraded. And many of us are put in this position when we form an opinion on some controversial issue while aware that others—who share our evidence and who seem as likely as we are to form accurate beliefs on the basis of such evidence—have reached a contrary opinion.

If we're serious about forming accurate opinions in any of these cases, it would seem that we must take seriously the evidence of our possible epistemic malfunction. This means, in the first place, assessing our likely reliability on the matter in question. And in many cases, it will also involve using that assessment to modify our confidence about the matter in question. I will assume here that this is true even in cases where our direct thinking about the matter in question happens to be perfectly fine—when it hasn't actually been distorted by sexism, political passion, or fatigue. So, for instance, even if a sleep-deprived medical resident has in fact figured out the dosage for his patient correctly, his confidence that the dosage is correct should be lowered by his knowledge of the cognitive effects of fatigue.¹

I want to look here at the first step—the assessment of one's own cognitive reliability—, which is where the awkwardness of acting as judge in one's own case arises. It seems clear that one must base this assessment on one's evidence, and that one must use one's own cognitive faculties to

¹ Though many would agree, this is not an uncontroversial assumption; for dissent, see, e.g., Kelly (2005), Lasonen-Aarnio (2014), or Titelbaum (2015).

do so. And it seems clear that the evidence on which this reliability assessment is based will often include evidence about the specific situation one is in: It's our medical resident's evidence that he's fatigued *today* that requires his estimate of his own reliability to be lowered. And knowing whether my disagreeing friend is drunk or not *now* affects the degree to which her disagreement is evidence against the reliability of my thinking about the disputed issue.

But some ways of invoking situation-specific reliability evidence don't seem kosher—precisely because of the way that they fail to separate the judger and the judged. For example, in the medical resident's case, it would seem irrational for him to be confident that his determination of the correct dosage had not been distorted by fatigue, on the grounds that the data in his patient's chart really do, after all, support prescribing that particular dosage. Or if my friend—whom I know to be more reliable than I am at weather prediction, and who has studied the data as carefully as I have—disagrees with me about the chance of rain tomorrow, it would seem irrational for me to dismiss her disagreement on the grounds that, after all, the meteorological data support my prediction, so it must be she who misinterpreted the data this time. In each case, it seems irrational for the agent to assuage the worry about their own reliability in the way described, even if the medical or meteorological data do in fact support the opinion that the agent initially reached.

For this reason, various people have proposed that the agent must assess her reliability in a way that's *independent* of the reasoning that the reliability evidence called into question.² So-called *independence requirements* have been stated in various ways—many of them pretty vague. But the basic idea is something like this: When our medical resident assesses his reliability, in light of the information that he's been awake for 27 hours, he should assess it in a way that's independent of (at least) his reasoning from the data on his chart to the correctness of a particular dosage. Or when I assess the likely reliability of my weather forecast, in light of my friend's disagreement, I should do so in a way that's independent of my reasoning from the meteorological data to the correctness of a particular prediction. So while the agent must in the end serve as judge of her own reasoning, the sort of independence wanted is aimed as separating the agent's judging of her reliability from the very thinking whose reliability is being judged.

Independence requirements have mostly been discussed as underlying Conciliatory accounts of disagreement. But the motivation behind them—avoidance of a certain kind of question-begging—is far more general. It applies equally in cases involving worries about sexist bias, political passion, hypoxia, and so on. In general, it seems plausible that some sort of independence requirement must be a part of any satisfactory account of the rational import of “higher-order evidence”: roughly, evidence that constrains our beliefs by way of bearing on the reliability of our thinking about the subject matters of our beliefs.³

² Much of the discussion of so-called independence requirements has focused on the disagreement problem. And the requirements have sometimes been formulated in ways that won't match what I'll discuss here. In particular, disagreement discussions sometimes are stated in a way that focuses on reliability-assessments of the disagreeing friend, not of the agent. I think that the root issues are the same as the ones discussed here, since assessments of the reliability of a disagreeing friend are important in large part because they bear on the reliability of the agent's own thinking on the given matter. I will not spend time here on delicate exegeses.

³ I've tried to motivate Independence principles in Christensen (2010, 2011)

Not surprisingly, this idea is problematic in several ways, and several writers have offered arguments against the very idea of independence principles. There are apparent counterexamples. And there are worries about whether we have sufficient motivation for precluding an agent from basing her judgment on evidence which she plainly has available to her—evidence which she may well be assessing impeccably. Here, I would like to examine some of these arguments, with an eye towards seeing how they may be resisted. This will, I think, give us some insight into how independence requirements should be understood. My hope is that this can take us some way toward meeting what I see as the biggest hurdle facing independence principles: the difficulty in giving a plausible formulation of such a principle with any decent degree of precision. I'll sketch some of the difficulties facing this project at the end.

1. Apparent Counterexamples

Let us begin with a very rough formulation of the relevant sort of Independence principle. It is inadequate in various ways, but a simple formulation will be adequate to discuss objections.

INDEPENDENCE: The rational response to higher-order evidence bearing on an agent's first-order thinking about P depends on an assessment of the reliability of her thinking about P—an assessment that's independent of that thinking.

A common sort of purported counterexamples to Independence principles involves cases where an agent forms an initial belief based on extremely strong evidence, and begins with a very high rational credence in the claim at issue, and then receives some higher-order evidence suggesting that her thinking about that claim may be unreliable. The worry has mostly been discussed in the context of disagreement, so the examples involve the agent's generally reliable friend disagreeing with such a claim. In many cases, it's plausible that two things may both be true: (1) that before learning of her friend's disagreement, the agent would rationally have thought her friend and herself equally likely to reach accurate opinions on matters of this type; and (2) that after hearing her friend's opinion, the agent should not lose much confidence in her initial opinion. The worry, then, is that in virtue of (1), the agent must dismiss the threat to her reliability in a way that's based on her reasoning to her initial opinion—which is exactly what independence principles are supposed to prevent.

For example, suppose my friend and I have worked very carefully through the same elementary addition problem. I've checked my answer, 17, in a couple of different ways with pencil and with a calculator. So I'm extremely confident in my answer. Then my friend, who has also been writing with a pencil and using a calculator, announces that she has arrived at 13! Even though I had no antecedent reason to think that my friend would be less reliable in this calculation, it seems that I should not lose a lot of confidence in my answer—instead, I should lose confidence in my friend's reliability on this occasion. If asked why I lost confidence in my friend's reliability on this occasion, I might well say, "She got 13, but the answer's 17." In other words, I seem to be assessing my friend's

reliability—and thereby assuaging the potential worry about my own reliability—in a way that’s based on the very reasoning whose reliability I’m trying to assess.⁴

However, it has been argued that there is a way of understanding such cases that’s fully consistent with independence requirements. Independence requirements, properly understood, allow for the following sort of response in the our math case: “It’s incredibly unlikely that two agents who were feeling clear-headed, had no memories of recent drug-ingestion or psychotic episodes, and were genuinely working through a simple arithmetic problem using repeated pencil-and-paper and calculator checking, would end up with incompatible answers. So it’s overwhelmingly probable that one of us is feeling fuzzy-headed, has some other evidence of likely cognitive malfunction, hasn’t really been working carefully with the pencil and calculator, or (most realistically, perhaps) is being insincere for a joke or some other reason. Since I know that none of those apply to me, it’s likely that one of them applies to my friend. So the answer my friend announced is unlikely to be correct, and I needn’t take her disagreement as strong reason to doubt the reliability of my own reasoning.”⁵

This response depends on interpreting “independent of the agent’s thinking about P” in a particular way. The reliability assessment may take into account *facts about the agent’s reasoning*, as long as it doesn’t *rely on* that reasoning, in the sense of depending on that reasoning’s correctness. One can see the distinction at work here by noticing that one could well make the relevant assessment in advance of seeing the arithmetic problem at all—one should expect in advance that if one soberly and carefully goes through a simple problem, only to have one’s friend announce a disagreement, it is very likely that one’s original answer will be correct. And of course, in the above presentation of the example, our intuition that I should not lose confidence could not have depended on the correctness of my initial reasoning: I never said what my reasoning was, or what the arithmetic problem asked.

This sort of response, which depends on clearly distinguishing between (a) assessing the reliability of some bit of one’s thinking by relying on that thinking, and (b) assessing the reliability of some bit of one’s thinking on the basis of facts about that thinking, will be crucial in examining some more recent purported counterexamples to Independence. For ease of reference, let us call this basic strategy the Facts About Reasoning (FAR) approach. With this in mind, let us turn to some more recent purported counterexamples to Independence principles.

a. Tricky Question Disagreements

Michael Arsenault and Zachary C. Irving (2013) argue that Independence fails in cases involving the following sort of disagreement:

Tricky Question: I’m thinking about the answer to a multiple-choice problem on a math test of the sort that I know sometimes utilizes “decoy answers”: wrong answers that seem right at first. As I think about problem 8, a problem involving dice, I first think that C is the right answer. Then I have an “aha!” moment, and decide that C was a decoy answer—it

⁴ For early versions of this type of objection, see Lackey (2010) or Sosa (2010).

⁵ This is a highly condensed version of the account given in Christensen (2011, 9ff), which draws on Lackey(2010), though Lackey does not intend to be defending Independence principles.

requires double-counting the 2—and that answer A is really correct. My friend, who's equally reliable at this sort of test, tells me she thinks that the answer is C.

As Arsenault and Irving (A&I) point out, it does not seem that I should lose much confidence in my answer—I should conclude that my friend has fallen for the decoy. But in dismissing my friend's disagreement, they argue, I use my first-order reasoning about the problem to conclude that 8 is a tricky question with decoy answer C. So I must violate Independence.

Now as A&I are aware, the FAR approach can also provide an Independence-respecting way of dismissing my friend's answer. After all, I can rely on the fact that I *seem to have recognized* a decoy answer and avoided it, while my friend picked the answer that *seemed to me* to be the decoy. Surely, in cases fitting that general description, I'd expect that I'd be much more likely to be right. Indeed (as we can see from the above description of the case), I could even have said, in advance of seeing problem 8 at all, that if I find myself seeming to recognize a decoy that my friend chooses, I'm much more likely to be correct than my friend is.

Nevertheless, A&I argue that the Independence-violating response is more intuitive:

We think our interpretation better captures how you'd intuitively dismiss your peer... When your peer announces that she got C, it would be perfectly natural to appeal to the first-order mathematical trick you think she's fallen for. For example, you might think to yourself, 'Aha! That sucker double-counted 2!' You'd feel no pressure to retreat to the second-order fact that you seem to [have avoided a decoy answer]. (188)

In support of the claim that the Independence-violating train of reasoning is legitimate, they point out that if I wanted to make sure my dismissal of my friend's belief was correct, I'd naturally double-check my own first-order mathematical reasoning, not double-check to see if it really did *seem to me* that I'd found a decoy answer.

In examining this argument, let us first consider the point about double-checking. The fact that I wouldn't look for evidence about how things seemed to me is obviously explained by the improbability of being wrong about whether C really seemed to me to be a decoy. But the more interesting question is about the significance of going again through the first-order mathematical reasoning. Why, we might ask, would it be rational to be more confident in my answer after rechecking my reasoning (assuming I get the same result the second time)? Not necessarily because I have any new first-order mathematical reason to believe that C is a decoy—I may well have simply carefully retraced the very same steps I did the first time. However, in rechecking I do get new *higher-order* evidence backing up the reliability of my original reasoning. (Again, one can see this by thinking about doing the reasoning in advance of actually doing it: if one reasons one's way to an answer to a problem, and then rechecks one's reasoning and gets the same answer, it's more likely that one's initial reasoning was correct.) So the FAR approach actually seems to explain quite naturally why I would double-check my first-order reasoning: having gone through the reasoning again, I'm rationally more confident in my answer. And the explanation does not involve becoming more certain that it really *seemed* to me that C was a decoy.

It's also worth noticing one reason that we should be careful in drawing conclusions from the naturalness of my explaining my friend's disagreement by imputing a particular mathematical error to her. Suppose that I religiously observe Independence. In thinking out loud about what to believe about the math problem in light of my friend's disagreement, I might reason follows: "I first got what seemed like the correct answer, C, then saw an apparent subtle problem with it—it seemed to involve double-counting 2—and concluded that A was right instead. My friend believes that C is the right answer. In situations when I reject an apparent decoy and my friend thinks that that same answer is correct, I'm usually right. So probably, I'm right this time. A is the correct answer, and C is a decoy—*my friend, poor sucker, double-counted 2!*"

The point is that, having dismissed my friend's disagreement via FAR, I *am* then in a position both to retain confidence in my original answer, and also to maintain confidence in the reasoning behind it—including that C is a decoy made plausible by double-counting 2. Nothing here violates Independence. Independence comes in at the stage of deciding how much to trust my own reasoning. If the verdict is good, I may then rely on that reasoning in various ways. This most obviously includes retaining my belief that A is the correct answer. But it also includes explaining my friend's having fallen into error.

A structural analogy may help make the point clearer. Suppose I have an excellent almanac, which tells me that the population of Barcelona is greater than 1.6 million. My friend Jocko has an almanac which says that the population of Barcelona is less than 1.6 million. If our almanacs are equally current, and produced by equally reputable publishers, I certainly cannot conclude that Jocko's almanac is wrong in this case simply because it disagrees with mine. But suppose I know that Jocko's almanac was produced by the Acme Company, which pays high-school students to gather its information by asking random Reddit threads. In that case, I should not lose confidence in what my almanac tells me. And, in that case, it would be perfectly natural for me to say to Jocko, "Your almanac is wrong about the population of Barcelona. I know from my almanac that it's over 1.6 million". Once I dismiss the disagreement provided by Jocko's almanac on grounds that are independent of what my almanac says about Barcelona, I can both believe what my almanac says, and use that to support the claim that Jocko's almanac is wrong about the population of Barcelona in particular.

Coming back to the case of peer disagreement, it's worth noticing that we do not always distinguish clearly between (1) explaining why I can disregard my friend's belief, and (2) explaining why she reached the wrong belief. Independence precludes me from relying on my initial reasoning for the first of these explanatory tasks, but once that is done, Independence poses no obstacle to my relying on my initial reasoning for the second. So in situations where I do legitimately dismiss my friend's disagreement, the intuitive legitimacy of my attributing a particular error to my friend does not, I think, give much support to the claim that Independence-violating reasoning is legitimate.

Is there a way, then, of getting a firmer intuitive grip on the issue? At this point, we have identified two lines of reasoning in support of dismissing my friend's disagreement: one which violates Independence by relying on the relevant first-order reasoning, and one which respects Independence by focusing on facts about the reasoning (in this case, my seeming to discover the decoy answer). A natural way to think about the legitimacy of these lines of reasoning is to think

about cases where these two sorts of factors come apart. So, for example, we might leave the first-order reasoning in A&I's case unchanged, but change the higher-order evidence in a way that affects the legitimacy of using the seeming decoy-discovery to dismiss my friend's opinion. For example, consider:

Extra-Tricky Question: I'm thinking about the answer to a multiple-choice problem on a test of the sort that I know sometimes utilizes decoy answers. In fact, I know that it not only uses questions that have ordinary decoy answers, it also uses an equal number of questions that involve "double-decoy answers": these are answers which are likely to seem like decoy answers, even though they're correct! (The tempting-but-unsound reasoning in these problems leads to the conclusion that a certain answer is a decoy.) I've done tests like this with my friend many times. When I think I've avoided a certain (regular) decoy answer, and my friend announces belief in that same answer, my friend has been right as often as I have.

Today, as I think about problem 8, a problem involving dice, I first think that C is the right answer. Then I have an "aha!" moment, and decide that C was a decoy answer—it requires double-counting the 2—and that A is really correct. Then my friend tells me she thinks that the answer is C.

In this case, given my track record in similar disagreements with my friend, it seems clear that I may not dismiss my friend's opinion. But notice that the Independence-violating reasoning A&I offer me in the original Tricky Question case is just as applicable here. So if it were legitimate there, it's hard to see why I could not similarly dismiss my friend's disagreement in the second case. By contrast, the Independence-respecting train of reasoning that allowed me to dismiss my friend's disagreement in the first case does not apply in the second case. In Extra-Tricky Question, my seeming discovery of a decoy does not have the same higher-order implications for my likely correctness as it did in the original case.

This strongly indicates that it is only the Independence-respecting train of reasoning that makes dismissing my friend legitimate in the original case. So it seems to me that we need not rely on questionable intuitions about which of the two trains of reasoning in the original case are legitimate. We can test for legitimacy using other cases. And when we do, it seems to me that the Independence-violating train of reasoning fails the test.

b. Asymmetrical Evidence Disagreements

In a sustained critique of Independence principles, Thomas Kelly (2013) offers a different sort of counterexample:

Holocaust Denier: I believe, on the basis of solid historical evidence (E), that the holocaust occurred (H). I meet Denier, who informs me that he disagrees. But I know that he's ignorant of much of the historical evidence on which I base my belief.

As Kelly points out, I should not lose confidence in my belief in this case. And Kelly then claims that in dismissing the Denier's belief, I may legitimately reason in the following way: "E, and E strongly supports H. But Denier is ignorant of E. So Denier is ignorant of evidence that strongly supports H. Therefore, Denier's opinion about H is untrustworthy." And this reasoning would violate Independence.⁶

Now Kelly is well aware of the FAR strategy, and acknowledges that it would work in some asymmetrical-evidence cases. So, for instance, suppose that I have a strong apparent memory of having eggs for breakfast, and my friend—who was not with me this morning—has a much lower degree of confidence in my having had eggs than I do. Clearly, I should not conciliate. But as Kelly notes, this case does not put pressure on Independence, since I can dismiss my friend's low credence without appealing to the evidence on which my belief is based. I'd reason as follows: "I have vivid memories about what I had for breakfast. My friend, not having been there, has no such memories. So his opinion is lacking in epistemic credentials." The key here is that I'm not assuming, in the course of dismissing my friend's belief, that my evidence really does support strong confidence in my having had eggs. I'm just relying on my having *some vivid apparent memories about what I ate*, which would be expected to reliably indicate what I really did eat.

The obvious question, then, is whether a similar approach can be taken to the Holocaust-Denier case. Kelly thinks not. He argues that my dismissal of Denier's opinion rests crucially on my assessment of E as *strongly supporting H*. He acknowledges that in certain versions of the case, an Independence-respecting way of dismissing Denier will be available: For example, if E came from my visiting the Holocaust Museum, and I know that Denier has not been exposed to any similarly reliable source of information about H, I'd have an Independence-respecting way of dismissing his disagreement. But Kelly claims that in his original case, there is no way for me to dismiss Denier's belief that doesn't rest on my claim that E strongly supports H.

But why should we think this? Kelly presents his case in schematic form, without describing the sort of evidence that I have and Denier lacks. In the one detailed description we have of a case where E can be identified—the Holocaust Museum case—it's easy to see how the exact sort of Independence-respecting reasoning we used in the eggs case can be employed to discredit Denier. If we changed the source of evidence from museum-visits to college courses, or independent historical study, or exposure to experts, it does not seem that the move would become any harder to make. So without seeing a concrete filling-out of the Holocaust-Denier case that would somehow preclude a similar move, I don't see why Holocaust Denier would pose a real difficulty for Independence.^{7,8}

⁶ I've condensed Kelly's presentation of the case, and the dismissal argument, for simplicity. The precise version of Independence he's targeting requires that my assessment of Denier's belief in this case be independent of my "assessment of the considerations that led [me] to believe" that H. Applied to the reasoning in the text, this would target the step "E strongly confirms H". I should also note that the particular Independence principles Kelly considers are particularly formulated to apply to disagreement, rather than to higher-order evidence in general. This won't matter right away, but will be important below.

⁷ See Matheson (2015, 95ff) for a similar take on Kelly's example.

Now I should point out that while Kelly believes that Holocaust Denier is a counterexample to Independence, he does not want to rest the main weight of his critique on such examples.⁹ But Errol Lord—in a piece entitled “From Independence to Conciliationism: an Obituary” (2014)—attempts to push Kelly’s style of counterexample harder. Instead of arguing that Independence-respecting trains of reasoning for dismissing disagreement will be *unavailable* in some asymmetrical-evidence cases, Lord, like Arsenault and Irving, claims that even if such lines are available, the Independence-violating trains of reasoning are legitimate. He writes, “If you can permissibly cite the fact that you have more evidence, you can permissibly cite that evidence. Since you can permissibly cite the fact that you have more evidence, you can cite the evidence. Thus, Independence is false.” (370).

However, as we saw above, we should be careful in how we treat intuitions about what one may reasonably *say* in failing to conciliate with dissenters’ opinions. Once one has dismissed the dissenter’s opinion, one will be in a position to cite one’s original evidence and reasoning, both to support one’s belief and to explain why the dissenter is mistaken. So we cannot refute Independence by cherry-picking cases where dismissing disagreement is legitimate on Independent grounds, and pointing out that it’s natural to support the retained belief by citing one’s original first-order reasoning. Again, we can test the legitimacy claim more carefully: by considering cases where Independence would *not* allow dismissing dissent.

There are many such cases. The Extra-Tricky Question case was one designed to parallel A&P’s example. But we can more closely parallel the Holocaust Denier case by considering a disagreement over a matter where we do not have entailing evidence. In order to make it a case where Independence-respecting dismissal will not be appropriate, we need only equalize the

⁸ As Alex Grossman pointed out to me, it’s not clear that the anti-question-begging motivation for Independence would really support the requirement that I bracket the evidential support E gives to H in this case. Let us suppose that Denier is completely ignorant of E. Then one might think that his disagreement does not cast significant doubt on my original assessment of E’s bearing on H.

There are some tricky issues here, but let us suppose that this is correct. In that case, one might argue that this sort of case should not be even a *prima facie* counterexample to Independence, if Independence is properly formulated. This line might suggest that a correct formulation of Independence should allow in certain cases that some of the reasoning behind my original belief—the reasoning purely from non-shared evidence— could after all be used in evaluating the reliability of my peer.

I’m not sure what to make of this line. One complication is this. It seems that evaluating the reliability of those who disagree is important at least in part because that bears on my own reliability. But the anti-question-begging motivation for Independence really does support not allowing me to evaluate the reliability of my own first-order reasoning about H in a way that depends on that very reasoning. So it might be argued that Independence, properly formulated, would after all prohibit my relying on my reasoning from E to H in assessing my own reliability after meeting Denier (or before, for that matter). If we take that line, then then a properly formulated Independence principle would after all prohibit me from relying on my reasoning from E to H in assessing my own reliability. In that case, Kelly’s case would be a counterexample to Independence —if there was no way of dismissing Denier that did not rely on my reasoning from E to H. So the fact that there does seem to be a way of dismissing Denier’s opinion in a way that doesn’t rely on that reasoning would be crucial to preserving Independence after all.

⁹ Kelly’s main argument will be examined below.

evidence possessed by the two parties. So consider an example of the sort often used to motivate conciliatory accounts of disagreement:¹⁰

Rain Denier: I believe, on the basis of solid, but complex, meteorological evidence E, that it will rain tomorrow (R). I discover that my colleague Wendy, who also has E as her meteorological evidence, believes $\sim R$. Wendy and I have long track-records in this sort of prediction. Neither of them is perfect, but Wendy's is just as good as mine. In cases where we've disagreed in the past, she's been right as often as I have.

In this case, it would not be rational for me to simply dismiss Wendy's disagreement—even if, as we may suppose, this happens to be one of the cases where I'm the one who assessed the meteorological evidence correctly.¹¹ But of course, there is an Independence-violating train of reasoning that would justify dismissing Wendy's disagreement, if it were legitimate: E really does support R, after all. The difference between Holocaust Denier and Rain Denier lies just in the higher-order consideration that in the former, I have more probative first-order evidence than Denier does, while in the latter Wendy and I share the same first-order evidence. Of course, in Rain Denier, I can cite the fact that we have the same evidence E. But as we can see, this does not translate into the permissibility of my dismissing Wendy's disagreement by directly citing E to support R. There is an important distinction between citing the fact that I have certain evidence, and directly citing that evidence in support of dismissing my friend's dissent. Erasing that distinction would leave us without any explanation for why my dismissing Wendy's disagreement would be illegitimate.

This point generalizes. Lord seeks to defend the legitimacy of Independence-violating reasoning in cases where Independence-respecting reasoning is also available. He complains that invoking Independence-respecting reasoning to defuse purported counterexamples like Holocaust Denier goes wrong by treating the purported counterexamples as 'extensional': that is, as cases where Independence is inconsistent with the correct verdict. But as we've seen, if the sort of reasoning he would defend were legitimate, it would be legitimate even in cases where Independence-respecting reasoning is *not* available. The literature on Conciliatory approaches to disagreement provides plenty of examples: Some involve disagreements with peers (e.g. disagreements about restaurant bills between equally-reliable mental calculators, or disagreements about rain between equally reliable forecasters). Perhaps more pointedly in the present context, there are also disagreements with epistemic superiors (e.g. when a logician tells me my proof is wrong, or when a superior meteorologist tells me it's not going to rain). And the literature on other higher-order reasons for self-doubt offers still more cases (e.g. medical diagnoses made in conditions of

¹⁰ I've changed the topic of disagreement to avoid contamination of our intuitive judgments by, e.g., background assumptions about the irrationality or immoral motivations behind actual cases of Holocaust denial.

¹¹ One could always simply reject this judgment. But a main motivation for invoking Independence principles is precisely to explain intuitively compelling conciliatory verdicts in this sort of case. Simply denying those intuitive verdicts would of course leave Independence unmotivated. But I take it that Lord (like Kelly and A&I) is after a more interesting way of arguing against Independence.

sleep-deprivation, conclusions about airplane fuel levels made when likely hypoxic). In many such cases—the ones that motivate invoking Independence principles in the first place—it seems illegitimate to dismiss worries about one’s own thinking by relying on the targeted reasoning. But why? If it’s not intuitively legitimate to use Independence-violating reasoning to dismiss doubts in these cases, and if the cases where dismissal is legitimate are exactly those cases where Independence-respecting reasons for dismissal are available, surely this suggests that it’s only the Independence-respecting reasoning that’s legitimate, even when both are available. To put the point another way: if Independence is false, it should be quite easy to find ‘extensional’ counterexamples to it.¹²

2. Worries about Motivation

Even if the apparent counterexamples to Independence principles do not in the end succeed, it is undeniable that the principles are awkward. So if a less awkward way could be found of explaining why dismissal is legitimate in some cases and not in others, Independence might well turn out to be unmotivated. Here, I’ll look at two arguments along these lines, and a third argument intended to show that the motivation for Independence principles only supports weak versions of the principle—versions which cannot support full-blown Conciliationism about disagreement.

a. Independence and explaining friends’ errors

Arsenault and Irving, having rejected Independence, do offer a way of distinguishing between cases where Independence-violating dismissal is legitimate, and ones where it is not.¹³ They contrast their Tricky Question case with a more standard case where I disagree with my peer about a non-tricky math question, and acknowledge that in that case, relying on my own reasoning to dismiss my friend’s dissent would illegitimately beg the question. So consider:

Ordinary Math Test: My friend and I are taking an ordinary math test, without tricky questions. We have long and equally-good track records on this type of test, and when we’ve disagreed, we’ve been right equally often. I think (correctly) that the answer to question 8 is ‘true’, but my friend thinks it’s ‘false’.

Why would it be illegitimate to use Independence-violating reasoning in this sort of case, but not in Tricky Question? A&I propose that the difference is that in the Tricky Question case, my use of my first-order reasoning is incorporated into an *explanation* of why my friend made a mistake.

It’s not clear to me that this criterion really will give the right answer in the general run of non-tricky-question disagreement cases. In the Ordinary Math Test example, my first-order reasoning would support the hypothesis that my friend made some calculational error or other, but

¹² Thus it seems to me that—to adapt a line from Mark Twain—Lord’s report of the death of Independence-based Conciliationism is something of an exaggeration (See White (1897)).

¹³ The account they offer is only designed to apply to disagreement cases—I won’t worry about this here.

would not point to a specific error. Perhaps the idea is that that's not specific enough an explanation. But we might consider a math test comprised of problems for which one particular kind of error is particularly common. Suppose the error is failing to carry the "1", but the problems are complicated enough (with the results of some calculations being subtracted from the results of others) that when two people's answers differ by 10 or 100, we don't have reason to think that the higher answer is more likely to be correct—just that someone forgot to carry a "1" somewhere. So suppose that my friend and I have long track-records of equally high reliability on this sort of problem, and that on the occasions that we have disagreed, we've been right equally often. If I reach a certain answer, and my friend announces that her answer differs from mine by 10, it seems clear that I cannot dismiss her disagreement. But of course here I would have a quite specific explanation of her error—that she forgot to carry a "1".

One might object that this is still a more general explanation than I had in the Tricky Question case, where I'd identified a specific place at which my friend made the error. But we can make our new case parallel in that way, too. Suppose I see some train of reasoning involving a particular carrying omission that would yield my friend's answer. But—to make the higher-order evidence parallel—suppose that she tells me she sees a train of reasoning involving a carrying omission that would produce my answer. (And we can suppose that we have a long history of being right equally often in *this* type of disagreement.) Again, it would be illegitimate for me to dismiss her via my original reasoning. But this time, my original reasoning would support a very specific explanation of her error.

In fact, the Extra-Tricky Question example reinforces this point. There, my original reasoning points to a particular error on my friend's part—double-counting the 2. In fact, it's the exact same explanation A&I invoke in their account of why the Independence-violating reasoning is legitimate in their original Tricky Question case. But again, once we take care to control the thought-experiment by balancing the higher-order evidence, Independence-violating dismissals of my friend's disagreement no longer seem legitimate. And this is unrelated to whether those dismissals involve a more or less specific explanation for my friend's error.¹⁴

b. Does the Requirement of Total Evidence make Independence Otiose?

Kelly, also concentrating on disagreement cases, agrees with the Conciliationist that some loss of confidence is often required, even when the agent in question began by reasoning correctly. He notes that in some such cases—e.g. when an expert disagrees with my belief that P—the Conciliationist is correct about one thing: it would be unreasonably dogmatic to dismiss the expert's disagreement by a train of reasoning such as: "Well P is true, and you believe \sim P, so despite your general reliability, you're wrong this time, and I needn't worry about your disagreement." Nevertheless, Kelly argues, we don't need to invoke anything like Independence to explain this

¹⁴ It's also worth noting that Independence-violating reasoning that involves a particular explanation of one's friend's error is equally available in cases where the friend is an epistemic superior—yet dismissing the friend's opinion in such cases is intuitively irrational.

phenomenon. He offers what he takes to be a better explanation: that in such cases, my *total evidence* no longer supports my belief that P; so it no longer supports my dismissing the expert by citing P.

[H]ere is the short story about why it will often be unreasonable and dogmatic for me to dismiss your contrary opinion in the envisaged way: after I add the fact that you believe as you do to my stock of evidence, it will no longer be reasonable for me to believe that p, given what is then my total evidence. And if it's no longer reasonable for me to believe that p, then I lack any rational basis for inferring that your sincere testimony is misleading evidence. (45)

As Kelly notes, the Conciliationist will say exactly the same thing. In cases similar to Ordinary Math Test, the Conciliationist will say that my total evidence—including the disagreement of a logic expert, or even the disagreement of a peer—no longer supports my belief that answer C is correct. But the Conciliationist will argue that we cannot explain this without appeal to a principle along the lines of Independence. And this is where Kelly dissents.

Someone who rejects such principles can still account for the badness of intuitively dogmatic reasoning simply by appealing directly to the normative requirement that one take into account one's total evidence (as opposed to some proper subset of one's total evidence). In short, the Principle of Total Evidence can do all of the work that needs doing. (46)

Of course, this raises the question of why it is—if we reject Independence principles—that the total evidence doesn't support P after I learn of my friend's disagreement. As Kelly realizes, the apparent need for something like Independence is sharpest in cases (like disagreements over math test answers) where my original belief was based on my seeing how the evidence *entails* the correctness of my answer. After all, what bearing could sociological facts about people's opinions have on the purely mathematical question of whether certain premises entail the correctness of a certain answer? Since my total evidence still entails the correctness of my original belief, it's hard to see why a peer's (or even an expert's) testimony should have any weight at all. Kelly writes:

[T]here is a genuine puzzle here, but it is a mistake to think that that puzzle motivates the adoption of Independence or Independence* [these are Kelly's formulations of disagreement-specific Independence principles]. After all, ... whenever one performs a non-trivial calculation, one should not be perfectly confident of one's answer even before another person comes on the scene (given one's awareness that one is fallible, etc.) But once it is granted that one should not be perfectly confident even before one's view is contradicted by a peer, there is no additional mystery or formal difficulty as to how acquiring that misleading testimonial evidence can push the credence that it is reasonable for one to have still lower. ... There is presumably some story to be told here. The crucial point is that there is no reason to think that the story in question entails Independence or Independence*, since

those principles explicitly concern how one should assess the beliefs of other people, and the phenomenon arises even before other people come on the scene. (46-47)

Kelly is certainly correct in pointing out that no *disagreement-specific* Independence principle can account for the general phenomenon in question. But disagreement is best seen as just one among many sources of doubt about the reliability of one's thinking. (Our question about total evidence comes up equally in other cases: what bearing does my tiredness, or facts about oxygen levels in my blood, have on the question of whether a certain conclusion is true, given premises that entail that conclusion?) So it's only natural to expect that whatever explains why disagreement evidence can require loss of confidence in propositions entailed by one's evidence will also explain why evidence that one is drugged, tired, biased—or even one's awareness that one is fallible—can do the same thing. And if something along the lines of Independence can provide the requisite general explanation, then the right way of applying the Principle of Total Evidence will actually *require* using an Independence principle.

Some support for this approach can be seen by contrasting two Entailment cases:

Entailment vs First-Order Inductive Evidence: I'm doing a multiple-choice math test, in which I know that the right answers were assigned randomly to letters A through E. So without looking at what problem 8 says, I have statistical evidence which makes rational low (.2) credence in answer C, which is "15." Then I look at the problem: " $2.5 \times (9 - 3) = x$ ". I rationally become extremely confident that C, "15," is correct.

Entailment vs Higher-Order Inductive Evidence: I'm doing a fill-in-the-blank math test. Problem 8 is " $2.5 \times (9 - 3) = \underline{\quad}$ ". I become extremely confident that the correct answer is 15. Then I learn that before doing the problem, I'd been given a drug which causes people to get wrong answers to this sort of problem 4/5 of the time, while seeming to themselves to be doing math clear-headedly. (I've been given the drug before, and have arrived at correct answers only 20% of the time.) My rational credence that "15" is the right answer drops dramatically, to around .2.

In both cases, I have some problem-given evidence, and I also have some inductive evidence. In the first case, the inductive evidence is that answer C is correct in only 20% of the problems; other things equal, it would support .2 credence that the answer is "15". In the second case, the inductive evidence is that an agent who'd been drugged in a way that results in getting correct answers 20% of the time got the answer "15"; other things equal, this also would support .2 credence in "15". But in each case, my total evidence also includes the problem-given evidence that entails that the correct answer is "15".

In the case involving ordinary inductive evidence, the entailing evidence acts just as we might expect it to: it swamps the inductive evidence, rendering it virtually irrelevant. But in the case involving significant higher-order reason for self-doubt, the entailing evidence cannot stand up to the inductive evidence that I'm likely to be reasoning badly. What explains the difference? Of

course, there are many complexities that arise in cases where different bits of evidence push in different directions, even when only first-order evidence is involved. But it seems to me that a very natural explanation for the dramatic difference between these two cases in particular lies in some aspect of rationality that, in some way, at least partially disables the entailing evidence in the second case, allowing the inductive evidence to gain some purchase. Independence principles—which disallow full reliance on otherwise-compelling reasoning when an agent has higher-order reason for self-doubt—seem a very natural way to explain this puzzling phenomenon.

So it seems to me that Kelly's observation does not in the end reveal any lack of motivation for Independence principles. But it does underscore an important point: that the motivation for such principles is not confined to disagreement cases. The correct understanding of rational responses to one's total evidence in disagreement cases should flow out of our more general understanding of how the Principle of Total Evidence applies in cases where some of one's evidence is higher-order. And this general understanding should include an explanation of how entailing evidence, while it typically swamps other evidence, is easily undermined by certain sorts of higher-order evidence. If an Independence principle can provide the requisite explanation, then we have reason to see Independence not as being a competitor to the Principle of Total Evidence, but instead, as being exactly what allows us to apply the Principle of Total Evidence correctly, in judging our own epistemic performance.

One final point is worth making about the relationship between Independence principles and the Principle of Total Evidence. One might worry that Independence principles are somehow incompatible with the basic idea behind the Requirement of Total Evidence: after all, don't they require an agent's response to her total evidence be independent of the support provided by (at least part of) her evidence? And doesn't render a portion of her evidence irrelevant?

To answer this question, it's important to note that Independence principles apply directly to just one component of an agent's rational response to her evidence: her assessment of her own reliability. And that assessment may combine with various other factors—factors which are also shaped by the agent's total evidence—to determine the credence that's rational for the agent to end up with. (We saw one aspect of this in thinking about cases where an agent brackets some reasoning in judging her friend unreliable, and then, having dismissed the friend's disagreement, rationally relies on that same reasoning in retaining her belief, and in explaining where her friend went wrong.) So there is no reason to think that accounts of rational belief that incorporate Independence principles must have the consequence that some portion of an agent's total evidence is epistemically irrelevant. True, some Independence-based accounts of higher-order evidence would have this consequence. But others, just as clearly, do not. For example, on some Independence-based accounts, the fact that my problem-given evidence really entailed that the correct answer in the second Entailment case above was 15 would be crucial in determining what credence it was rational for me to have in that case.¹⁵ So Independence principles need not conflict with the Requirement of Total Evidence in any way—in fact, they may help constitute it.

¹⁵ For examples of both sorts of accounts, and detailed discussion of this issue, see Sliwa and Horowitz (2015), Schoenfield (2015) and Christensen (2016).

c. Can only a Weak Form of Independence be Motivated?

Han van Wietmarschen (2013) does not argue for rejecting all versions of Independence principles. Instead, he argues that attention to the motivation behind imposing such principles will show that the strong versions needed to underwrite robustly conciliatory views of disagreement are false. The weaker principles (which he thinks can be properly motivated) allow that agents who hold onto their original beliefs in certain disagreement situations remain justified to believe their original beliefs, despite the disagreement.

Van Wietmarschen begins by pointing out that the loose formulations of Independence found in the literature admit of two sorts of interpretation when they require, for example, that when I'm in a disagreement situation, I evaluate the epistemic credentials of the person who disagrees with me "in a way that doesn't rely on the reasoning behind my initial belief".¹⁶ Focusing on cases where an agent initially forms a belief that P on the basis of first-order evidence E, an "Unrestricted" Independence principle would require the agent to bracket all reasoning from E to P in assessing her own and her friend's reliability. A "Restricted" principle would only require the agent to bracket her own original actual *process* of reasoning.

His argument focuses on disagreement situations where two epistemic peers do some fairly quick reasoning, then discover that they disagree—situations like Ordinary Math Test, where I believe that the answer to question 8 is 'true', then discover that my equally-reliable friend believes it's 'false'. Van Wietmarschen argues that in this sort of situation, Unrestricted versions of Independence are unmotivated, and we should only require the agent's assessment to be independent of the actual episode of reasoning she used in coming to her initial view. This is because my friend's disagreement does not impugn my general ability to reason to correct answers on math tests—it only provides evidence that I made an error in my particular original reasoning process on this occasion. So while it would beg the question to rely on that particular process—hence motivating the imposition of a Restricted form of Independence—the same kind of motivation is not available to support imposing Unrestricted Independence.

Given this claim, van Wietmarschen argues that agents in disagreements such as Ordinary Math Test do not lose justification to believe their initial views (assuming that those views were justified to begin with), despite the disagreement. That's because such an agent's first-order evidence supports the claim that she was right this time, and this support is not disabled by Restricted Independence, which only applies to a specific past episode of reasoning. So if we understand Independence correctly, an Independent assessment would not, after all, cast doubt on the claim that the agent's original belief was correct, and the friend was mistaken this time.¹⁷

¹⁶ The ambiguous formulation is from Christensen (2011).

¹⁷ Van Wietmarschen also offers a more detailed positive argument that on plausible theories of evidential support, Restricted Independence allows agents in cases like Ordinary Math Test to be justified to believe the proposition they originally believed (supposing, still, that those beliefs were justified before the disagreement). He illustrates his argument with the following theory of evidential support, which he credits to Richard Foley:

(ES): A body of evidence E supports the belief that p for subject S if and only if S would believe p on ideal reflection on whether p, given access to E.

Van Wietmarschen goes on to argue that in cases like Ordinary Math Test, even if the steadfast agent is still justified *to believe P*, she is no longer justified *in believing P*—her belief in P is no longer *well-grounded*. That’s because the higher-order evidence does target the episode of reasoning that the agent used in coming to believe P. Restricted Independence applies, so to be properly grounded, the agent’s belief would have to be based on, or grounded in, reasoning independent of that targeted reasoning episode. We might put this in terms of propositional and doxastic rationality by saying that, on van Wietmarschen’s account, P is still *propositionally rational* for the agent to believe, but the agent no longer has a *doxastically rational* belief in P. And van Wietmarschen rightly points out that much of the disagreement literature does not distinguish between the two notions. On this view, our intuitive rejection of the rationality of steadfast responses to cases like Ordinary Math Test should be seen merely as responding to a lack of well-groundedness, or doxastic rationality.¹⁸

One important thing to notice about this position is that the proposed narrowing of conciliatory views should not provide much comfort to those disturbed by the views’ skeptical implications. True, it would have the consequence that those who had reasoned correctly prior to learning of disagreement, and then stood by their beliefs, still believed the propositions that were supported by their evidence. But their beliefs would not be held rationally, any more than would be the beliefs of those who form beliefs through wishful thinking, and just happen to arrive at the beliefs their evidence supports. In maintaining our controversial beliefs, we would still be believing irrationally. Indeed, van Wietmarschen holds that agents in such situations are required to suspend judgment on the disputed topic.

Nevertheless, the argument does pose an interesting question for defenders of Independence. Is it true that the motivation for Independence principles only supports requiring independence from a particular episode of reasoning?

To examine this more general question, we might look at examples beyond the sort of disagreements van Wietmarschen concentrates on. We can start with a case that doesn’t involve disagreement at all. Consider the small-plane pilot who has good evidence from her altimeter that she’s likely to be hypoxic. Even before she reads the dials and gauges that provide the data relevant to calculating whether she has enough fuel to reach the more-distant airport—that is, before she forms any judgment about whether she has enough fuel—she has reason to distrust whatever calculations she might perform from this data. In this case, what’s targeted by her higher-order

Van Wietmarschen’s idea is that ideal reflection is deep and thorough, so in the disagreement situation, an agent who had reflected ideally would realize that while she had reflected deeply and thoroughly, her friend who disagreed about the logic answer had only thought about it briefly during the test.

However, as Chad Marxen points out, this argument is problematic. If we imagine such an agent reasoning that way, we’re imagining her to have evidence about her reasoning process that is absent in the original case. So even if such an agent would arrive at a rational belief that P, she would *not* do so on the basis of the original evidence E. And the argument is supposed to show that P is rational to believe *on E*, not on a batch of total evidence that differs with respect to crucially relevant claims about the two agents’ reasoning. I will not go more deeply into this problem here; see Marxen (ms.) for details.

¹⁸ I will not discuss van Wietmarschen’s treatment of well-groundedness, but will note that a somewhat similar position is defended in Smithies (2015). For some reasons to worry about holding that disagreement destroys well-groundedness without affecting propositional justification, see Marxen (ms).

evidence is not some specific episode of reasoning that took place in the past. The pilot's higher-order evidence targets any reasoning the pilot might do from the dials and gauges to the conclusion that she has enough fuel.

Given that the pilot's higher-order evidence targets any such reasoning she might do, a properly formulated Independence principle will have to cover any such reasoning. So the motivations for imposing an Independence requirement in this sort of case will support something stronger than Restricted Independence—in fact, something more along the lines of Unrestricted Independence would seem to be exactly what's motivated. And it seems clear that the same will apply in many other cases of higher-order evidence—for example, ones involving bias, drugs, or fatigue. It should even apply in many cases of disagreement—especially cases of persistent disagreement on controversial issues that give the issue much of its interest. After all, in these cases one often knows about the disagreement before one studies the issue.¹⁹ In many such cases, the reasoning targeted by the higher-order evidence would include any first-order reasoning the agent might do on the relevant topic. If that's right, then in many cases, something along the lines of Unrestricted Independence seems to be appropriate.

With this in mind, we might now ask: What does this show about the suggestion that higher-order evidence in these cases is relevant only to the well-foundedness (or doxastic justification/rationality) of an agent's beliefs, and not relevant to what the agent has justification to believe (that is, to propositional justification/rationality)? Of course, the answer to this sort of question will depend on how we understand the relevant notions, and different writers have offered interestingly different accounts. But without taking a stand on the 'right' way of understanding them, it's worth noting a sharp difference between standard cases where agents form beliefs that are propositionally, but not doxastically, rational, and cases where an agent ignores higher-order evidence and ends up believing the proposition that would be supported by her first-order evidence alone.

Suppose that Jocko has strong evidence supporting the belief that Xena is guilty. But instead of thinking through the evidence, Jocko just jumps to the conclusion that she's guilty, because he dislikes her and wants her to be guilty. In such cases, the fact that Jocko's evidence makes it rational to believe that Xena is guilty—that it makes "Xena is guilty" propositionally rational for Jocko to believe—seems closely tied to the fact that, were Jocko to form a belief about this matter in the most rational way, given his evidence, he'd believe Xena to be guilty.

Now there may be difficulties with general principles that *define* propositional justification in this way. For example, in some strange cases, the agent's forming a belief about the relevant matter would itself provide the agent with evidence against the believed proposition.²⁰ But putting this sort of complication aside, it's quite plausible that something like this relationship is behind our idea of propositional rationality: in general, the propositions that are rational for an agent to believe, given

¹⁹ I should note that van Wietmarschen himself thinks that deep, persistent disagreements can undermine propositional justification.

²⁰ I'm not sure that these examples are as worrisome as they may seem. See Turri (2010) for useful discussion.

certain evidence, are those that would be believed as part of an ideally rational doxastic response to that evidence.

But notice that nothing like this is true in cases where agents disregard higher-order evidence. As we've just seen, higher-order evidence does not always target a particular past episode of reasoning; in many cases, it targets whatever first-order reasoning the agent might do on the relevant matter. Let us consider an agent in such a situation—for example, our pilot. Suppose that despite the strong evidence that she's hypoxic, she decides not to worry about that. Instead, she just consults the relevant dials and gauges, and forms the belief on that basis that she has enough fuel. Since the reasoning on which her belief is based is targeted by her higher-order evidence, her belief will not, on the suggested view, be doxastically rational. But if that's right, then it seems that she *cannot* reason her way to a doxastically rational belief that she has enough fuel, even if she manages to do her first-order reasoning impeccably. And it's not just that she lacks some epistemic skill: no one with her total evidence—which includes the higher-order evidence of hypoxia—could rationally believe that they had enough fuel. And it's worth noting that the barrier to our pilot's forming a doxastically rational belief that she has enough fuel has nothing to do with her forming such a belief providing evidence against the relevant proposition. It's her higher-order evidence, which is part of her original total evidence, which precludes doxastically rational belief that she has enough fuel.

To my mind, that constitutes a good reason to treat cases of ignoring higher-order evidence quite differently from standard cases in which an agent believes a proposition that is rational for her to believe, but where her belief is not doxastically justified/rational or well-founded. Indeed, it seems quite unintuitive to me to say that while our pilot's *evidence itself* precludes her rationally believing that she has enough fuel, nevertheless that evidence makes the proposition that she has enough fuel rational for her to believe! Of course, one could define a notion of 'evidential support' that worked this way—after all, one can simply define evidential support as not taking into account higher-order considerations. But a notion of "rational to believe" which divorces it so completely from what could be believed rationally seems to me to have left the topic of rational belief too far behind.²¹

The motivations for Independence, then, do not really limit us to Restricted-Independence-style principles which would only prohibit dependence on particular past episodes of reasoning. They extend to prohibiting dependence on certain reasoning that the agent might do—and in some cases extend to prohibiting dependence on any train of reasoning from the agent's first-order evidence to whatever belief that first-order evidence, considered in isolation, would support. In such cases, the agent is in an epistemic position where, *precisely because of her evidence*, she cannot rationally believe the relevant proposition. It seems quite unnatural to say, in such situations, that the agent's evidence makes the relevant proposition rational for her to believe.²²

²¹ I should note that van Wietmarschen would identify justification to believe with evidential support. Given that identification, I would be inclined to resist the claim that higher-order evidence is irrelevant to evidential support.

²² It's also worth noting that the effect of the higher-order evidence is not just to make it *irrational* for agents such as our pilot to have the beliefs that would be supported by their first-order evidence alone. Our intuitions in such cases also support *positive* rational evaluations of alternative beliefs (or degrees of belief). So, for instance, in disagreement situations such as Ordinary Math Test, when the agent learns of the disagreement of her peer, a credence of .5 seems more rational for the agent to have than the agent's pre-disagreement credence would be. In a similar disagreement with

3. Conclusion...the way forward

If all this is on the right track, some sort of Independence principle is a legitimate, and perhaps necessary, element in understanding rational responses to higher-order evidence. Saying that does not, of course, tell us what such a principle will look like in detail. But in addressing some of the concerns that have been raised about the very possibility of legitimate Independence principles, I think we have gotten some clues as to what a good formulation might look like.

First, it must allow reliability assessments to be informed by facts about certain trains of reasoning. Paradigm examples would include the fact that a train of reasoning involved apparent discovery and rejection of a decoy answer, or the fact that it was based on evidence likely to prove strongly probative on the relevant issue. The challenge is to allow this while precluding reliability-assessments that depend on certain other facts about the reasoning, such as the fact that it involves *actual* discovery and rejection of a decoy answer, or the fact that it was based on evidence which *does strongly support* a given belief about the relevant issue. As a first step, we should require that the reliability-assessment be independent of the targeted reasoning in the sense of not relying on the cogency of that reasoning.

We've also seen that the targeted reasoning should be understood to encompass more than a particular past reasoning episode. In many cases, the bearing of higher-order evidence is on possible reasoning the agent might do, not just on reasoning she has done. In light of this, I suspect that the right way of formulating Independence will involve reference to epistemic support relations, rather than episodes of reasoning. As a first pass, we might try something like this: the relevant reliability-assessment is one that's supported by the agent's evidence, when one excludes whatever support relations between the agent's first-order evidence and the proposition in question are such that the higher-order evidence suggests that the agent is likely to mis-assess them.

Unfortunately, this first pass is not nearly good enough—a lot of questions remain. For example: Does it really make sense to divide first-order from higher-order evidence? I think it does not, because there will be many cases where the same piece of evidence can play both roles. The beliefs of other people, for example, are often evidence at both levels. This suggests that the right way to describe the evidential support relations that need to be bracketed is not by reference to what pieces of evidence they involve. Instead, we probably need to focus on differentiating first-order from higher-order evidential support relations. Roughly, evidence bears on an agent's doxastic attitude toward P in a higher-order way insofar as it bears on that attitude via bearing on how reliable an agent should think she is in assessing how evidence bears on her attitude toward P in a direct way. That's the intuitive idea—and it's not hard to illustrate it with examples. But it's not really a clear formulation.

an epistemic superior (e.g. an expert mathematician), an even lower credence would be more rational. So the role of higher-order evidence is not just that of interfering with the rationality of the doxastic response that would be appropriate absent the higher-order evidence. Higher-order evidence contributes to the rationality of alternative doxastic responses. This does not seem explicable on an account where the higher-order evidence interferes with well-foundedness without changing what the agent's total evidence supports.

A related question is whether one can neatly circumscribe the evidence whose first-order support for the relevant proposition needs to be bracketed by the Independent reliability-assessment. I suspect that again, the answer is no. For example, prospects for circumscribing the relevant evidence seem pretty dim when one thinks about disagreements over big issues, such as we find in economics or philosophy. And the more holistic we take evidential support to be, the dimmer the prospects. My current inclination is to think that the right way of handling this is to include all first-order support of the relevant proposition. But I'm not confident that this won't cause trouble.

Another question that arises is this: can we cleanly divide cases where the agent has higher-order reason for self-doubt from ones where the agent has no apparent evidence at all relevant to her reliability? Again, the answer seems to be no: surely there is a continuum here. And further along the continuum are cases where an agent has positive evidence of her own reliability. So if there isn't really a discrete class of cases where higher-order evidence is relevant, can we find a smooth, unified way of handling all cases? I hope so. At this point, the approach that seems most promising to me is to posit a degree of self-trust as being warranted independently of any evidence. If that worked out, it could provide for rational reliability-assessments even in the absence of specific evidence bearing on reliability.

Yet another problem concerns cases where higher-order evidence is indiscriminate—in particular, when it targets the agent's ability to reason about higher-order evidence. Some such cases seem so close to paradoxical that we might skip over them with relatively clear conscience—but others do not.

So those are some of what seem to me to be significant obstacles that we face in formulating an Independence principle. I've sketched some of the directions that seem to me most promising for meeting those obstacles. But I'd be lying if I said I felt clear or confident about any of them—there's clearly a lot of work remaining to be done. I personally suspect that in the end, the biggest threat to Independence principles does not lie in the sort of counterexample-based or motivational challenges we looked at earlier. I suspect that it lies in the difficulty of coming up with a precise version of Independence that can address these sorts of questions in a satisfactory way.²³

At this point, I'm convinced that finding an acceptable Independence principle is necessary, if we are to describe the most rational way of acting as judge in one's own epistemic case. Given that it's necessary, I very much hope that it's not also impossible.

Acknowledgements

Arguments forming the basis of this paper were previously made in talks at the University of Texas at Austin, the University of Cologne, and New York University, and were discussed in my graduate and undergraduate seminars at Brown; thanks to the participants at all of these occasions. Thanks also to the audience at the APA Eastern Division Meeting in Savannah. And thanks especially to Nathan Ballantyne, Anna-Maria Eder, Alexander Grossman, Chad Marxen, Josh Schechter, Declan Smithies, and Jonathan Vogel for helpful discussion of these issues and/or comments on earlier

²³ I try to meet some of these challenges in Christensen (forthcoming).

drafts of this material. Finally, thanks to the Marc Sanders foundation for sponsoring the Sanders Lectures, and for all the different ways they support contemporary Philosophy.

References

- Arsenault, M. and Z. Irving (2013), "Aha! Trick Questions, Independence, and the Epistemology of Disagreement," *Thought* 1: 185-194.
- Christensen, D. (2010), "Higher-Order Evidence," *Philosophy and Phenomenological Research* 81: 185-215.
- . (2011), "Disagreement, Question-Begging and Epistemic Self-Criticism", *Philosophers' Imprint* 11, no. 6.
- . (2016), "Disagreement, Drugs, etc.: from Accuracy to Akrasia," *Episteme*: 397-422.
- . (forthcoming), "Formulating Independence," to appear in M. Skipper and A. Steglich-Petersen, eds., *Higher-Order Evidence: New Essays* (Oxford University Press).
- Kelly, T. (2005) "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: 167-96.
- . (2013), "Disagreement and the Burdens of Judgment," in D. Christensen and J. Lackey, eds., *The Epistemology of Disagreement: New Essays* (Oxford University Press).
- Lackey, J. (2010), "A Justificationist View of Disagreement's Epistemic Significance," in A. Haddock, A. Millar and D. Pritchard, eds., *Social Epistemology* (Oxford University Press).
- Lasonen-Aarnio, M. (2014), "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research* 88.2: 314-345.
- Lord, E. (2014), "From Independence to Conciliationism: An Obituary," *Australasian Journal of Philosophy* 92: 365-77.
- Matheson, J. (2015), *The Epistemic Significance of Disagreement* (Palgrave Macmillan).
- Marxen, C. (ms.), "A Defense of Propositional Conciliationism".
- Schoenfield, M. (2015), "A Dilemma for Calibrationism," *Philosophy and Phenomenological Research* XCI No. 2: 425-455.

- Smithies, D. (2015), "Ideal Rationality and Logical Omniscience," *Synthese* 192 (9): 2769-2793.
- Sosa, E. (2010), "The Epistemology of Disagreement," in A. Haddock, A. Millar & D. Pritchard, eds., *Social Epistemology* (Oxford University Press).
- Sliwa, P. and S. Horowitz (2015), "Respecting *All* the Evidence," *Philosophical Studies* 172 (11): 2835-2858.
- Titelbaum, M. (2015), "Rationality's Fixed Point (Or: In Defense of Right Reason)", *Oxford Studies in Epistemology* 5: 253–294.
- Turri, J. (2010), "On the Relationship between Propositional and Doxastic Justification," *Philosophy and Phenomenological Research*, 80(2), 312–326.
- van Wietmarschen, H. (2013). "Peer Disagreement, Evidence, and Well-Foundedness. " *Philosophical Review* 122(3): 395–425.
- White, F. M. (1897), "Mark Twain Amused / Humorist Says He Even Heard on Good Authority That He Was Dead / Cousin, Not He, Sick / New Book Just Finished, and It Will be Called *Another Innocent Abroad*," *New York Journal*, June 2: p. 1.