Bottom Up justification, asymmetric epistemic push, and the fragility of higher order justification

Kappel, Klemens

Published in:
Episteme

DOI:
[10.1017/epi.2017.19](https://doi.org/10.1017/epi.2017.19)

Publication date:
2019

Document version
Peer reviewed version

Citation for published version (APA):
Kappel, K. (2019). Bottom Up justification, asymmetric epistemic push, and the fragility of higher order justification. *Episteme*, 16(2), 119-138. <https://doi.org/10.1017/epi.2017.19>

Bottom Up justification, asymmetric epistemic push, and the fragility of higher order justification[#]

Forthcoming in Episteme, 2017

Abstract

When a first order belief accurately reflects the evidence, how should this affect the epistemic justification of a higher order belief that this is the case? In an influential paper, Kelly argues that first order evidential accuracy tends to generate more justified higher order beliefs (Kelly 2010). Call this *Bottom Up*. I argue that neither general views about what justifies our higher order beliefs nor the specific arguments that Kelly offers support Bottom Up. Second, I suggest that while we can reject Bottom Up, we can still accept that justified higher order beliefs significantly affect the justification of first order beliefs. Third, I argue that the epistemic justification of higher order belief is *fragile* in the sense that it tends to dissipate when a subject is confronted with certain defeaters, including notably the sort of defeaters arising from disagreement, precisely when higher order justification depends on first order success in the ways that one may think support Bottom Up.

1. Introduction

Suppose that my object level belief accurately reflects the strength of the evidence available to me: my evidence e warrants credence n in a proposition p , and upon

[#] Earlier versions of this paper were presented at workshops in Leuven, Cologne and Copenhagen. Thanks to audiences for helpful comments, in particular Chris Kelp, Thomas Grundmann, Jon Matheson, David Christensen, Mikkel Gerken, Bjørn Hallsson, Emil Moeller, Frederik Andersen, Giacomo Melis, Fernando Broncano-Berrocal, and Josefine Pallavicini. Thanks to an anonymous referees for Episteme for no less than two sets of detailed and helpful comments to earlier drafts.

appreciating the evidence accurately I believe p with just that credence. Am I then justified in my higher order belief that my object level belief correctly reflects the evidence? Conversely, suppose that I got it wrong at the object level; I accord some proposition higher or lower credence than the evidence actually warrants. Does this render my higher order belief (that I did get it right) less epistemically justified? In an influential paper in the epistemology of disagreement, Kelly argues that first order evidential accuracy tends to bring about higher order justification (Kelly 2010):

'... when one correctly responds to a body of evidence, one is typically better justified in thinking that one has responded correctly than one is when one responds incorrectly' (156).

And similarly:

'... when E is genuinely good evidence for H , this very fact will contribute to the justification for believing the epistemic proposition that E is good evidence for H that is available for those with the relevant competence.' (159).

For convenience, I will refer to this view as *Bottom Up*, though Kelly does not use that term.¹ Roughly, Bottom Up asserts that appreciating first order evidence

¹ Titelbaum uses 'bottom-up' for the view that when some first order belief is rationally impermissible, then 'no amount of testimony, training, or putative evidence about what's rational can change what is rationally permitted or what the agent is rationally permitted to believe about it' (Titelbaum 2013, 279). In the same paper he asks 'How is the justificatory map arranged such that one is never all-things-considered justified in both an attitude A and the belief that A is rationally forbidden in one's current situation?'. The most obvious answer, Titelbaum suggests, is that 'every agent possesses a priori, propositional

accurately tends to make one epistemically justified in believing the epistemic proposition expressing what one's evidence is evidence for, and in believing that one has indeed responded correctly to one's first order evidence. Bottom Up also asserts that misapprehension of evidence does not contribute to higher order justification in similar ways. So, accuracy at the first order level makes a distinct difference to higher order justification.

The complement to Bottom Up is the more familiar idea that being justified in believing that there is something epistemically amiss with one's first order belief should make one less justified in that belief. I will call the view that there is this connection running from higher order belief to first order belief *Top Down*, and will say a bit more about it below.

Bottom Up is quite important for a particular argument in the epistemology of disagreement. Suppose that you and I know one another to be approximate epistemic peers looking at the same body of evidence. I now realize that we disagree. It strongly seems to me that p , whereas you are confident that p is false. Suppose that, in fact, I got it right: my belief correctly reflects the evidence, and you are wrong in your assessment of the evidence. As Kelly points out, it is tempting to think that though there is an asymmetry at the object level, you and I are even at the higher order level. Assuming that we both believe that we appreciate the first order evidence correctly, these higher order beliefs are equally justified. If this were true, then, as Kelly notes, 'it would be unreasonable to favor

justification for true beliefs about the requirements of rationality in her current situation. An agent can reflect on her situation and come to recognize facts about what that situation rationally requires. Not only can this reflection justify her in believing those facts; the resulting justification is also empirically indefeasible." (Titelbaum 2013, 276). This view is related to what I call Bottom Up. Titelbaum defends this as a part of his general view on akratic conflicts, and does not offer details about how a priori propositional justification may come about by reflection on first order evidence, or why it would be empirically indefeasible. A discussion of Titelbaum's interesting views is beyond the scope of the present paper.

one's own belief over the belief of one's peer' (155), despite the asymmetry at the object level. So, the argument goes, since we are equally justified in our higher order beliefs, you and I should accord equal weight to the first order credences involved in our disagreement, which in turn is a strong consideration in favor of the Equal Weight View. Kelly claims that this argument for the Equal Weight View fails, because 'Asymmetry at the lower level tends to create an asymmetry at the higher level, an asymmetry that otherwise would not have existed' (160). So, according to Kelly, the above argument for the Equal Weight View is blocked if Bottom Up is true, as this view implies that evidential asymmetries at the object level generate asymmetries at the higher level.

Actually, Bottom Up is important not only for the epistemology of disagreement, but for a wider question concerning how first order evidence and higher order evidence should be integrated. Consider a case of *epistemic akrasia*, that is, a case where a subject *S* rationally believes *p* on some evidence *e*, but also rationally believes that there is no good evidence for believing *p*. Until fairly recently, most epistemologists thought that *S* ought to integrate her beliefs, by adjusting her credence in her first order belief, or in her higher order belief, or both to avoid such conflicts. Borrowing a label from (Christensen 2007), I call this view *Integration*. As I understand Integration, all it says is that object level and higher-level belief must somehow be aligned. Integration is open to both Bottom Up and to Top Down, but could also accept one without the other, or could maybe even deny both and suggest a different way of aligning the levels. It should be noted that Integration is far from uncontroversial. Indeed, in recent literature a number of contributors have questioned Integration or specific versions of it, see for example (Lasonen-Aarnio 2014; Weatherson Ms). Others have defended certain versions of it, see in particular (Sliwa and Horowitz 2015; Horowitz 2014).²

My discussion concerns the details of Integration. According to Top Down, justified higher order beliefs affect the justification or rational credence of

² For further views on akratic conflicts, see see (Titelbaum 2013; Christensen 2013; Worsnip 2015)

first order belief. But according to Bottom Up, first order success may directly affect the level of justification of certain higher order belief. As will emerge below, one can accept Bottom Up, while rejecting Top Down, or conversely, just as one can, of course, accept both Top Down and Bottom Up, and thus hold that the direction of influence goes both ways. Finally, one can reject both views. How we understand the interactions between first order belief and evidence and higher order belief and evidence is decisive for how we should understand Integration.

In this paper I will argue for three main claims. First, I discuss a number of different arguments for Bottom Up - some derived from Kelly's work, others not - and argue that it is on the whole difficult to find compelling arguments for significant versions of Bottom Up (section 3). Second, I suggest that while we can reject Bottom Up, we can still accept the Top Down view that justified higher order beliefs significantly affect the justification of first order beliefs, and I elaborate on this peculiar asymmetry in direction of influence (section 4). Third, I argue that the epistemic justification of higher order belief is *fragile* in the sense that it tends to dissipate when a subject is confronted with certain defeaters, including the sort of defeaters arising from disagreement, exactly when higher order justification depends on first order success in the ways that one may think support Bottom Up. By contrast, when the justification of higher order belief depends on other sources, it is not as fragile, though this varies as a function of the source of justification (section 5). In section 6, I provide some concluding remarks.

2. Higher Order Beliefs and Bottom Up

We first need some preliminary remarks about higher order beliefs and Bottom Up. Here I consider *epistemic* higher order beliefs, that is, beliefs about epistemically relevant features of object level beliefs, though I will continue to simply refer to these as higher order beliefs. We might think of higher order beliefs in a variety of ways. One significant kind of higher order belief involves epistemic propositions. An *epistemic proposition* is a proposition that some body of evidence e warrants a particular doxastic attitude, for example full belief in a proposition or credence n in a certain proposition. For the most part I will for convenience use ' $\langle e$ warrants p to degree $n \rangle$ ' to refer to epistemic propositions, but none of the

arguments I discuss depend on the fine-grainedness or the nature of epistemic propositions that this suggests. One could instead run the arguments in terms of epistemic propositions of the form $\langle e \text{ warrants full belief in } p \rangle$.

Suppose that S believes the epistemic proposition $\langle e \text{ warrants } p \text{ to degree } n \rangle$, and that this proposition is true. Yet, for S to evaluate her own object level belief in p , it is clearly not enough that S believes this proposition. S must also believe that her first order belief in p is held with some credence or other so that she can assess whether her object level belief correctly reflects the epistemic proposition. Obviously, higher order beliefs need not be this complicated, and typically they are not. S may have a higher order belief that a first order belief has a particular desirable epistemic property, for example, that it is epistemically justified, or known. Or S may simply believe of a first order belief that it correctly reflects the evidence available to S, whatever level of credence that requires.

Return now to Bottom Up. The rough idea is that assessing the first order evidence correctly contributes in a distinct way to the justification of higher order beliefs, whereas misinterpreting evidence at the object level does not. Bottom Up does not, of course, assert that object level beliefs are always accompanied by higher order beliefs. Bottom Up only claims that when object level beliefs accurately reflect first order evidence, a subject is thereby in a position to form a highly justified positive higher order belief, whereas this is not the case when a subject misjudges evidence.

It is important to note that Bottom Up does *not* assert that there is a direct link between having *strong evidence* for object level belief and a high level of justification of higher order belief. Rather, what influences justification of higher order belief, is the *proper appreciation* of evidence at the first order level, be the evidence strong or not. Kelly concurs to this in the first of his remarks quoted above, though not in the second, where he seems to assert that strong first order evidence *itself* contributes to higher order justification. Clearly, however, we need the stricter second reading of Bottom Up. Suppose that Albert and Bertrand face the same evidence e pertaining to a proposition p , and Albert correctly grasps the evidential import of e , whereas Bertrand does not. In this case, the proponent of Bottom Up would want to say that Albert's higher order beliefs are more justified

than Bertrand's. The difference between them is due to differences in the accuracy at the first order level, not due to difference in strength of their evidence, as this is the same. It might often be true, however, that when some body of first order evidence is genuinely good evidence for a proposition, then it is easier to appreciate the force of the evidence correctly, and this would tend to yield a link between having good evidence at the first order and being more justified in one's higher order beliefs. But still, the decisive association is between first order accuracy and higher order justification, not between the strength of first order evidence as such and higher order justification.

3. Why we should not accept Bottom Up

In this section I argue that we should not accept Bottom Up. First, I consider the some obvious sources of justification for higher order beliefs and argue that they do not support Bottom Up. Next, I turn to Kelly's specific arguments for Bottom Up, and argue that they fail.

The most obvious source of justification for higher order beliefs is what we might call *track record evidence*. One form of track record evidence is generated by our justified expectations about our cognitive system. We may reasonably assume that our cognitive system works well under normal circumstances, and that under those conditions it generally leads us to accept beliefs that arise as proper responses to the evidence.³ If we can reasonably assume that our circumstances are indeed normal, then we have some evidence for thinking that our first order beliefs properly reflect the evidence.

Track record evidence hardly supports Bottom Up, however. The problem is that the strength of track record evidence is independent of evidential accuracy or other epistemic qualities of the particular object level beliefs that they

³ One can surely question the truth of the crucial premise as there is plenty of evidence indicating that we are subject to many forms of biases that incline us to misinterpret or neglect evidence, and that this is how our cognitive system works under normal circumstances. I will set this complication aside.

concern. Consider two cases that only differ by the quality of a subject's response to first order evidence. In one case, Albert responds properly to first order evidence. In an otherwise identical case, Bertrand mistakes the force of the first order evidence. The track record evidence available to Albert and Bertrand is exactly similar, and thus licenses equal degrees of justification of their respective higher order beliefs. So, the justification of higher order beliefs provided by this form of track record evidence does not underwrite Bottom Up.

It might be thought that there are other features that tend to make higher order beliefs justified. One idea mentioned in passing by Kelly is that when object level beliefs accurately reflect the evidence this is typically not due to mere chance, but due to cognitive faculties that generally produce accurate assessments of the evidence.⁴ One may wonder, however, how much this does for the plausibility of Bottom Up. Even if object level beliefs generally reflects the evidence in non-accidental ways, this does not show that higher order beliefs tracks first order evidential success in any way relevant for the justification of higher order beliefs. In reply to this, it may be said that this only shows only that higher order beliefs do not necessarily track first order accuracy. It could still be the case that higher order beliefs track first order accuracy well enough.⁵ Let us now consider this possibility. Generally, when we confidently believe a proposition on the basis of some evidence, we are strongly inclined to believe that our object level belief is a proper response to the evidence. Otherwise we would, on reflection, revise our object level belief, or be subject to an irrational form of cognitive dissonance. There is, however, a reason to doubt that this disposition to form positive higher order beliefs typically serves to make higher order beliefs highly justified. The

⁴ Cf. Kelly's remark: 'in paradigmatic cases in which one takes up the view that is best supported by one's evidence, it is no mere accident that one has done so (although lucky accidents are of course possible, they are atypical). Rather, one takes up the belief in question precisely because it is supported by one's evidence' (156).

⁵ Thanks to a reviewer for suggesting this.

reason is that this disposition tracks first order confidence, not actual accuracy. So, normal subjects would, when prompted, form such positive higher order beliefs about any confidently held object level belief, even in cases where object level processes fail in ways that are not discernible to them.

In response, one might point out that for all this, we might be reasonable good at tracking first order accuracy in virtue of being very good at tracking a proxy for first order accuracy, namely first order confidence. If true, this might generate some level of justification of higher order belief. It is doubtful, however, if this will support Bottom Up. Suppose that Anton and Bertrand are equally good though imperfect trackers of first order accuracy in virtue of tracking confidence as a proxy for first order accuracy, and that this is how their higher order beliefs get their justification. Suppose now that in a particular case, Anton responds accurately to a certain body of first order evidence, while Bertrand does not. They both equally confidently believe that their first order belief is accurate. If their higher order justification derives from their capacity to track first order success, then since this capacity is the same for both, their higher order beliefs will be equally justified. Accordingly, this view does not support Bottom Up.

3.1 The Argument from Recognition

The above review of some obvious sources of justification for higher order beliefs does not suggest that Bottom Up should be plausible. So, let us now consider Kelly's main argument for thinking that Bottom Up holds. Kelly writes:

'Indeed, in a given case, one might very well take up the belief because one *recognizes* that this is what one's evidence supports. Plausibly, *recognizing that p* entails *knowing that p*. Assuming that that is so, then any case in which one recognizes that one's evidence supports a given belief is a case in which one knows that one's evidence supports that belief. Clearly, if one knows that one's evidence supports a given belief, then one is justified in thinking that one's evidence supports that belief; if one were not justified, one would not know. But, even if recognizing that *p* does *not* entail knowing that *p*, one would in any

case not be able to recognize that p if one were unjustified in thinking that p . It follows immediately from this that, whenever one recognizes that one's evidence supports such-and-such a conclusion, one is justified in thinking that one's evidence supports that conclusion' (156).

Call this The Argument from Recognition. The argument can be spelled out in the following way (for simplicity I omit modifiers that Kelly may want to add that this holds only in some cases, not all cases):

(i) When S responds correctly to evidence e by forming credence n in p , then S recognizes that $\langle e$ supports p to degree $n \rangle$.

(ii) Recognizing a proposition entails knowing the truth of that proposition.

(iii) So, when S responds correctly to evidence e , S knows the truth of the epistemic proposition that $\langle e$ supports p to degree $n \rangle$.

Of course, knowing the truth of an epistemic proposition is a significant case of having a highly justified higher order belief, and S will have this highly justified higher order belief in virtue of having responded correctly to the first order evidence. So, if (iii) is correct a noteworthy version of Bottom Up holds.

An initial worry about the Argument from Recognition is that it is not obvious why responding properly to a body of evidence should be construed as involving entertaining an epistemic proposition. Think for example about cases in which we routinely and with no great effort respond correctly to some body of evidence, yet in which we cannot explain to ourselves what evidence we have actually responded to. Why think that appreciating evidence and responding by forming a belief involve entertaining epistemic propositions? This question is difficult to settle, however, as much will depend on what we think is involved in entertaining a proposition.

Setting this aside, suppose that appreciating evidence involves entertaining epistemic propositions. Then why think that a demanding epistemic state such as knowing is required? As we have seen, Kelly asserts that responding correctly to evidence involves recognizing a proposition and that recognizing entails knowing, as asserted in premise (ii). Kelly doesn't offer a distinct argument for these claims, but simply assumes that 'recognize' takes a propositional complement, and he also assumes a success-dependent use of 'recognize' according to which *S* recognizes the truth of *p* only if *S* knows that *p* is true. Certainly, there are uses of 'recognize' that conform to this. The question, however, is whether this particular use of the verb 'recognize' is appropriate for what normally goes on when we appreciate evidence at first order level, even in cases where we correctly and confidently respond to evidence. Whether this is so is a substantial philosophical question, and it cannot be settled simply by a choice of word for the relation. Kelly offers no independent argument to substantiate why 'recognize' as it is employed in the argument is the proper way to describe what normally happens in the kind of phenomenon we are trying to understand.

Independently of this, there are reasons to worry about the idea that correctly grasping the evidential import of some body of evidence *e* involves knowing the truth of an evidential proposition such as $\langle e \text{ supports } p \text{ to degree } n \rangle$. Knowing on the basis of what evidence, one may ask? One idea would be that to fully grasp of the evidential import of *e*, one needs independent evidence of the truth of $\langle e \text{ supports } p \text{ to degree } n \rangle$ strong enough for knowing this proposition. Apart from being independently questionable, we should note that this would not support Bottom Up, as it does not show assert a link between grasping first order evidence and higher order justification. Anton and Bertrand could both have the same independent evidence for an evidential proposition, yet Anton believes a correct evidential proposition and Bertrand a false one. So, the most promising view would claim that when *S* correctly grasps *e*, then $\langle e \text{ supports } p \text{ to degree } n \rangle$ is known by *S* on the basis of her grasp of *e*. If this is right, then the underlying assumption is that the following evidential supporting relation holds:

(E1) When e supports p to degree n , then e supports $\langle e$ supports p to degree $n \rangle$ to a degree high enough for knowledge.

Surely, we might question the prima facie plausibility of (E1). Why should evidence about some mundane state of affairs (say evidence about the whereabouts of my car keys) also be evidence for the truth of evidential propositions? Moreover, note that (E1) threatens to lead to a regress. Consider the iterated epistemic proposition implied in (E1), that is, the evidential proposition expressing the relation between e and the higher order proposition supposedly supported by e , $\langle e$ supports p to degree $n \rangle$:

$\langle e$ supports $\langle e$ supports p to degree $n \rangle$ to a degree high enough for knowledge

Consider now again premise (i) and (ii) in the Argument from Recognition. If S gets to know the evidential proposition $\langle e$ supports p to degree $n \rangle$ in virtue of grasping e , and e is the evidential support for that proposition, then it seems that by premise (i) S should recognize the iterated evidential proposition above, and by (ii) it follows that S should know this proposition. Again, we might wonder about the evidence upon which S knows the iterated proposition? Once more, it would seem that the only plausible answer is would be: e . So, in addition to (E1) we get the following:

(E2) When e supports p to degree n , then e supports $\langle e$ supports $\langle e$ supports p to degree $n \rangle$ to a degree high enough for knowledge

Again, this looks odd. Can mundane pieces of evidence really be this powerful? And, it seems, the list goes on. So, it looks like an infinite regress in levels of iterated epistemic propositions that firmly supported by a single piece of evidence.

This plays out differently if we assume that knowledge of epistemic propositions is a priori. So, on such a view (which I merely consider but do not attribute to anyone) when S grasps the import of e , this involves recognizing the

epistemic proposition $\langle e \text{ supports } p \text{ to degree } n \rangle$ where this epistemic proposition is a priori knowable. Consider the suggestion that S gains a priori knowledge that $\langle e \text{ supports } p \text{ to degree } n \rangle$ without even having scrutinized e . This, of course, is implausible. A priori knowledge does not come out of nothing. So, a credible version of the idea must be that S 's a priori knowledge of $\langle e \text{ supports } p \text{ to degree } n \rangle$ is based on S 's grasp of e . Yet, the relation between e and $\langle e \text{ supports } p \text{ to degree } n \rangle$ cannot be an ordinary evidential relation supporting a posteriori knowledge. Rather, there must be some other basing relation allowing S to gain a priori knowledge of $\langle e \text{ supports } p \text{ to degree } n \rangle$ in virtue of grasping e . This would seem to require the general truth of something like:

(E3) When e supports p to degree n then e grounds a priori knowledge that $\langle e \text{ supports } p \text{ to degree } n \rangle$

(E3) does not lead to the same regress as above, and maybe (E3) is plausible in other ways too. However, one may wonder why a different regress does not arise. The basic claim is that grasping e by forming the appropriate credence in p requires knowing the truth of an epistemic proposition such as $\langle e \text{ supports } p \text{ to degree } n \rangle$. Compare to obtaining a priori knowledge that $\langle e \text{ supports } p \text{ to degree } n \rangle$ by grasping the evidential import of e . Why doesn't this similarly require knowing the truth of $\langle e \text{ grounds a priori knowledge of } \langle e \text{ supports } p \text{ to degree } n \rangle \rangle$. If it does, then a new regress arises.

To avert these difficulties, one might consider running the argument in terms of epistemic justification, rather than knowledge, which is an option Kelly mentions in the above quote. One way of interpreting this remark is that recognizing the truth of a proposition requires being all-things-considered justified

in believing that proposition, rather than knowing it.⁶ We then get the following version of the Argument from Recognition:

(i) When *S* responds correctly to evidence *e* by forming credence *n* in *p*, then *S* recognizes that $\langle e \text{ supports } p \text{ to degree } n \rangle$

(ii*) *S* recognizes that $\langle e \text{ supports } p \text{ to degree } n \rangle$ only if *S* is all-things-considered propositionally justified that $\langle e \text{ supports } p \text{ to degree } n \rangle$

(iii*) So, when *S* responds correctly to evidence *e*, *S* is all-things-considered propositionally justified that $\langle e \text{ supports } p \text{ to degree } n \rangle$

Again premise (ii*) is, of course, demanding. As above, we might consider what could account for the truth (ii*). How, in the general case, does *S* become all-things-considered propositionally justified in epistemic propositions? One option would be that *S* needs independent evidence that make her all-things-considered propositionally justified, but for the reasons we have already seen, this would not work to support Bottom Up. So, again, the only workable idea would be that merely by grasping some body of evidence *e* one becomes all-things-considered propositionally justified in $\langle e \text{ supports } p \text{ to degree } n \rangle$. This would seem to presuppose a general principle along the following lines:

(E4) When *e* supports *p* to degree *n*, then *e* supports $\langle e \text{ supports } p \text{ to degree } n \rangle$ to a degree high enough for all-things-considered propositional justification.

⁶ Thanks to a reviewer for pressing this point. It's immaterial for the argument whether we read 'justified' as referring to all-things-considered propositional justification or something weaker.

(E4) probably does not lead to a regress, but it is still difficult to see what makes (E4) plausible. Note that (E4) should not be conflated with the plausible anti-akratic intuition that when rationally believing p on evidence e , one cannot at the same time believe (or be propositionally justified in believing) that e does not support p . This anti-akratic intuition does not support the basic idea in Bottom Up, which is that correctly grasping first order evidence *generate* higher order evidence. Rather, the anti-akratic intuition backs Top Down - the idea that higher order beliefs may destroy first order justification, even when first order justification reflects a correct grasp of first order evidence. Note also that, in E4 the relation between e and $\langle e \text{ supports } p \text{ to degree } n \rangle$ can hardly be the ordinary one that obtains between e and p . Grasping the evidential relation between e and $\langle e \text{ supports } p \text{ to degree } n \rangle$ is not like acquiring a posteriori justification in believing p on the basis of e . Maybe we should think of (E4) in terms of a priori propositional all-things-considered justification based on the grasp of e .

Now, even if we accept (E4), this would at most succeed in showing one half of Bottom Up. Recall, that Bottom Up asserts that correctly appreciating evidence makes one justified in believing the relevant epistemic proposition, but also that appreciating evidence correctly makes one justified in believing that one has indeed correctly appreciated the evidence. If a version of the Argument from Recognition based on (E4) is successful, this would only address the first claim. As noted above, knowing or being epistemically justified in believing an epistemic proposition does not entail knowing or being epistemically justified in believing that one's object level belief correctly reflects one's first order evidence. Hence, to get full support of Bottom Up, we would need to underwrite the idea that S 's grasping e correctly at t makes S at t justified in believing that S has grasped e correctly at t . As above, we can ask: by what means, or based on what evidence, does S become justified in this belief. Again, the answer has to be: e . For this to be true something like the following relation should hold:

e (and S at t has considered e) supports $\langle S$ at t has correctly appreciated $e \rangle$ to a significant degree.

Why should evidence about mundane fact in the world also serve as evidence that I have appreciated the evidence correctly when I have? While this is hard to believe, I will consider some arguments in favor of this view in the next section.

3.2 Evidentially transparent cases

Suppose that a body of evidence entails a proposition, and therefore warrants maximal credence in that proposition. Clearly, if the entailment is very complex and hard to understand it seems that one's higher order belief that one has responded correctly to the first order evidence might not be highly justified, even if one has in fact responded correctly. But consider now Kelly's remark that there are cases in which 'one's recognition that one's evidence supports a given belief is based on an unmediated appreciation of that evidence itself. Thus, in such cases, one's evidence not only confirms the belief in question, it also confirms a proposition to the effect that it is reasonable for one to hold that belief.' (160)

Consider first what I will refer to as *evidentially transparent* cases, by which I mean cases where one's object level belief is based on an unmediated appreciation of the evidence pertaining to that belief. Think, for example, of cases where one considers a fairly easy mathematical problem and clearly sees what the solution is, or perceptual cases where a certain visual input makes some factual proposition obvious. Kelly suggests that in evidentially transparent cases it is generally the case that one's evidence supports the truth of the relevant evidential propositions, and the proposition that one's object level beliefs accurately reflect the evidence. If this is right, it may provide some support to Bottom Up.

Though plausible at first sight, I nonetheless want to question that the appeal to evidential transparency can provide the defense of Bottom Up that we are looking for. To see this, suppose that Alf is a complete novice to music. A friend of his presses a key on the piano, and he tells Alf that the tone voiced is an F. A couple of hours later, someone again hits a key on the piano. Without seeing which key is pressed, Alf immediately think to himself: F. Suppose that the tone Alf hear this second time indeed an F, and that his spontaneous belief, far from being a mere guess, was due to a highly reliable capacity of his. In fact, Alf has perfect pitch.

Now, does the fact that Alf accurately appreciate the first order evidence, and does so in a way where the first order evidence is transparent to Alf, serve to make his higher order belief that his first order belief was indeed an accurate response to the evidence epistemically justified? It appears not. It seems that Alf would need more background evidence to be justified in his higher order belief. In particular, Alf would need additional evidence that he possesses this rather rare gift of perfect pitch, and maybe evidence that circumstances are normal, and so on.

I suggest that we should see Alf's case in the following way. There is a set of local evidence *e* relating to the tone played when Alf's friend strikes a key on the piano. When appreciating *e*, Alf immediately forms the belief that the tone played on the piano is an F, and Alf's accurate appreciation of *e* provides epistemic justification for Alf's belief that the tone played is an F. And we can say that *e* is evidentially transparent to Alf in the sense that upon grasping *e* it is immediately obvious to Alf that the tone played is an F. Clearly, however, it does not seem plausible to say that *e* thus construed, or Alf's immediate appreciation of *e*, also contributes to justification of Alf's higher order belief that his first order belief correctly reflects *e*. This is because *e* is evidence that the tone played is an F, and not evidence for the entirely different fact that Alf has perfect pitch. For Alf's higher order belief to be justified, Alf's total set of evidence would need to include additional evidence for Alf's believing that he is rather reliable in detecting pitch under the relevant circumstances.

In response to this one might opt for a more internalist construal of first order evidence. Suppose that Alf's first order evidence is not the tone played as such, but rather Alf's impression as of hearing an F being played, or Alf's auditory seeming that the tone played is an F. Construed this way we can plausibly say that Alf, upon appreciating his evidence, can form a justified higher order belief that his belief accurately reflects his evidence. After all, Alf's evidence is now his seeming that the tone played is an F, while Alf's belief is that the tone played is an F. Reflecting on this should be enough on its own for Alf to be highly justified in believing that his first order belief correctly reflect the first order evidence. So, arguably, on this internalist construal correctly appreciating first order evidence

makes one justified in higher order beliefs that one's first order beliefs accurately reflects the evidence.

Clearly, however, this internalist construal of evidence is not what we are after. To see this, recall the dialectical role of Bottom Up in the disagreement debate. Suppose that Albert responds properly to first order evidence, but Bertrand mistakes the force of the first order evidence. The question is whether this generates an asymmetry in the degree of justification at the higher order level. On the internalist construal of these cases, if both Albert and Bertrand respond in accordance with how things seem to them, then their higher order belief that they have responded correctly are equally justified, even if Albert is right and Bertrand wrong about their first order beliefs. If this is the way that Bottom Up becomes true, then it doesn't have the significance we took it to have.

One might worry that Alf's music case is very special for two reasons. First, it is unusual in the sense that it features a reliable perceptual faculty that most of us do not have, and where it therefore might be natural to think that one would need additional evidence for thinking that one has this faculty. Second, it is a perceptual case, as distinct from a case in which one reasons about the evidential import of a body of propositions.

Consider an ordinary perceptual case, such as visual perception. Surely, there are many features making us justified in believing that our visual perception is reliable under normal circumstances. But apart from this, the cases seem to be analogous. What makes me justified that a particular visual belief is a correct response to a body of visual evidence is not this visual evidence itself, but the more general reasons I have to believe that my visual perception is accurate. The same, I suggest, holds for cases involving reasoning rather than perception. Suppose that e is a body of mathematical evidence strongly suggesting the truth of a mathematical proposition p . Suppose that the evidential relation between e and p is a non-trivial entailment relation that it takes some effort and mathematical aptitude to see, but which is nonetheless transparent for those who invest the effort and are sufficiently skilled. Could it be that the mathematical evidence e is also evidence for the truth of higher order beliefs such as the belief in $\langle$ my belief that p correctly reflects the evidence e $\rangle$? It is hard to see why this should be the

case. In the ordinary case, where e is some mathematical evidence pondered by a subject S , e does not have the appropriate entailment relations to any facts about what doxastic states S ought to have. So, if we think of unmediated appreciation of mathematical evidence as immediately grasping an entailment relation, it is not as if there is an additional entailment relation the grasp of which makes one justified in a higher order belief that one's object level belief correctly reflects the evidence.

So, in general, it is difficult to see why, when some body of evidence e warrants credence n in p , then e is also evidence for $\langle e \text{ warrants credence } n \text{ in } p \rangle$, or the belief that any particular doxastic attitude regarding p is epistemically justified.⁷ This appears also to be true in cases where some evidence e in a vivid and immediate way supports p .

Yet, when considering simple cases of elementary reasoning, it surely seems that there is something right to the idea that Bottom Up may hold in cases of evidential transparency. Suppose that Andy, while driving his car, wonders how many apples he has in his basket in the back of his car. Andy then remembers for sure that he collected 5 green apples and 7 red apples, and that they are the only apples in his basket. As he accurately appreciates this evidence, Andy becomes rationally justified in his belief that there are 12 apples in his basket. Andy now considers his belief that there are 12 apples in the basket. Reflecting on this belief, Andy becomes rationally confident that his belief is indeed the proper response to the evidence at hand. So, by grasping the evidence correctly, and reflecting on this, Andy earns justification for believing that he has responded properly to his evidence. So, at some version of Bottom Up must be correct for at least some class

⁷ There are obviously special cases. Suppose I confidently believe that at least one of my beliefs is confidently held. It would seem that the very content of this belief makes it reasonable for me to believe that this belief of mine is highly justified. This case is special, of course, as it involves a first order beliefs whose very content makes it directly evidentially relevant for a certain higher order belief. Ordinary cases are not like that, so I set these special cases aside.

of evidentially transparent cases involving reasoning. Call this the Argument from Simple Reasoning.

I will argue, however, that even this case does not lend support to Bottom Up. It is clear how Andy's object level belief <there are 12 apples in the basket> becomes justified by being based on Andy's proper appreciation of the evidence. But the question is how Andy gets from there to being justified in his higher order belief that this object level belief correctly reflects his evidence. Clearly, Andy knows or justifiably believes <there are 12 apples in the basket>, and he knows that if there are 12 apples in the basket, then his belief <there are 12 apples in the basket> is true. But it doesn't follow from this that Andy knows or is justified in believing that his belief <there are 12 apples in the basket> is true. This is because, for any proposition *a* and *b*, knowing or justifiably believing that *a* entails *b* in conjunction with the mere *truth* of the antecedent *a*, does not entail that one knows justifiably believe the consequent *b*. Here is a way to see exactly what is missing in Andy's reasoning:

- (1) Andy at *t* accurately appreciates *e* by adopting credence *n* in *p*
- (2) Andy knows that (1) entails that his higher order belief in <*S* at *t* accurately appreciates *e* by adopting credence *n* in *p*> is true.
- (3) Transmission of justification across known entailments. For all propositions *a* and *b* and agents *S*: if *S* is justified in believing *a*, and *S* knows that *a* entails *b*, then *S* is justified in believing *b*.
- (4) From (1), (2), (3): At *t*, Andy is justified in his higher order belief in <that Andy at *t* accurately appreciates *e* by adopting credence *n* in *p*>.

We can grant that premise (1) is true in virtue of how the case is described. Surely premise (2) in the argument is plausible, and let us for the sake of argument accept premise (3) also, though a fully plausible transmission principle would need a

number of qualifications. Yet, as it stands, the argument is not valid, and the reason should be clear. For the argument to be valid it is not enough that (1) is true; Andy also needs to be justified in believing premise (1) - this is what the transmission principle in (3) requires. So, for desired conclusion (4) to follow, we need not the truth of (1), but rather the truth of

(1*) Andy at t is justified in believing that he at t accurately appreciates e by adopting credence n in p .

But clearly, (1*) is controversial in this context. While it is clear that Andy's correct appreciation of his evidence makes him justified in believing <there are 12 apples in the basket>, it not obvious that Andy's correct appreciation of that same evidence makes him justified in believing that he has correctly appreciated this evidence. But this is what the argument requires, if it is to be based on the transmission principle in premise (3).

I have argued that it is far from clear that evidentially transparent cases underwrite Bottom Up. But consider now what we might call *evidentially super-transparent cases*. By this I mean cases where S's immediate appreciation of first order evidence e makes obvious an epistemic proposition concerning the extent to which e evidence supports p , and where e is also is evidence that for S's higher order belief that her first order belief regarding p correctly reflects e .⁸ Clearly, Bottom Up, or a version of it, holds for super-transparent cases, as this is simply a matter of how they are defined. Now, this shifts the question to whether our actual evidential situation often of ever involves evidentially super-transparency. I don't think that it is obvious at all that there are many super-transparent cases, or perhaps any at all. In part this is for the reasons mentioned earlier: in general, it is hard to make sense of the idea that when e supports p to

⁸ I read Kelly's remark above as asserting a claim about evidentially transparent cases. But clearly, one might also read him as asserting the existence of evidentially super-transparent cases.

degree n then e supports $\langle e$ supports p to degree $n \rangle$, and the idea that evidence e can support both a belief, and the higher order stance that one has appreciated e correctly.

Let me conclude this section by addressing two worries. First, the format for epistemic propositions that I have used throughout (expressed by ' e supports p to degree n ') suggests that beliefs in epistemic propositions are fine-grained. Surely, fine-grainedness might pose an independent problem for how we might become highly justified in such beliefs. Seeing your scarf makes me justified in believing that your scarf is blue, but not justified in believing that it has some very specific shade of blue that I can hardly distinguish from other shades. It might be thought that the challenges I have raised for Bottom Up depends on epistemic propositions being fine-grained.⁹ In response to this, note that the arguments in the paper do not make use of the intuitively attractive idea that it would be difficult to be highly confident about the truth of a fine-grained epistemic proposition. The recurring problem is that it is difficult to make sense of the idea that appreciating object level evidence contributes to higher order justification in the way that Bottom Up requires.

The second worry is that while I may have shown that first order accuracy does not confer doxastic justification to higher order beliefs, it still yields propositional justification. I assume that doxastic justification concerns the epistemic justification of one's actual doxastic states, whatever it exactly means for a doxastic state to be epistemically justified. Propositional justification, by contrast, concerns what one would be doxastically justified in believing if one were sufficiently rational, attended to one's evidence, and formed the appropriate beliefs on the basis of that evidence. So, roughly, a subject S is propositionally justified with respect to a proposition p if and only if S possesses a body of evidence e , such that if S properly grasped e and based her belief that p on e , then S would be doxastically justified in her belief that p . There are, of course, many questions of the detailed interpretation of this. I suggest, however, that if what I

⁹ Thanks to a reviewer for bring up this issue.

have argued is correct, then we should reject the idea that first order accuracy constitute evidence making higher order beliefs propositionally justified. It is not as if we were subjects only to focus on the first order evidence available to them, or were they a little smarter, then they would tend to become more doxastically justified in their higher order beliefs whenever and because they appreciate their first order evidence correctly. Rather, if what I have argued is correct, there are no evidential routes from first order accuracy to increased justification of higher order belief, so first order accuracy makes us neither doxastically nor propositionally more justified in our higher order beliefs.

4. Asymmetric epistemic push

If what I have argued so far is correct we should reject Bottom Up. If we still accept Top Down - the idea that justified higher order beliefs can have a significant impact on the justification of first order belief - then there seems to be a significant asymmetry in that upward connections and downward connections work in quite different ways. If correct, this may have significant implications for the epistemology of disagreement and for how Integration is to be understood. In this section, I will elaborate on this asymmetric epistemic push.

Consider first downwards connections. It is impossible here to discuss all types of cases in which one might think that higher order beliefs exerts some epistemic influence on the justification of first order belief, so I will focus only on the most intuitively compelling. Consider a case modified from Christensen (2011): Adrian believes that p upon inspecting a body of mathematical evidence e . Now a source that Adrian fully trusts tells him that he has ingested a pill that contains a reason-distorting drug. The drug will affect a limited set of his cognitive faculties in ways that he cannot detect, and this will lead him to make very severe mistakes in his assessment of e , though Adrian will seem to himself to be thinking as clearly and discerningly as ever.

Intuitively, it seems that Adrian should reduce confidence in his object level belief. But why don't first order beliefs and higher order beliefs affect one another in both directions? After all, both first order beliefs and higher order beliefs figure in the same epistemic subject, and in cases like these they conflict in

ways that call for some form of rational adjustment. Why isn't the adjustment going in either direction? Surely, Adrian's first order mathematical belief that p and the evidence for this belief do not imply or provide direct evidence that Adrian has not ingested the reason-distorting drug. But similarly, Adrian's higher order belief that he is subject to the drug does not imply or provide evidence that p is false. So, why is there a downward connection when there is no upward connection?

A suggestion is that the asymmetry is related to how we rationally reflect about our beliefs. By reflecting about a belief I simply mean wondering and reasoning about the truth of the belief and about the implications of the belief for one's other beliefs. Intuitively, it seems very clear that if one has a justified negative higher order belief and reflects about the implications of this, then one should respond by reducing confidence in one's object level belief. So, the case above suggests that we should accept something like the following:

(Top Down Rationality) If S has evidence e' that makes her justified in believing that her object level belief that p is not a proper response to S 's object level evidence e for p , then S should modify her rational credence in p .

To be sure, (Top Down Rationality) is not uncontroversial and this and related ideas have come under attack in recent discussions, though for reasons that are unrelated to the problems affecting Bottom Up discussed in this paper.¹⁰

However, reflecting on an object level belief does not similarly create an upward push. Consider a variant of the above case. Beth has considered a body of mathematical evidence e , and correctly formed credence n in p . Beth then

¹⁰ For discussion and defenses of similar principles, see (Horowitz 2014; Sliwa and Horowitz 2015b). For critical discussions of the idea that higher order evidence can affect first order rationality, see for example (Lasonen-Aarnio 2014; Titelbaum 2013; Weatherson ms).

happens to overhear a conversation in the philosophy department in which the possibility of a reason-distorting drug is mentioned. Beth accepts the idea of a reason distorting drug that affects a limited set of her cognitive faculties in ways that one cannot detect, and that such a drug would lead her to make very severe mistakes in her assessment of *e*. Beth now starts reflecting about her first order belief in *p*, and specifically wonders if that belief or the evidence *e* supporting it gives her any reason to think that she is not under the influence of a reason distorting pill of the kind in question. Clearly, it is very hard to imagine that Beth can become justified in her belief that she is not under the influence of a reason-distorting drug by reflecting about her first order belief, or the evidence supporting that belief. There are two basic reasons why this is so. The first is that the Beth's first order belief *p* and her evidence *e* does not bear on the truth of the relevant higher order belief. The second reason is that even if the first order evidence did bear on the truth of the higher order belief, there would seem to be something question-begging about using the evidence and belief whose very propriety is in question as the starting point for a chain of reasoning that purportedly shows that Beth is justified in believing that that belief is a proper response to the evidence.

So, one might suggest that what grounds the asymmetry is this. Even when, as in Beth's case, higher order beliefs do not imply anything about the *truth* of first order beliefs, they may have implication for the *rationality* of holding those higher order beliefs. Ordinary first order beliefs, by contrast, do not in similar ways have implications for the *rationality* of higher order beliefs, as there is no principle similar to (Top Down Rationality) connecting first order beliefs to the rationality of holding higher order beliefs. To be sure, I haven't provided any argument for (Top Down Rationality), except for what is involved in presenting a couple of intuitively compelling cases, and neither have I responded to the criticism of such principles.¹¹ But my aim is the more limited one of pointing to an

11

See earlier footnote.

asymmetry suggesting that one can reject Bottom Up while still accepting Top Down.

5. The Fragility of Justification of Higher Order Beliefs

I now want to consider what I will refer to as the fragility, or lack of robustness, of justification of higher order beliefs. Though I rejected this view above, suppose for the sake of argument that accurate assessment of first order evidence directly contributes to higher order justification in evidentially transparent cases. Clearly, not all cases are as transparent as those discussed above, but we might think of evidential transparency as a matter of degree, all cases displaying some even if only a minimal degree of evidential transparency.¹² We could then accept a graded version of Bottom Up according to which when S believes p on the basis of e , and this reflects an accurate appreciation of e 's evidential import for p , then S 's higher order belief that S has accurately grasped e is justified proportionally to the degree to which e is evidentially transparent to S .

I want to focus on what I will refer to as the fragility of higher order justification derived from this source. Suppose that a body of evidence e for a proposition p is evidentially transparent to me, such that by accurately assessing the evidence I generate a significant degree of higher order justification. I then receive some defeating evidence e' . Suppose, for example, that I am told by a trustworthy source that I have ingested a drug making my assessment of e wildly inaccurate, though everything appears normal to me. Or suppose that I come across evidence indicating that e doesn't support p after all. Surely, in both cases I will thereby have received an undercutting defeater of my object level justification. Now, I want to argue that that very same undercutting evidence would *also* seem to defeat my higher order justification in so far as this derives entirely from the transparency of my first order evidence. In this sense, my higher order justification, when originating in evidential transparency, fails to be robust.

¹²

Thanks to a reviewer for making this suggestion.

To see this, consider the specifics of evidential transparency and defeating evidence. Take evidential transparency first. Assume that *S* is in an evidential transparent situation with respect to *e* and *p*. Suppose, moreover, that because the situation is evidentially transparent, it further holds that *e* is evidence for $\langle e \text{ warrants credence } n \text{ in } p \rangle$, and that *e* is evidence for the truth of *S*'s higher order belief that her object level belief in *p* reflects an accurate assessment of *e*. Consider then undercutting evidence at the object level. Suppose that *S*'s evidence *e* warrants credence *n* in *p*. Relative to this, *e'* is undercutting defeater if either (i) *e'* is evidence for the falsity of $\langle e \text{ warrants credence } n \text{ in } p \rangle$, or (ii) *e'* is evidence that *S*'s capacity to appreciate *e* is not reliable in the relevant circumstances.

Now, attending to the details of this, we can see that when *e'* is an undercutting defeater of *S*'s first order belief, then *e'* is also a defeater of *S*'s higher order justification derived by the evidential transparency of *e*. This is because *e'* undercuts either the evidential link between *e* and $\langle e \text{ warrants credence } n \text{ in } p \rangle$ or the evidential link between *e* and the truth of *S*'s belief that her assessment of *e* is accurate. In other words: when *e'* undercuts first order justification by indicating the falsity of $\langle e \text{ is evidence for } p \rangle$, then *e'* also undercuts higher order justification by indicating the falsity of $\langle e \text{ is evidence that } \langle e \text{ is evidence for } p \rangle \rangle$, which is precisely what we assumed to be true of evidentially transparent cases. Similarly, if *e'* undercuts justification of first order belief by indicating that *S*'s capacity to assess *e* is not reliable, then *e'* also indicates that *e* is not good evidence for it being true that *S*'s capacity to assess *e* is reliable in the circumstances, again a relation that we now suppose to hold in evidentially transparent cases. So it seems that when one gathers undercutting defeating evidence for first order evidence, then one thereby gets evidence undercutting the higher order justification that one might receive through evidential transparency of that first order evidence. As I said, this is a sense in which higher order justification is fragile.

I now want to suggest an implication of this for disagreement cases. It will be helpful first to consider a simpler case. Suppose that Aron is highly confident that *p* on the basis of *e*. Aron then encounters additional rebutting and undermining evidence, that is, evidence that his object level belief that *p* is false, or evidence that his object level belief is based on a process that is unreliable or in

some way flawed. Now, as it happens, Aron has a very solidly justified higher order belief his first order beliefs almost always accurately reflect the evidence in these types of cases. Assume, for example, that Aron's ability to understand and evaluate the type of evidence in question has been thoroughly tested, and that it has turned out that Aron has an outstanding and highly reliable ability to assess evidence accurately in these cases, and that Aron knows this. It seems natural to assume that Aron's highly justified higher order belief can defeat the object level defeaters he encounters, and when it does, it is rational for Aron to retain his high confidence that p .

Suppose now that Albert and Bertrand consider some evidence e independently of one another. Assume that their cases are evidentially transparent. It turns out that, upon considering the evidence Albert firmly believes p , whereas Bertrand rejects p . Assume that Albert is right: p is true, and e does indeed support high confidence in p . Given what we have assumed about evidential transparency, since Albert appreciates e correctly and his situation is evidentially transparent, this tends to contribute to his higher order justification, whereas this is not true for Bertrand. So, Albert and Bertrand are not on a par on the higher order level. Assume next that they discover the disagreement they have with one another. It might seem that Albert and Bertrand's situations are now asymmetrical, since Albert has correctly appreciated the force of e , but Bertrand has not. So, the asymmetry at the first order level is reflected by an asymmetry at the higher order level. If this were true, then Albert's situation would be like Aron's - Albert's high degree of higher order justification would entitle him to dismiss the evidence constituted by his disagreement with Bertrand, whereas Bertrand would not be similarly entitled.

However, if what I have argued above is correct, we might deny that the asymmetry at the higher order level remains in full force after the disagreement is disclosed. The reason is that the disagreement is evidence of some weight for Albert that his first order process might be defective in some way, either because e doesn't support p to the extent that Albert's confidence requires, or because Albert might have made a mistake in assessing the evidence. I argued that evidence that tends to defeat the justification of first order belief also tends to

defeat higher order justification regarding that belief, when the higher order justification derives from evidential transparency. If this is correct, then Albert's higher order justification will be affected by the disagreement. Even if Albert's higher order belief is more justified than Bertrand's prior to disclosing their disagreement, this might change when they learn of their disagreement. If correct, Albert may no longer be in a position like Aron's, whose high degree of justification of a higher order belief entitles him to dismiss incoming first order defeating evidence.¹³

I have argued that higher order justification may be defeated by first order defeaters when it depends on first order evidence. Compare to higher order justification that is due to track record evidence. A general form of track record evidence appeals to normal circumstances. In so far as we can assume that circumstances are normal and our cognition functions properly, we have some evidence suggesting that our beliefs are accurate responses to the evidence. A more demanding form of track record evidence appeals to specific facts about more finely individuated belief-forming processes. I may know that a belief of mine is formed in circumstances C by a particular process M, and have independent evidence that M tends to respond very accurately to the evidence in those circumstances. Both forms of track record evidence are independent of first order accuracy in a particular case, and for this reason they do not underwrite Bottom Up. However, for exactly that very same reason higher order justification deriving

¹³ This argument presupposes that disagreement is a first order defeater, a claim that some might reject. My impression, though, is that most people do accept that evidence of disagreement has the form suited for a first order defeater. What they disagree about is whether this defeater is then itself defeated or otherwise upset by other features, which might be asymmetrical in various disagreement cases. One such asymmetrical defeater-defeater could be differences in degrees of higher order justification. My claim is that this particular claim is not plausible if differences in higher order justification are held to be due to evidential transparency.

from track record evidence is more robust in certain respects. To see this, suppose S has solid independent evidence e^* for her general capacity to assess evidence correctly in a particular type of situation. S is now in that type of situation, considers some body of evidence e which she takes to support p to degree n . Suppose now that S receives some new bits of evidence e' indicating either that $\langle e$ warrants p to degree $n \rangle$ is false, or that S is not capable of assessing the evidence accurately in the particular case. S's situation is now that she has some evidence e^* indicating that she is capable of assessing e correctly, and some evidence e' indicating that she is not. The balance of these two conflicting bodies of evidence determines her justification of her higher order belief. However, it is not that e' undercuts e^* , like what we saw above in the case above assuming that higher order beliefs were justified by evidential transparency. Rather, this new evidence e' is a rebutting defeater relative to e^* . So, evidence of first order failure provided by cases of disagreement does not undercut higher order justification that depends on track record evidence in the way that it undercuts higher order justification that depends on evidential transparency. In that sense, track record evidence is more robust, though of course even such evidence can be overridden.

Before concluding the paper, I want to address the general worry that much of what I have argued depend on *Independence*, or related controversial ideas that one is rationally required to bracket certain forms of evidence.¹⁴ As stated by Christensen, Independence is the following principle:

‘In evaluating the epistemic credentials of another’s expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn’t rely on the reasoning behind my initial belief about P.’ (Christensen 2011)

As it has often been pointed out, Independence seems implausible for a variety of reasons. Independence seems to commit agents to simply ignore their first order

¹⁴

Thanks to a reviewer for bringing up this worry.

evidence in cases of disagreement, and may also be implausible in cases involving high degrees of rational certainty, or when one's beliefs depend on a wide range of other beliefs that a simultaneous subject to disagreement (cf. (Lackey 2010; Sosa 2010; Kelly 2013)). It might therefore be considered a weakness if the arguments above somehow relied on Independence, or similar requirements that evidence be bracketed.

In response to this, note that the reasons for rejecting Bottom Up advanced in section 3 concerned the difficulty of establishing appropriate relations between correct appreciation of object level evidence and higher order beliefs. The asymmetry discussed in section 4 was suggested to depend on (Top Down Rationality). One might worry that (Top Down Rationality) actually commits to Independence. But clearly it does not. (Top Down Rationality) says that if one has evidence suggesting the impropriety of a first order belief, then one should modify one's rational credence in that belief. This does not commit to the antecedent being true in cases of disagreement. Neither does (Top Down Rationality) say anything about bracketing. Another worry might be that the arguments that there are no reasoning principles similar to (Top Down Rationality) depend on Independence. But this worry is unfounded, as this rejection was based on two ideas not committing to Independence. The first idea was that there generally is evidential import of first order evidence on the truth of higher order belief. The second idea was that if you believe p on the basis of e , and wonder whether your belief correctly reflects the evidence, you cannot reason your way to an answer by simply repeating to yourself that you believe p on the basis of e . This is, of course, a sort of bracketing, but it is much more limited and circumscribed than Independence, and it doesn't commit to the implausible implications of Independence. Finally, one might be concerned that the idea that higher order justification may be fragile again depends on Independence, or similar controversial assumptions about bracketing. But here the argument was the quite specific observation that evidence undercutting first order justification also constitute evidence undercutting higher order justification in the special circumstance assumed to hold in evidentially transparent cases: that the justification of higher order belief depend on correct appreciation of first order

evidence. This specific argument does not depend on Independence or controversial ideas of bracketing.

6. Concluding remarks

Let me conclude by summarizing the main findings. I have argue that we should reject Bottom Up as it is difficult to explain how first order accuracy could bring about higher order justification. These problems with Bottom Up leave Top Down unaffected. Thus, we should reject Bottom Up, but we might still accept Top Down. There might be a general asymmetry in the way that object level and higher order levels affect one another. Object level beliefs accurately reflecting the evidence do not in any clear and convincing way impact the degree of justification of positive higher order beliefs. For all this, negative higher order beliefs can still significantly impact the justification of object level beliefs.

Of course, our higher order beliefs might arise and be justified in a variety of ways. Often our higher order beliefs arise from our spontaneous inclination to form positive higher beliefs as a complement of confidently held first order belief. This process itself does not provide higher order beliefs with any high level of justification. Though I have argued that this is implausible, someone might insist that at least in evidentially transparent cases, higher order beliefs earn some degree of epistemic justification from first order evidence being evidentially transparent. Even if this were true, higher order justification obtained in this way would not be very robust as object level undercutting evidence would generally also be higher order undercutting evidence. When higher order justification depends on track record evidence, it is more robust vis-à-vis the defeating effect of object level defeating evidence, and this is because track record evidence justifying positive higher order belief does not depend on object level evidence. However, just how robust track record evidence is seems to depend on details of the nature of the track record evidence.

References

Christensen, David. 2007. "Does Murphy's Law Apply to Epistemology? Self-Doubt

- and Rational Ideals." *Oxford Studies in Epistemology* 2: 3–31.
- . 2011. "Disagreement, Question-Begging and Epistemic Self-Criticism." *Philosophers Imprint* 11 (6). www.philosophersimprint.org/011006/.
- . 2013. "Epistemic Modesty Defended." In *The Epistemology of Disagreement*, edited by David Christensen and Jennifer Lackey, 77–97. Oxford: Oxford University Press.
- Horowitz, Sophie. 2014. "Epistemic Akrasia." *Nous* 48 (4): 718–44. doi:10.1111/nous.12026.
- Kelly, Thomas. 2010. "Peer Disagreement and Higher Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted A Warfield, 111–74. Oxford University Press.
- . 2013. "Disagreement and the Burdens of Judgment." *The Epistemology of Disagreement: New Essays*, 31–53. doi:10.1093/acprof:oso/9780199698370.003.0003.
- Lackey, Jennifer. 2010. "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford University Press.
- Melis, Giacomo. 2014. "Understanding Undermining Defeat." *Philosophical Studies* 170 (3). Springer Netherlands: 433–42. doi:10.1007/s11098-013-0238-z.
- Sliwa, Paulina, and Sophie Horowitz. 2015a. "Respecting All the Evidence." *Philosophical Studies* 172 (11): 2835–58. doi:10.1007/s11098-015-0446-9.
- . 2015b. "Respecting All the Evidence." *Philosophical Studies* 172 (11). doi:10.1007/s11098-015-0446-9.
- Sosa, E. 2010. "The Epistemology of Disagreement." In *Social Epistemology*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford University Press.
- Sturgeon, Scott. 2014. "Pollock on Defeasible Reasons." *Philosophical Studies* 169 (1). Springer Netherlands: 105–18. doi:10.1007/s11098-012-9891-x.
- Titelbaum, Michael G. 2013. "Rationality's Fixed Point (or: In Defense of Right Reason)." *Oxford Studies in Epistemology* 5.
- Weatherson, Brian. 2016. *Normative Externalism*.
- Worsnip, Alex. 2015. "The Conflict of Evidence and Coherence." *Philosophy and*

Phenomenological Research, September, n/a-n/a. doi:10.1111/phpr.12246.

Aarnio, Maria Lasonen. 2014. "Higher-Order Evidence and the Limits of Defeat."
Philosophy and Phenomenological Research LXXXVIII (2).
doi:10.1111/phpr.12090.