

Models of Rational Inference: Incorporating Higher-order Evidence

JAKE STONE

AUSTRALIAN NATIONAL UNIVERSITY

Abstract

The aim of this paper is to develop a model of rational inference incorporating higher-order evidence (HOE). I conduct this analysis in three stages. First, I show that HOE is evidence of a distinct type as it is agent-relative and undermines an agent's rationality. Here I largely follow the position adopted by David Christensen, which I take to give strong *prima facie* plausibility to the idea that there is HOE and that it should modify our initially-formed beliefs by forcing us to bracket our first-order evidence (FOE). Secondly, I discuss some possible models of rational inference incorporating HOE. In particular, steadfastness, in which one prefers FOE, and calibrating, which favours HOE. Finally, I mediate between steadfastness and calibration in order to determine whether Christensen's position can be sustained. The key question is: how exactly should first-order and higher-order evidence interact? If we cannot give a good account of the interaction, then Christensen's position looks less tenable. I argue that if we keep in mind the different perspectives from which we can evaluate rationality – internally and externally – a defensible model emerges.

1. Introduction

To understand higher-order evidence (HOE), let's first consider an example of such evidence:

Fatigue

You are a highly trained army commando conducting operations behind enemy lines. After three sleepless days and nights, you have located the enemy's secret headquarters. Unfortunately, it is near a school. Using a map and your geographical knowledge, you determine (correctly) the precise coordinates of the headquarters. Having become highly confident of the location you prepare to radio in an airstrike. Before you do, your second-in-command reminds you that you haven't slept for days and that fatigued officers often make errors of calculation.

In this scenario it seems prudent to re-evaluate the relationship you perceive to hold between your evidence and the hypothesised location of the base. This holds even where you have determined the location correctly. The higher-order evidence about being fatigued appears to undermine your ability to utilise your first-order evidence, E, in support of a hypothesis, H, about the location of the enemy base. To be precise, it isn't evidence about the actual relationship between E and H. Rather, it is evidence that causes you to re-evaluate whether you have perceived the relationship between E and H correctly. The prevailing consensus is that higher-order evidence raises the possibility that you may have made a cognitive error in your calculations utilising E in support of H.¹ As a result, the higher-order evidence of being fatigued seems to be evidence that is relevant to the hypotheses under consideration.

But then again, there is some doubt over whether HOE is even evidence at all. That is, one might doubt that higher-order evidence is the type of thing that has any bearing on your beliefs in the hypotheses under consideration. Why should being fatigued affect the location of the enemy headquarters? If HOE is evidence, it is at the very least evidence of a unique kind. Consequently, the distinction between HOE and FOE gives rise to two interesting questions. First, is HOE really a unique kind of evidence? And secondly, if it is unique, what relationship does HOE bear to FOE? This essay will answer these questions in three sections. First, Section 2 explains why there is a strong *prima facie* case to conclude that HOE really is evidence of a unique kind. Section 3 then considers the different models which have been proposed to deal with HOE. The original contribution of this paper occurs in Section 4 where I mediate between these proposed models. I argue that, if we keep in mind the different perspectives from which we can evaluate rationality, a defensible model emerges. More specifically, I argue for a model I call J*-calibration which allows us to vindicate our desire to judge in accordance with our first-order evidence while taking the presence of higher-order evidence seriously.

2. A *Prima Facie* Case for Higher-order Evidence

I am largely in agreement with Christensen, and other authors, that HOE is a unique kind of evidence which is epistemically significant.² Consequently, the following discussion is expository and relies upon Christensen's arguments as *prima facie* evidence that HOE is unique in that it is (1) agent-relative, and (2) appears to undermine an agent's rationality by requiring us to "put aside or bracket" our FOE.³

2.1 Agent Relativity

¹ Christensen 2010, p. 186; Lasonen-Aarnio 2014, p. 315. It is also worth noting that some HOE might raise the possibility that you are thinking especially accurately. For example, an agent who has taken a neurotropic.

² Christensen 2010. For other authors who also argue that HOE can cause us to bracket our FOE see Feldman 2005, Lasonen-Aarnio 2014, Leydon-Hardy 2016, and Schechter 2013.

³ Christensen 2010, p. 195.

HOE is unique as it is agent-relative. In the ‘Fatigue’ case, the higher-order evidence undermines your belief that H. However, for any other agent (who is well rested) with belief H and evidence E, the higher-order evidence will have no such effect. For why should someone else being fatigued affect their belief in H? This, argues Christensen, is contrary to our normal conception of evidence where the effect does not depend on who the agent is (though it may differ depending on an agent’s prior beliefs or background evidence).

2.2 HOE Appears to Undermine an Agent’s Rationality

2.2.1 Bracketing of FOE

HOE is also unique for its ability to undermine an agent’s rationality by preventing them from giving first-order evidence its full due, by undermining their belief in H based on E. Once aware that you are fatigued, you can no longer update your beliefs on the basis of E even though it seems rational to do so. That is, you can’t give E the full consideration that rationality seems to require. This is so whether the agent has committed an error or not. Recall that in the ‘Fatigue’ case, despite your lack of sleep, you calculated the coordinates of the enemy headquarters correctly. However, as you are unaware of your correctness, it still seems prudent to re-evaluate your credence in the coordinates (especially considering the innocent lives at stake). Nevertheless, this does not necessarily require you to reduce your credence. Reduction is problematic as it requires you to increase your credence in the alternative hypotheses and the HOE is silent on these alternatives.⁴ So, HOE prevents you from giving your FOE its full due, even where that FOE appears to support your original conclusion, and where it may actually do so. In this way, HOE requires an agent “to at least to some extent, *put aside* or *bracket*” their evidence.⁵

2.2.2 Contradicts Conditionalisation

Conditionalisation is the process by which an agent updates their credences in light of new evidence. For example, suppose you have a coin which you assume to be fair. Suppose you then learn that the coin was been flipped 100 times and never landed on heads. In light of this new evidence, E, it is rational to revise your credence in the hypothesis that the coin is fair. In particular, an agent’s new credence should match the credence she would have assigned to her hypothesis had she initially supposed that E was the case. More formally, the principle of conditionalisation requires that: a rational agent possessing an initial credence in some hypothesis $Pr_{\text{initial}}(H)$ who subsequently gains new evidence E, should update their beliefs so that their new credence $Pr_{\text{new}} = Pr_{\text{initial}}(H \mid E)$.

So how does the principle of conditionalisation relate to HOE? Well, once we come to possess HOE, we must re-evaluate our credence in the coordinates of the enemy

⁴ My thanks to Katie Steele for bringing this point to my attention.

⁵ Christensen 2010, p. 195.

headquarters we determined on the basis of E. However, we also know that prior to possessing the HOE, conditionalisation would recommend something different. Suppose that prior to departing on your mission you learnt the same evidence, E, used in the original ‘Fatigue’ case. Let us also suppose that E entails the coordinates H of the enemy headquarters. If E entails those coordinates, then conditionalising on E should lead us to increase our credence in H. Similarly, if in three days’ time, you come to learn E plus evidence, F, that you are fatigued, conditionalisation recommends you adopt the same credence in H, as F has no impact on the relationship between E and H (E entails H whether you are fatigued or not). But we have already said that in three days’ time when you learn E and F, rationality requires you to re-evaluate your confidence in H. Schoenfield and Christensen agree that this result contradicts conditionalisation, which says that the credence we should adopt in H, supposing we learn E & F, should be the same as the credence we adopt when we actually learn E & F.⁶

3. Possible Models for Dealing with HOE

Before considering whether Christensen’s characterisation is defensible, I first want to canvas the possible models of rational inference incorporating HOE. When considering how first-order and higher-order evidence interact, we have two models to choose from: *steadfasting* and *calibrationism*. In what follows, I adopt the definitions of steadfasting and calibration utilised by Schoenfield.⁷ Steadfasters maintain that an agent’s total evidence (E+HOE) supports the same credence that the agent’s first-order evidence (E) alone supports. Conversely, calibrationist theories are those theories which require your credence to match your expected degree of reliability. Specifically, calibrationists recommend that

if, independently of the first order reasoning in question, your expected degree of reliability concerning whether P at time t is r, r is the credence that it is rational for you to adopt at t.⁸

This requires an agent to align their credences with their expected degree of reliability. In much of the literature on higher-order evidence, it isn’t entirely clear what one’s expected degree of reliability is. However, I think it makes the most sense if we take r to be spelled out in terms of a higher-order probabilistic belief function⁹ which assigns a probability to the hypothesis under consideration supposing that your first-order reasoning supports that hypothesis. This might be spelt out more formally as: $\text{Pr}_{\text{HOE}}(\text{H}$

⁶ Schoenfield 2016; Christensen 2010.

⁷ Schoenfield 2015, 2018.

⁸ Schoenfield 2015, p. 428.

⁹ A function is a mathematical device that takes inputs and assigns to each input a single output. A belief function is a kind of function which takes propositions about the world as inputs and assigns to each of those propositions a probability that it is true (epistemologists call these probabilities ‘credences’ or ‘degrees of belief’). A higher-order belief function is just a function that takes your degrees of belief as inputs and assigns to them a probability of being correct.

$\mid \Pr_{\text{FOE}}(H) > 0.5) = r$.¹⁰ Or more simply, the probability that H really is true given the calculations you conducted using your FOE led you to conclude that H was true. In 'Fatigue', this would be the probability that the target is at x supposing you have calculated that the target is at x: $\Pr_{\text{HOE}}(\text{target at } x \mid \Pr_{\text{FOE}}(x) > 0.5) = r$.

Calibrationist views can be further divided into what Schoenfield calls evidential calibration ('E-calibration') and judgment calibration ('J-calibration'). J-calibration requires an agent in possession of HOE to adopt a credence equal to r (the agent's expected degree of reliability) in the judgment they have made, even if the evidence does not actually support that judgment. E-calibration, however, requires an agent with HOE to take on a credence of r in the hypothesis that is supported by their FOE. And what does it mean to say that an agent makes a judgment which is supported by the evidence? Well, a judgment is supported by the evidence if it is the judgment an agent would make upon reflection when in circumstances conducive to good reasoning.¹¹

These proposals may be elucidated by considering some examples of the beliefs J-calibration and E-calibration recommend. Let's assume that having an interest in knowing such things, the military has conducted extensive research into the influence of fatigue on calculating capacity. Consequently, we know that soldiers in situations like 'Fatigue' tend to reach the correct conclusion 60% of the time. Therefore, you know that your expected reliability, r, in 'Fatigue' is 0.6. That is, $\Pr_{\text{HOE}}(\text{target at } x \mid \Pr_{\text{FOE}}(x) > 0.5) = 0.6$. Assuming that your calculation of the coordinates is correct, J-calibration and E-calibration deliver the same verdict. J-calibration recommends you adopt a credence equal to r in the judgment you made on the basis of your evidence, which would mean adopting a credence of 0.6 in H. E-calibration recommends you adopt a credence equal to r in the proposition which is supported by your FOE. In this case, that would mean adopting a credence of 0.6 in H. Conversely, the two models give different recommendations when the agent has committed a cognitive error concerning the relationship between E and H. If you incorrectly judge $\sim H$ to be supported by E, then J-calibration recommends 0.6 credence in $\sim H$, while E-calibration's recommendation is unchanged. Having canvassed the models of rational inference incorporating HOE, we can now determine whether Christensen's position is sustainable.

4. Models of Rationality

¹⁰ $\Pr_{\text{FOE}}(H)$ is the probability that your first-order belief function assigns to the proposition H. In this case, $\Pr_{\text{FOE}}(H)$ must be greater than 0.5 as any probability lower than 0.5 (50%) indicates that your first-order reasoning does not support the conclusion that H is true.

¹¹ My thanks to an anonymous reviewer for pointing out that a pedant might object that an agent may reason poorly even if in circumstances conducive to good reasoning, in which case the judgment would not necessarily be supported by the evidence. In response to such a pedant I would argue that to be in circumstances conducive to good reasoning an agent must be free of cognitive defects. Therefore, the fact that an agent has made a cognitive error indicates they are not in circumstances conducive to good reasoning.

When determining whether Christensen's position can be sustained, the key question is: how exactly should first-order and higher-order evidence interact? If we cannot give a good account of this interaction, then Christensen's position looks less tenable. When answering this question, we should keep in mind that there are different perspectives from which we can assess rationality. In particular, we can employ internalist or externalist standards. I take an internalist standard of rationality to be one which requires consistency only between an agent's doxastic states and evidence to which an agent has access, where doxastic states and evidence will be sufficiently accessible to count as internal if they can be accessed through reflection alone. Conversely, an externalist standard of rationality requires consistency between an agent's doxastic states and evidence to which that agent does not necessarily have access. As the externalist versus internalist debate in epistemology is a long and storied one I don't intend to mediate between the positions.¹² Instead I utilise this distinction as a means of clarifying the debate in the higher-order evidence literature. Making these perspectives explicit demonstrates why the models of rational inference incorporating HOE differ as well as the benefits and shortfalls of each of the proposed solutions. As a result, a discussion of the internal and external requirements of rationality helps clarify how to incorporate HOE into our model of rational inference.

4.1 An Internalist Standard

Christensen's model of HOE gives clear guidance regarding what behaviour is rational from the perspective of an agent possessing HOE. We said that possessing HOE means you can't give your evidence the full consideration that rationality requires. Once you become aware of the fact that you are intoxicated, fatigued, etc., you can no longer update your beliefs on the basis of E even though it seems rational to do so. That is, HOE denies an agent the ability to determine through reflection alone the true relationship between their evidence and the hypotheses under consideration. Consequently, the agent no longer has sufficient access (by internalist standards) to the true relationship between E and H. Conversely, the relationship an agent *judges* to hold between E and H and the fact that the agent possess HOE remains accessible. Consequently, calibrating (which type I'll consider shortly) is the behaviour which is rational from the perspective of an agent with HOE. As put by Christensen:

After all, if I could give all my evidence its due, it would be rational for me to be extremely confident of my answer, even knowing that I'd been drugged...Yet it seems intuitively that it would be highly irrational for me to be confident in this case..., I must (at least to some extent) bracket the reasons this evidence provides, if I am to react reasonably to the evidence that I've been drugged.¹³

So, it is rather puzzling to say that in the face of higher-order evidence one should maintain the same credence in a proposition about which your reasoning is called into

¹² For readers who would like to know more about this debate both Pappas (2017) and Poston (2008) provide accessible introductions.

¹³ Christensen 2010, p. 195.

doubt. Consequently, a denial of steadfastness seems a modest claim. However, steadfastness has some intuitive appeal of its own: why should my being drugged, delusional, fatigued etc., have any bearing on the likelihood that some unrelated proposition is true? And it isn't clear which one of these intuitions should prevail. Consequently, both sides need to justify their models.

So if the internalist wishes to advocate for calibrationism they must provide a plausible account of this model of rational inference. From an agent's perspective, E-calibration appears to be a non-starter for at least two reasons. First, E-calibration commits agents to Moore-paradoxical epistemic states.¹⁴ That is, E-calibration requires agents to contradict their own judgments so that they must think to themselves 'I judged that H but I believe that \sim H.' Many authors have considered it unlikely that such a state could ever be rational.¹⁵ And secondly, from an agent's perspective, E-calibration can't be a guide to their behaviour. The presence of HOE means that an agent can no longer be certain which proposition is supported by the FOE; remember that HOE causes us to question whether we have correctly assessed the relationship between our FOE and hypothesis under consideration. That is, in the presence of HOE you should reconsider whether you have used your FOE appropriately; maybe the hypothesized location of the enemy base is wrong. But then again, it is also possible that you calculated the location correctly. In the presence of HOE you can't be sure what conclusion is actually licensed by your FOE. And since E-calibration requires you to adopt a credence which is supported by your FOE (not merely what appears to be supported by the FOE) agents with HOE have no way of knowing whether they have complied with E-calibration. Therefore, any agent with HOE will be epistemically stranded, with no guidance on to how to proceed. Consequently, E-calibration isn't an acceptable standard of internalist rationality.

On the other hand, J-calibration is a reasonable standard from the perspective of an agent with HOE. J-calibrating requires agents to adopt beliefs which are consistent with their judgments. Judgments which, unlike the relation of evidential support, are accessible. Furthermore, J-calibration also avoids the Moore-paradoxical epistemic states which make E-calibration seem untenable. Finally, by assigning credence r to the proposition an agent takes their FOE to support J-calibration accommodates the requirement that we bracket our first-order reasoning. Thought of in this way, J-calibration is an acceptable model of internalist rationality.

However, as put by Schoenfield J-calibration "makes rationality too cheap."¹⁶ This cheapness arises in cases where an agent's judgment regarding the relationship between E and H is incorrect. In the case that E supports H but an agent incorrectly judges that E supports \sim H, J-calibration recommends one adopt a credence r in \sim H. But this requires approval of credences which our evidence does not support. Consequently, if rationality only requires coherence between internal criteria, it appears to allow agents to adopt absurd yet rational beliefs. Agents will be deemed rational no matter how clearly their beliefs contradict their first-order evidence,

¹⁴ For a discussion of these issues see Sliwa and Horowitz 2015.

¹⁵ See for example: Feldman 2005, Silins 2012, and Titelbaum 2015.

¹⁶ Schoenfield 2015, p. 431.

provided their beliefs are based on genuinely held judgments. Surely this is too permissive: we don't want any old judgment to count as sufficient for rational belief.

One could argue that internal coherence constraints are insufficient for rationality regardless of any appeal to an agent's HOE. I think this is correct, but I also think that J-calibration exacerbates the permissiveness of internalism. J-calibration designates agents as 'rational' simply because they recognise that they are in an impaired state.¹⁷ Merely recognising that one is impaired, however, doesn't negate epistemic errors in judging the relationship between E and H.

So even where an agent is hindered by HOE, we feel that their judgments shouldn't be divorced from the degree of evidential support provided by their FOE. But what the FOE supports isn't accessible to an agent with HOE. So, intuitively, we feel that rationality should have some externalist requirements. Therefore, I'd now like to turn to a consideration of externalist evaluations of rationality.

4.2 An Externalist Standard

Evaluating rationality from a purely external third-person point of view is problematic for its inability to take HOE into account. Presumably, an external third-person perspective is unaffected by HOE given that the effects of such evidence are agent-relative. Consequently, on an externalist view, it is correct to adopt a credence in the hypothesis which is supported by the evidence unadulterated by higher-order considerations. So if an agent is affected by HOE and wrongly concludes that E supports $\sim H$, the agent is irrational and should change their credence to favour H.¹⁸ Where an agent correctly concludes that E supports H, the agent should steadfast. Therefore, all that counts for a determination of rationality is if our agent is correct in their evaluation of E and H. In this way, an external model of rationality fails to deal with HOE at all. I echo the opinion of Feldman that such a model is deficient.¹⁹ The fact that fatigued agents often make mistakes is a genuine source of doubt concerning whether the relationship between E and H had been perceived correctly. That is, an agent who moderated their credence in the hypotheses under consideration as the result of possessing HOE would be regarded as more rational than one who did not. Consequently, a theory of rational inference should be developed in a way that acknowledges the existence of HOE. Therefore, a model that can't incorporate these doubts is lacking.

Furthermore, as with E-calibration, such a model leaves agents epistemically stranded. When an agent has higher-order evidence, they can't be certain that they have perceived the relationship between E and H correctly. We said that if an agent wrongly concludes that E supports $\sim H$, the agent should change their credence to favour H. Or, in the event that an agent is correct about the relationship of evidential support, they should steadfast. However, In the presence of HOE an agent won't

¹⁷ Again, my thanks to Katie Steele for raising this point.

¹⁸ Note that this is neither steadfasting nor calibrating.

¹⁹ Feldman 2005, p. 107.

know which situation they find themselves in. Consequently, they won't be able to determine which course of action to take. Even worse, agents will often find themselves epistemically stranded. Christensen begins his paper by acknowledging that: "spending a few hours reading a book like Kahneman reveals a depressing plethora of unexpected and unobvious factors that subtly but significantly disrupt reliable cognition."²⁰ So, if we adopt a purely externalist evaluation of rationality, we should expect to find agents upstream without the proverbial epistemic paddle most of the time.

4.3 *A Hybrid Standard*

It is important to note that internalist and externalist models are simply different viewpoints from which we can evaluate rationality. However, this does not necessitate that these viewpoints are in competition. Consequently, it is possible to propose a standard which incorporates both internal and external constraints on rationality. This is particularly important in the case of HOE as a hybrid model can overcome the difficulties we have identified so far. We said that internalism rendered rationality too cheap and that externalism's failure to deal with HOE was a deficiency. Interestingly, the advantages of each model appear to complement the disadvantages of the other—internalism gives weight to one's HOE and externalism isn't cheap as it requires agents to accurately assess the relationship between E and H.

As discussed, our intuitions point in favour of taking the presence of HOE into account. Consequently, we should calibrate in one form or another. We also noted that E-calibration is a non-starter given it requires Moore-paradoxical epistemic states and leaves agents epistemically stranded. J-calibration then seems to be our only option. But then again, we said that J-calibration cheapened rationality. However, this objection isn't fatal and Schoenfield recognises that it is possible to respond by arguing that this only shows that J-calibration is necessary but not sufficient for rationality.²¹ For why, says the J-calibrationist, should we expect J-calibration to remedy an agent's initial irrational judgment? J-calibration only requires coherency between our judgments and our belief states. Consequently, it only speaks to that coherence and is silent on whether one's initial judgment involved rational errors.

While acknowledging that this reply is open to the J-calibrationist, Schoenfield (2015: 438) argues that it also fails to preserve a robust version of rationality. Even if J-calibration is only a necessary requirement, it still obliges agents to believe hypotheses not supported by their FOE. If J-calibration is necessary for rationality, adopting a credence which doesn't conform to your expected degree of reliability is irrational. So, to be rational an agent must adopt a credence equal to r in H. Consequently, if a rational credence does exist it must assign r to H even if H isn't supported by one's FOE. So, argues Schoenfield, the retreat to necessity rather than sufficiency doesn't avoid our original concerns with J-calibration.

²⁰ Christensen 2016, p. 397; Kahneman 2011.

²¹ Schoenfield 2015, p. 437.

Nevertheless, all hope is not lost. If the J-calibrationist wishes to accommodate FOE they can argue that where an agent has incorrectly judged the relationship between E and H there is no rational belief state that can be adopted. J-calibration is only concerned with coherence between our judgments and our beliefs, it does not require that the initial judgments are rational. If, then, the judgments made by an agent are irrational, it shouldn't come as a surprise that complying with J-calibration fails to make the agent rational. Rationality on this model would require coherence between the agent's judgments and their beliefs (J-calibration) plus rational judgments to begin with!²²

For these reasons, J-calibration won't secure rationality by itself. However, it is still a necessary principle; surely an agent isn't rational if they don't abide by their own judgments. We concluded that in order to be rational we require compliance with J-calibration plus judgments which are initially rational. Consequently, our next question arises: what would make our initial judgments rational? The answer is indicated by our preceding analysis. We said that where an agent has misjudged the relationship between E and H there is no rational belief state they can adopt. Therefore, in order to be rational, an agent must judge the relationship correctly. More precisely, an agent must make judgments which are supported by their first-order evidence. So the two conditions required for rationality are:

J-FOE – One should make judgments that are supported by one's first order evidence.

J-calibration – One should assign a credence r (equal to one's expected degree of reliability) to the hypothesis H they judge to be correct.

J-FOE then serves as our external requirement – an agent must judge in accordance with the evidence, even where HOE denies the access needed to determine what the evidence supports. While J-calibration requires coherence between internally-accessible doxastic states – the agent's judgments and beliefs. These conditions taken together I call J*-calibration. If J*-calibration is correct it places tight constraints on what will count as a rational belief. In order to be rational, an agent is required to both judge correctly and form beliefs which are consistent with those judgments. Therefore, rather than cheapening rationality, J*-calibration makes rationality more difficult to satisfy.

One could argue that J-FOE is an odd requirement given that HOE prevents an agent from assessing whether they have complied with J-FOE. However, this requirement isn't odd if we remember that we are allowing a partially external evaluation of rationality. Even though it may not be possible to know whether one has complied with J-FOE, all things considered, an agent who judges in accordance with their FOE is more rational than one who does not. So, from an external perspective, compliance with FOE is a requirement of rationality. As a result, an agent's inability to determine whether they have satisfied J-FOE is consistent with J-FOE being a necessary requirement of rationality.

²² Christensen 2016, p. 409.

One might also worry that J^* -calibration, in particular J-FOE, leaves agents epistemically stranded in the same way that the purely external evaluation of rationality and E-calibration left agents stranded. I disagree. Unlike E-calibration or purely external rationality, J^* -calibration still provides some guidance to agents with HOE. Specifically, an agent with HOE should assign r to the hypothesis they judged to be best supported by their evidence, whereas an agent in the other scenarios is unable to determine what to do.²³

Another worry regarding the ability to follow this model concerns cases in which an agent receives HOE before forming a credence in H . In such a case the required inputs for J^* -calibration do not appear to be available. According to Christensen such cases make it impossible to form a credence in H independently of your HOE.²⁴ However, I argue that an independent assessment is still possible. To see why recall that HOE isn't evidence about the actual relationship between E and H . Rather, it is evidence that causes you to re-evaluate whether you have perceived the relationship between E and H correctly. Consequently, the presence of HOE does not make it seem as though the FOE supports some other hypothesis. Instead your HOE motivates you to re-examine the hypothesis which seems to be supported by your FOE. That is, the way the FOE appears to you is unchanged by the presence of HOE. For example, if you looked at the sky and concluded it was orange and I then told you that I had secretly drugged you with a chemical that makes blue things look orange, the sky would still appear orange to you. Accordingly, the fact that your FOE appears to support a particular hypothesis is not affected by the presence of HOE. And this is true whether you receive the HOE before or after making an appraisal of your FOE. Therefore, it is still the case that you can make an appraisal of the FOE independently of your HOE. Consequently, in cases where you receive HOE before appraising your FOE an application of J^* -calibration is straightforward – determine which hypothesis your FOE appears to support, then bracket that hypothesis as required by your HOE.

So J^* -calibration doesn't leave agents epistemically stranded, but it is still impossible to consciously comply with when one has HOE. An agent might J-calibrate, but still fail to satisfy J-FOE and therefore be irrational. Furthermore, because of the presence of HOE, the agent would be unaware that they had behaved irrationally. So, in this sense, J^* -calibration is not a useful guide to rectifying rational errors in our judgement of the relationship between our FOE and H . In response to this objection I am in agreement with Christensen who, when defending a similar objection against his own model of rational inference incorporating HOE, had the following to say:

...acknowledging that we can't always get evidential relations right or be certain that we have gotten them right is compatible with holding that failure to get evidential relations right is a rational failure. In epistemology, as in ethics,

²³ Of course, from an agent's point of view J^* -calibration prescribes the same course of action as J-calibration; that is, an agent in possession of HOE should adopt a credence in the hypothesis they judged to be best supported by their evidence which is equal to their expected degree of reliability. Consequently, from an agent's perspective J^* -calibration doesn't provide different guidance to J-calibration. Nevertheless, J^* -calibration is an improvement on J-calibration as it acknowledges that we should adopt a credence equal to our expected degree of reliability *and* assess our FOE correctly.

²⁴ Christensen 2016, p. 408.

the correct norms may put perfection beyond what ordinary humans can generally achieve.²⁵

What I have provided is an ideal theory of rational inference incorporating HOE. It is unreasonable to expect that such a theory also ensures we no longer make rational errors when assessing our FOE. By way of analogy, consider a diet that contains all necessary nutrients and has been shown to substantially improve one's health. It is lamentable that some people may find it difficult to stick to this diet. However, it is no objection to the desirability of such a diet that some agents prefer junk food to healthy meals. Rather than showing the diet is theoretically deficient this highlights a separate problem concerning an agent's willpower or preferences. Similarly, ensuring that agents correctly utilise their FOE is a separate problem to developing a theoretically sound model incorporating HOE. For this reason, it is no objection to J*-calibration that agents cannot always comply with this standard. Of course, it is also desirable to have some means of ensuring we evaluate our FOE correctly, however, this is a task for another theory.

5. Conclusion

The aim of this paper was to defend a model of rational inference which could accommodate HOE. The preceding analysis has done just this. This analysis helps clarify why HOE presents a problem for our models of rational inference and why agents often fall short of ideal rationality. Once these issues are made clearer a defensible model that explains how HOE and FOE interreact becomes apparent. In particular, J*-calibration allows us to vindicate both internal and external constraints on rationality – we can accommodate our desire to judge in accordance with our first-order evidence while taking the presence of higher-order evidence seriously. Of course, the theory is not a panacea for human irrationality as it cannot rectify human fallibility in the evaluation of first-order evidence, but this is not something such a theory is required to do.

References

- Christensen, David. (2010) 'Higher-Order Evidence.' *Philosophy and Phenomenological Research* 81: 185-215.
- Christensen, David. (2016) 'Disagreement, Drugs, etc.: from Accuracy to Akrasia.' *Episteme* 13 (4): 397-422.
- Feldman, Richard. (2005) 'Respecting the Evidence.' *Philosophical Perspectives* 19 (1): 95–119.

²⁵ Christensen 2016, pp. 411-2. Christensen defends what he calls the Idealised Thermometer Model which is related to, but different from, my own model of rational inference incorporating HOE.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Lasonen-Aarnio, Maria. (2014) 'Higher-Order Evidence and the Limits of Defeat.' *Philosophy and Phenomenological Research* 88, no. 2: 314-345.

Leydon-Hardy, Lauren. (2016) 'Some Resistance to the Idealized Thermometer Model.' *Episteme* 13 (4): 423-426.

Pappas, George. (2017) 'Internalist vs. Externalist Conceptions of Epistemic Justification.' <<https://plato.stanford.edu/archives/fall2017/entries/justep-intext/>> accessed 29 August 2019.

Poston, Ted. (2008) 'Internalism and Externalism in Epistemology.' Internet Encyclopedia of Philosophy. <<https://www.iep.utm.edu/int-ext/>> accessed 29 August 2019.

Schechter, Joshua. (2013) 'Rational Self-Doubt and the Failure of Closure' *Philosophical Studies* 163 (2): 428-452.

Schoenfield, Miriam. (2015) 'A Dilemma for Calibrationism.' *Philosophy and Phenomenological Research* 91 (2): 425-455.

Schoenfield, Miriam. (2018) 'An Accuracy Based Approach to Higher Order Evidence.' *Philosophy and Phenomenological Research* 96 (3): 690-715.

Silins, Nicholas. (2012) 'Judgment as a Guide to Belief.' In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar. Oxford University Press.

Sliwa, Paulina, and Sophie Horowitz. (2015) 'Respecting All the Evidence.' *Philosophical Studies* 172 (11): 2835-2858.

Titelbaum, Michael. (2015) 'Rationality's Fixed Point (or: In Defense of Right Reason).' *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne. Oxford University Press.