

University of Groningen

Higher-order evidence and losing one's conviction

Henderson, Leah

Published in:
Noûs

DOI:
[10.1111/nous.12367](https://doi.org/10.1111/nous.12367)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Henderson, L. (2021). Higher-order evidence and losing one's conviction. *Noûs*.
<https://doi.org/10.1111/nous.12367>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Higher-order evidence and losing one's conviction

Leah Henderson 

Department of Theoretical Philosophy,
University of Groningen, Groningen, The
Netherlands

Correspondence

Leah Henderson, Department of Theoretical Philosophy, University of Groningen, Groningen, The Netherlands.
Email: l.henderson@rug.nl

Funding information

Netherlands Organisation for Scientific Research (NWO), Grant/Award Number: 275-20-058

Abstract

There has been considerable puzzlement over how to respond to higher-order evidence. The existing dilemmas can be defused by adopting a 'two-dimensional' representation of doxastic attitudes which incorporates not only substantive uncertainty about which first-order state of affairs obtains but also the degree of conviction with which we hold the attitude. This makes it possible that in cases of higher-order evidence the evidence sometimes impacts primarily on our conviction, rather than our substantive uncertainty. I argue that such a two-dimensional representation is naturally developed by making use of imprecise probabilities.

1 | INTRODUCTION

So-called 'higher-order evidence' presents a puzzle in epistemology. It undermines our confidence in our own reliability or rationality, or our confidence that our first-order evidence really supports the doxastic attitude that we have formed. Can it then be rational for us to maintain that first-order doxastic attitude? If we do, there will be tension between our attitudes at the first-order and at the higher-order level – so-called 'Level-Splitting'. Is this acceptable? Or are we rationally required to revise our first-order attitude in order to eliminate the tension between levels? There appear to be problems either way. Level-Splitting brings the threat of epistemic akrasia and it has implausible consequences for decision-making. Yet, existing Revision views fail to account for the fact that in certain higher-order evidence cases, the evidence does not seem to tell us anything substantial about first-order matters. Its effect seems to be rather to make us less confident about our own judgment.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2021 The Authors. *Noûs* published by Wiley Periodicals LLC.

In this paper, I will argue that this dilemma can be resolved if we represent doxastic attitudes in a way which captures two distinct components or dimensions to our uncertainty. On the one hand, we have uncertainty about which particular first-order state of affairs obtains – I will call this our ‘substantive’ first-order uncertainty. On the other hand, we may have some uncertainty about our substantive judgments themselves. I will refer to this as the amount of ‘conviction’ we have in those judgments. The overall representation of our doxastic attitude then should incorporate both of these dimensions. This opens the possibility that higher-order evidence may in some circumstances primarily impact on the conviction dimension of the doxastic attitude, rather than on the substantive dimension. An advantage of this approach is that it allows us to accommodate the intuitions that motivate Level-Splitting views, while avoiding the problems that are associated with them.

I will argue that a suitable framework for developing such a two-dimensional view is that of imprecise probability. Roughly the idea is the following. In this framework, a doxastic attitude can be represented by a set of probability measures, rather than one particular probability measure. Such a credal set has a degree of precision which is greater when there is a smaller difference between the highest and the lowest probabilities in the set. The degree of precision can serve as a representation of the degree of conviction. The proposal then is that cases of higher-order evidence may induce a loss of precision of the agent’s credal state. Loss of precision of the credal state can be shown, using existing decision theories for imprecise probabilities, to produce the behavioural implications we tend to associate with learning higher-order evidence.

2 | THE HIGHER-ORDER EVIDENCE DEBATE

The concept of higher-order evidence is usually informally characterised in the literature, and there are various formulations. For example, Christensen describes it as ‘evidence of my own rational failure’ (Christensen, 2010a), as ‘evidence concerning the reliability of our own thinking about some particular matter’ (Christensen, 2016). It has also been described as ‘evidence about what your evidence supports’ (Sliwa & Horowitz, 2015) and evidence that ‘induces doubts that one’s doxastic state is the result of a flawed process’ (Lasonen-Aarnio, 2014). The phenomenon of higher-order evidence and the puzzles it presents are usually introduced by giving one or several paradigmatic examples. One of the classic examples is:

Hypoxia: I have just achieved a difficult first ascent in the Himalayas. As the weather turns, I have to abseil down a long pitch. I have gone through a sequence of reasoning several times to check that I have constructed my anchor correctly, that I haven’t under-estimated the length of the pitch, and that I have threaded the rope correctly through my belay device and carabiner. I then acquire evidence that I am in serious danger of being affected by a mild case of hypoxia caused by high altitude. Such hypoxia impairs one’s reasoning while making it seem perfectly fine. I know that mountaineers have made stupid but fatal mistakes in the past as a result of being in such a condition. (Lasonen-Aarnio (2014), p. 315)

Another example is:

Fatigued doctor: I’m a medical resident who diagnoses patients and prescribes appropriate treatment. After diagnosing a particular patient’s condition and prescribing

certain medications, I'm informed by a nurse that I've been awake for 36 hours. I know that people are prone to make cognitive errors when sleep-deprived (and perhaps I even know my own poor diagnostic track record under such circumstances). (adapted from an example in Christensen (2010a)).

The question that is posed is: what should happen to my doxastic state on receiving higher-order evidence in cases like these?

Before getting to the proposed answers to this question, it will be useful first to provide some definitions. I will characterise the distinction between first-order and higher-order evidence, by making use of a distinction between first-order and higher-order propositions.¹ First-order propositions concern ordinary subject matter in the world. Examples are the proposition that the rope is correctly tied, or the proposition that the appropriate medicine for the patient is X. Higher-order propositions, on the other hand, are propositions that concern an epistemic agent, her doxastic states or doxastic attitudes. They may concern the relation between evidence and first-order doxastic attitudes. For example, a higher-order proposition is that the medical resident's belief is well-supported by her evidence, or that I have processed certain evidence in a rational way in forming my doxastic attitude towards a first-order proposition. Thus, first-order and higher-order propositions are distinguished by their subject-matter.

An agent can be in a certain doxastic state, which consists of taking a certain doxastic attitude towards a proposition. If the proposition is a first-order one, we say they are in a 'first-order doxastic state', and have a 'first-order doxastic attitude'. If the proposition is a higher-order one, we say they are in a 'higher-order doxastic state', and have a 'higher-order doxastic attitude'. How exactly should we represent doxastic attitudes? There are several different frameworks which are commonly used: the framework of categorical or full belief, and the Bayesian framework of partial belief or credence. In the framework of full belief, there are usually taken to be three doxastic attitudes that one might take towards a given proposition: believe the proposition, disbelieve the proposition, or suspend judgment about it. In a Bayesian framework, the agent's doxastic attitude is a credence represented by a probability for the proposition.

Now consider the notion of evidential support. I will say that evidence supports which doxastic attitude towards a proposition you have reason to take.² We can have evidence which bears on the attitude that we should take towards a first-order proposition. We will call this 'first-order evidence'. Evidence which bears on the attitude we should take towards a higher-order proposition we will call 'higher-order evidence'. This way of talking means that the same piece of evidence may serve as both first-order and higher-order evidence in cases where it bears on attitudes we should take towards both kinds of propositions.

The higher-order evidence debate concerns the question of what kind of revision of the overall doxastic attitude of an agent is called for in examples like those given above. The overall doxastic attitude consists of first-order doxastic attitudes and higher-order doxastic attitudes. In examples like Hypoxia and Fatigued Doctor, it seems clear that higher-order doxastic attitudes need to be revised. For example, in the Hypoxia case, suppose I have a belief in the higher-order proposition that I am reasoning well. The evidence that I may have hypoxia has bearing on this proposition, and should cause some revision in my attitude towards it, such as losing confidence in it. Similarly, in the doctor case, I have some initial doxastic attitude towards the proposition that my first-order evidence from examining the patient supports my belief that X is the correct medicine. The evidence that I am underslept should make me lose some confidence in that proposition.

An important axis of the debate concerns whether, as well as revising one's attitude towards higher-order propositions, one is also rationally required to revise one's attitudes towards the

first-order propositions. Given our set-up, this question can be posed as: Is the evidence in these kinds of cases also first-order evidence? In Hypoxia, am I required to revise my first-order belief that the anchor is correctly constructed? In Fatigued Doctor, am I required to revise my belief that X is the appropriate medicine?

The view that such revision to first-order attitudes is not required has been called ‘Level-Splitting’, because it allows for a certain kind of tension between first-order and higher-order attitudes (Horowitz, 2014). For example, there is some tension between believing the proposition q that X is the correct medicine and being doubtful about whether my evidence supports belief in q . Views which allow that Level-Splitting may be rational in least some cases have been put forward by a number of authors (Coates, 2012; Hazlett, 2012; Lasonen-Aarnio, 2014, 2015; Worsnip, 2018; Wedgwood, 2012; Williamson, 2011; Weatherson, ms).

Others argue that this is not permissible, and that I must also revise my first-order attitude in order to remove the tension with my revised higher-order attitudes. I will call this alternative view ‘Revise’. Concrete answers to how this should be done have been proposed in two distinct frameworks: the framework of full belief, and the framework of credences.

In the full-belief framework, the role of higher-order evidence has been conceptualised as a kind of undercutting defeater. Suppose one has a first-order belief about a proposition p on the basis of evidence E . Two kinds of defeaters are usually distinguished (Pollock, 1986). A rebutting defeater for a belief in p consists of evidence for not- p . An undercutting defeater, on the other hand, consists of evidence that E does not support p . A number of authors have suggested that higher-order evidence operates as an undercutting defeater for existing first-order beliefs, since it casts doubt on the connection between the first-order evidence and the belief concerning p (eg. Feldman, 2007, 2009; Christensen, 2010a; Lasonen-Aarnio, 2014). When the original belief is undercut, the usual suggestion is that the agent abandon the belief and suspend judgment instead.

In the framework of partial belief, the main suggestion has been that the original first-order credence should be revised according to some kind of ‘calibration’ principle for probabilities, which aligns the agent’s credences with the information about reliability supplied by the higher-order evidence (White, 2009; Sliwa & Horowitz, 2015; Schoenfield, 2015, 2018; Roush, 2009, 2017). When this is recommended, examples are often formulated to include more specific frequency information. For example, Sliwa and Horowitz discuss a version of the doctor example formulated as follows:

Drugs: Anton is an anesthesiologist, trying to determine which dosage of pain medication is best for his patient: A or B. To figure this out, Anton assesses some fairly complex medical evidence. When evaluated correctly, this kind of evidence determines which dose is right for the patient. After thinking hard about the evidence, Anton becomes highly confident that dose B is right. In fact, Anton has reasoned correctly; his evidence strongly supports that B is the correct dose. Then Sam, the chef at the hospital’s cafeteria, rushes in. “Don’t administer that drug just yet”, he says guiltily. “You’re not in a position to properly assess that medical evidence. I slipped some reason-distorting mushrooms into your frittata as a prank. These mushrooms make you much less reliable at determining which dose the evidence supports: in the circumstances you presently face – evaluating this type of medical evidence, under the influence of my mushrooms – doctors like you only tend to prescribe the right dose 60% of the time!”. In fact, Sam is mistaken: the mushrooms he used were just regular dried porcini, and Anton’s reasoning is not impaired in the least. But neither he nor Anton knows (nor has reason to suspect) this (Sliwa and Horowitz (2015), p. 2836).

The key idea of those advocating recalibration is that Anton should replace his initial credence with his expected reliability, which is in this case taken to be 60%. Thus Anton should reduce his initially high credence that B is the correct dose to 60% (Sliwa & Horowitz, 2015).

2.1 | Motivations for level-splitting

It has been argued that there are certain classes of higher-order evidence cases where Level-Splitting is an appropriate response. This has, for example, been suggested in cases of ‘misleading higher-order evidence’. Allen Coates presents the following case:

Sherlock Holmes. Sherlock Holmes judges how well Watson reasons about a case. Generally his assessment of Watson’s performance is very reliable. Suppose Watson assesses the evidence in a particular case and comes to the conclusion that the butler did it. However Holmes tells Watson that he reasoned in an irrational fashion. On this occasion Holmes happens to be wrong (even though in general, he isn’t), and Watson did reason correctly Coates (2012).

According to Coates, in this case Watson may reasonably (though wrongly) conclude that his conclusion is irrational, but nonetheless it can also be rational for Watson to maintain his first order belief that the butler did it. The idea here, and in similar examples, is that the first-order and higher-order ‘levels’ can be addressed separately. On the one hand, there is evidence which bears on the first-order proposition that the butler committed the crime. This consists of the evidence which Watson has inspected. On the other hand, there is also evidence that, when considered alone, only seems to bear on higher-order attitudes. In this case, the evidence presented by Holmes’ testimony clearly bears on Watson’s attitude towards the higher-order proposition that he has reasoned rationally. Yet, one might think, this testimonial evidence does not have any direct bearing on whether the butler did it or not. It doesn’t tell us that the butler is any more or less likely to be guilty. It is only evidence about the rationality or otherwise of Watson’s own thought processes.³

One way people have attempted to bring out the point that in cases like these, the evidence does not, without further assumptions, appear to tell us anything about the first-order proposition itself is to point out that in such cases, the agent could actually be right, and could still in fact have functioned rationally. This would happen when the evidence which serves as higher order evidence to the contrary is actually misleading. It supports a doubtful attitude towards my own rationality, but in actual fact, I have reasoned in a rational fashion. This line of thought can be bolstered by appeal to externalist views about justification. Ralph Wedgwood, for instance, argues that the existence of an evidential support relation is sufficient for justification, whether or not the agent has correct beliefs about whether they have identified it (Wedgwood, 2012).

In some higher-order evidence cases, it seems clear that the evidence does provide reason to revise one’s substantive judgment in a particular way. This could happen, for example, if it consists of psychological evidence of a tendency to over-estimate credences or to be over-confident. In such a case, this can provide good reason for the agent to revise their credence downwards. But in other cases, the higher-order evidence may only indicate some general failures, without providing any information about the specific way one is likely to have gone wrong.⁴ These have been called ‘non-lopsided cases’, in contrast to the ‘lopsided cases’ where the evidence serving as higher-order evidence also indicates the direction in which one’s first-order attitude has erred

(Steglich-Peterson, 2019). One of the motivations for Level-Splitting has been that in non-lopsided cases, it seems strange to insist on changing the first-order attitude when evidence is presented which offers no reason for making any particular change.

2.2 | Motivations for Revise views

Although Level-Splitting appears to have some intuitive appeal in the kinds of cases just described, it does struggle with two other considerations, which together form an important part of the motivation for adopting a Revise view instead. The first difficulty is that Level-Splitting may involve epistemic akrasia. In standard practical cases of akrasia, one thinks one ought not to do a certain action a , but one does it anyway. You think you shouldn't smoke, for instance, but you smoke in any case. Epistemic akrasia then is understood as a situation where you think you ought not have a certain doxastic attitude towards a proposition p , but you have that attitude towards p anyway. For example, you might think that you should not believe p , perhaps because you think that your evidence does not support believing p , yet you believe p anyway.⁵ If one takes the view that an agent in a higher-order evidence cases is permitted to enter into a Level-Splitting state, then this may mean that their overall epistemic state is akratic. For example, Watson in the Sherlock Holmes case ends up in a state of epistemic akrasia because after hearing Holmes' testimony, he might reasonably think that his evidence does not support believing that the butler did it, yet he believes that nonetheless. Although there have been a number of attempts by Level-Splitters to argue that epistemic akrasia is rationally permissible in certain cases (Coates, 2012; Christensen, 2016; Lasonen-Aarnio, 2014), there are also a number of arguments that epistemic akrasia should be regarded as irrational (Greco, 2014; Smithies, 2012; Titelbaum, 2015). Revise views typically avoid any commitment to epistemic akrasia because they recommend that the first-order attitude be revised in a way which eliminates the tension between first-order and higher-order attitudes that arises with Level-Splitting. Consider for example, the Revise view in the full-belief framework. In the Sherlock Holmes case, a Reviser recommends treating Holmes' testimony an undercutting defeater for Watson's belief that p , the butler did it. Then Watson does not end up in a state where he thinks that he shouldn't believe p , and believes p anyway. Rather he suspends judgment as to p .

Another major difficulty for the Level-Splitting view is its implications for action (Horowitz, 2014). According to standard views in decision theory, our doxastic attitude towards first-order propositions guides our decision-making. If those attitudes are expressed as credences, one of the main principles that is employed is to maximise expected utility, given the credences. In the Fatigued Doctor case, suppose my credence that X is the right medicine for the patient remains high, despite my lack of confidence that my evidence supports such a high level of credence. Then, according to the usual view of decision-making, I should be just as inclined to act on my opinion as I was before, because the part of my doxastic state which is relevant for decisions has not changed. Yet, we might think this evidence *should* make some difference to what the doctor is inclined to do. In particular it should make the doctor hesitate, and to avoid making immediate decisions if they can be postponed. I might be well-advised to call off duty and to postpone further prescriptions until I have had enough sleep. Of course, the situation may not allow for postponing action. I might not be able to take myself off-duty. In that case, it seems perhaps the doctor should be less inclined to take risks – avoiding prescribing any medicine for which there are significant negative consequences in the case that she is wrong. In the case of the hypoxic hiker, there is also no luxury of putting off important decision-making. In such situations where choice is forced, it seems reasonable that the effect of the higher-order evidence should be to induce a tendency

TABLE 1 Decision table for Anton. The utilities for his three options A, B and C are shown. Possible states of the world are s_1 and s_2

	s_1	s_2
A	20	-80
B	-13	20
C	0	0

to err on the side of caution, by opting for the less risky option. The problem for Level-Splitting is that none of these apparently reasonable ways of responding to the higher-order evidence is recommended, since the basic action-guiding doxastic attitude has not been revised.

On the other hand, Revise views will typically produce the needed change in the action-guiding first-order attitude. Consider the Sherlock Holmes case. If Watson revises to a state of suspension of judgment, then this will be more conducive to hesitation and postponing of decision-making than retaining full belief in p . Similarly, in a partial belief framework, the Reviser recommends that the credence be revised, for example by recalibration. This also avoids the consequence that the higher-order evidence has no implications for decision-making. To give a concrete example, suppose, for example, that Anton's decision table is as represented in Table 1.

Anton faces a choice between administering dose A and administering dose B to his patient. Suppose that Anton also has a third option, C, which is not to administer any dose. Let s_1 be the state of the world where the patient has disease X, and s_2 be the state where she has disease Y. There is positive utility for administering dose A if she has X, and for administering B if she has Y. Giving her B, when she has X is quite undesirable, and giving her A, when she has Y, is even more so. Not giving any medicine, option C, is taken to preserve the status quo. Suppose Anton has initially a high credence, 0.9, in s_1 . After hearing Sam's warnings, he recalibrates to a credence of 0.6. According to expected utility theory, if Anton has any precise probability over 0.8 for s_1 , he should choose to give A. If he has probability between 0.61 and 0.8, he should choose C (do nothing), and if his probability is below 0.61 he should choose B. Thus, in this case, Anton becomes so unsure about s_1 that he decides to take the less risky option B, where the consequences of choosing the wrong medicine are less detrimental. In a case where the evidence is such as to produce not such a large revision – say to a new credence in between 0.61 and 0.8, then Anton should, according to expected utility theory, choose the option C of not acting. Thus, in a case like this, the Recalibration view does imply that learning higher-order evidence of this kind will prompt either hesitation to act, or choosing a less risky option.

In summary, then, Revise views can avoid epistemic akrasia and also have more plausible consequences for decision-making than Level-Splitting views. However, they still struggle with the intuition that in cases which motivated Level-Splitting, the higher-order evidence appears to give no reasons in favour of any substantive shift in first-order opinion. I will now suggest that a way around these difficulties is to adopt a what I will call a 'two-dimensional' representation of doxastic attitudes. In fact, such representations have already been advocated on independent grounds, as we will see.

3 | TWO-DIMENSIONAL REPRESENTATIONS OF DOXASTIC ATTITUDES

It is useful to distinguish between two kinds of uncertainty which a doxastic attitude may need to represent. When the agent thinks about first-order propositions, they will typically entertain

as possible a range of mutually exclusive alternatives. The simplest case is a proposition p , which may be either true or false. But there may also be several alternatives: s_1, s_2, s_3 , etc. Part of what the doxastic attitude of the agent represents is how she regards the alternatives in relation to each other. For example, if she believes p , then she does not believe $\neg p$. Where the attitudes are credences, these must be distributed over the alternatives according to the agent's view of how likely each is. When the agent makes a judgment about which of the alternative propositions is true or more likely to be true, we will say that the agent is making a 'substantive' judgment. But we can also consider another dimension, which is the amount of conviction that we have in the substantive judgment. The amount of conviction is governed by a number of factors, including the amount of evidence on which we base our substantive judgment, how rational we think we were in forming it, and so on. I propose that higher-order evidence can impact on either or both of these two dimensions, and therefore it is helpful to have a representation of doxastic attitudes which captures both. Do the commonly used frameworks for doxastic attitudes do this?

Let us first consider the full-belief framework. Here there are three doxastic attitudes that one might take towards a given proposition: believe the proposition, disbelieve the proposition, or suspend judgment about it. Although the options are limited to only three, I would argue that these should be placed on a two-dimensional space. Changing from belief to disbelief in a proposition p constitutes a change in one's substantive judgment. But moving to a state of suspended opinion is a move along the conviction dimension, rather than the substantive dimension. Suspending judgment in a proposition is not the same as 'half-believing' the proposition. Rather, it is an alternative attitude we use to represent a loss of conviction in both or either of the attitudes representing substantive judgments. Although the full-belief framework can represent a loss of conviction then, it is limited in its ability to represent states of graded opinion between full belief and disbelief.

On the other hand, the Bayesian framework does allow for representation of attitudes like partial belief. But arguably it is a one-dimensional representation. Although we can represent an agent becoming increasingly doubtful about whether a proposition is true or not by a decreasing credence in that proposition from 1 to zero, we have no extra dimension along which to represent a loss of conviction. The analogue of a state of suspended opinion, in this framework, is identified with a particular substantive judgment – usually 0.5 credence in the truth of the proposition.

What we need is a representation of doxastic attitudes which is fully two-dimensional, and which allows for graded attitudes not only along the substantive dimension, but also along the conviction dimension. We may gain or lose conviction to a greater or lesser extent, so it will be good if the attitude of suspended opinion is not the only option. There is already a well-developed framework which will allow for such representations of doxastic attitudes, namely imprecise probability. We will turn to this in section 4. But before that, I will first sketch how, in general terms, a two-dimensional representation of doxastic attitudes can in principle help to resolve the puzzles that have arisen in the higher-order evidence debate.

First, we have seen that there appear to be some cases of higher-order evidence where it is not so plausible to think that the evidence produces a change in one's substantive judgment. This is particularly true for the non-lopsided cases. Rather what happens is that the evidence undermines our conviction in the substantive judgment. In a two-dimensional representation, this can be represented as a move along the conviction dimension, rather than the substantive dimension. This does not, of course, preclude the possibility that there are other cases where one does move along the substantive dimension, or indeed cases where there is a shift along both dimensions.

Second, we can avoid epistemic akrasia. This is because the first-order attitude is not just belief or credence, but also contains some representation of degree of conviction. Thus when one's confidence in higher-order propositions like 'my evidence supports a high credence in p ' is

undermined, this is reflected in a change to the conviction dimension of the first-order attitude. In the case of Hypoxia, for example, the evidence of hypoxia does support the idea that there is less basis for my first-order belief that the harness is correctly tied than I previously thought. The fact that I become less committed in my belief reflects the doubt that I have come to have in the higher-order proposition that my evidence supports that belief. It seems consistent to say that you can rationally hold a belief on the basis of little evidence, as long as you hold that belief with a suitably low level of conviction. Something like this idea has also been put in terms of epistemic modesty or humility. For example, Alan Hazlett has argued that there is an appropriate level of humility which you should have in your beliefs (Hazlett, 2012). What would be genuinely irrational then would be to fail to be humble where humility is called for. But you can hold a belief in a humble fashion while still having low confidence in the higher-order proposition that my evidence supports the belief, or that I formed it in a rational fashion.

Thirdly, we have seen that we should expect that the evidence in cases like Hypoxia or Fatigued Doctor should have some impact on the decision-making of the agent. This presents a difficulty for any views where the first-order doxastic attitude is not revised in any way. This problem is avoided, given a two-dimensional picture, because the first-order attitude is revised, albeit in some cases only or primarily along the conviction dimension. What is required then is that we have an account of how decision-making depends on the first-order attitude which makes it responsive to both dimensions of the doxastic attitude. As we will see, decision theories for imprecise probabilities do have precisely this feature.

I will now argue that imprecise probability does provide a natural way to develop a two-dimensional account of the representation of doxastic attitudes, and thus lends itself as a framework for considering problems raised by higher-order evidence. Notice that it has already been proposed that imprecise probability may be helpful in modeling closely related problems associated with peer disagreement (Elkin & Wheeler, 2018).

4 | IMPRECISE PROBABILITY

Imprecise probability (IP) is a generalisation of standard Bayesianism. In IP models the doxastic state of the agent is represented not by a single probability measure, but by a model which allows for the possibility of multiple probability measures. Imprecise probability is now a well-developed field, in which there are a number of different approaches (Augustin et al., 2014; Walley, 1991). For simplicity, we will focus on the case where the imprecise probability model consists of a set of probability measures. Following Levi, we will refer to this as a ‘credal set’ P (Levi, 1980). A lower and an upper probability are defined as the lowest and highest probabilities in the set respectively:

$$\underline{p}(q) = \inf\{p(q) : p \in P\}$$

$$\bar{p}(q) = \sup\{p(q) : p \in P\}$$

For a Bayesian, the upper and lower probabilities always coincide, giving a unique or ‘precise’ probability distribution. In an imprecise probability model, the upper and lower probabilities can come apart, giving a set of probabilities whose degree of imprecision about a proposition q can be measured by the difference between the upper and lower probability: $\bar{p}(q) - \underline{p}(q)$.

The important point for us is that the degree of imprecision in an IP credal set can be used to represent greater ‘uncertainty’ about the probability.⁶ In the terms I have used before, the degree

of precision of the credal set represents our ‘conviction’ about our first-order attitude. The amount of conviction that you have in your attitude can be governed by a number of factors, in particular the amount and quality of the evidence that it is based on.

There is in fact a long-standing history of recognition that there is another dimension to our doxastic state beyond the credence itself. Charles Sanders Peirce commented that:

to express the proper state of belief, not *one* number but *two* are requisite, the first depending on the inferred probability, the second on the amount of knowledge on which that probability is based (Peirce (1932), p. 421).

In Keynes (1921), Keynes emphasised the importance of the notion of ‘weight of evidence’ in decision-making. Suppose a Bayesian assigns probability to the proposition a that a certain coin will land heads when tossed. She assigns probability $p(a) = 0.5$ when she has no evidence from tosses but expects the coin to be fair. Suppose she now observes thousands or even millions of tosses giving statistical evidence of a symmetrical distribution between heads and tails. Her probability distribution will still be $p(a) = 0.5$. The first case and the more informed case arguably differ in the ‘weight of evidence’ that supports them. But this dimension is not represented in the probability distribution itself since it is the same in both cases.⁷

In an IP framework, the uninformed state of opinion about the coin may be represented by the set $\{p(a) : 0 \leq p(a) \leq 1\}$ whereas the more informed belief, based on much experience with the coin, could be represented by the singleton set $\{p(a) = 0.5\}$. The uninformed opinion has a greater degree of imprecision than the informed opinion. Thus, with this representation, we can distinguish between a state of complete ignorance and the state of having 0.5 credence in whether the proposition is true. The key point then is that an IP representation gives us a representation of a doxastic attitude towards a proposition which is in an important sense two-dimensional. Along one dimension is the degree of uncertainty that you have in a first-order proposition. Along the other dimension is the degree of conviction that you have in your first-order attitude. The degree of conviction you have in your first-order attitude is sensitive to factors like the amount or the quality of the evidence that you have.

A number of other authors have recognised the importance of another dimension of uncertainty accompanying the credence, and it has been variously called ‘ambiguity’, ‘epistemic risk’, and ‘Knightian uncertainty’ (Ellsberg, 1961; Knight, 1921; Sahlin & Persson, 1994). An empirical case has also been made for the claim that people’s decision-making behaviour is in fact sensitive to this extra dimension. Perhaps the best known example of such empirical studies concern the Ellsberg paradox, but there are a number of studies suggesting that people are inclined to demonstrate ‘ambiguity aversion’, that is, they are inclined to avoid acting in situations where they lack confidence about their probability judgements (Camerer & Weber, 1992). One attitude one might take to the empirical finding that people are sensitive to ambiguity could be to think that they are mistaken to do so. However, it seems more reasonable to think that the pattern of people’s actual preferences is quite rational. If we are not confident about our beliefs, this *should* provoke hesitation to act. It should also provide the stimulus for active attempts to improve one’s epistemic situation, either by removing impediments to one’s own rational processing powers, or, particularly in the case of disagreement, by stimulating further inquiry, deliberation and consultation. Therefore, sensitivity to confidence about one’s probability distribution is something which a normative framework for inquiry should represent. It has been suggested then that the Bayesian framework, which represents doxastic states as credences, is failing to represent a certain element of our rational decision-making. This has been a significant part of the motivation for the

development of the more general framework of imprecise probability and decision theories which reflect some of these apparently rational features of people's decision-making behaviour.

Making use of the IP framework allows us to accommodate both the intuitions which led to Level-Splitting, as well as the motivations for Revise views. Level-Splitting was motivated in particular by cases where higher-order evidence seemed to give rise to a loss of conviction in one's first-order attitude, rather than a substantive change to that attitude. In an IP framework, this is represented by revision consisting of an increase in the imprecision of the credal state, rather than in a shift to a different level of credence. But the framework also allows for other kinds of cases which do produce a substantive shift in credences as well as a change in conviction. Epistemic akrasia can be avoided because the first-order attitude is revised so that the different levels are not in tension. The higher-order uncertainty that we have is reflected in the lack of precision of the revised first-order doxastic attitude.

Furthermore, standard decision theories associated with IP do produce the kind of behavioural effects we typically associate with the impact of higher-order evidence. The amount of precision in a credal state has implications for how you are prepared to behave. One way to see this is to appeal to the behavioural interpretation of lower and upper probabilities (Walley, 1991). The lower probability of a can be interpreted as the maximum price you would pay for a gamble that pays you one unit if a is true and nothing otherwise. The upper probability is the lowest amount you would sell such a gamble for. Thus, for prices below $\underline{p}(a)$ you are definitely disposed to buy the gamble, and for prices above $\bar{p}(a)$ you are definitely inclined to sell it. For prices in between there is no definite disposition either way for either buying or selling. For a Bayesian working with only precise probabilities, the lower and upper probabilities always coincide and the probability then gives the 'fair price' that you would have for the gamble and you would be prepared to either buy or sell it at that price (De Finetti, 1931). Suppose then you are offered some particular bet on a – say where you will get 30 euros if a is true, and will lose 20 euros if a is false. If you have a precise opinion represented by $\{p(a) = 0.5\}$, you will have a definite preference to take the bet, since it gives you positive expected utility. However, if $\{p(a) : 0 \leq p(a) \leq 1\}$ is your state of opinion, you actually have no definite dispositions to either take or reject the bet. You can think of your set of probability measures here as like a 'credal committee' consisting of different members (Joyce, 2010). For prices in between the lower and upper probability, you are in a situation where some members of your committee advocate buying, but others advocate selling.

Nonetheless, since you may need to act even when your credal set is imprecise, decision theories have been developed for these models. There are a number of these. But a common decision principle is that you should avoid a bad worst-case scenario. As we will now see, the use of this kind of principle can give rise to the behavioural consequences of hesitating to take action (if delaying is an option), or 'playing-it-safe' with respect to the risk of different options.

For definiteness, I will illustrate using the prominent decision theory due to Gärdenfors and Sahlin (1982).⁸ According to this theory, one should first assess which of the probability measures in the credal set are 'admissible', that is, which count as 'serious possibilities' as opposed to just possibilities. One should then choose from the options the one which has the largest minimal expected utility, given the admissible probability measures.

Suppose now that the effect of higher-order evidence is to increase the degree of imprecision of the credal state of the agent, reflecting their loss of conviction in their opinion. Consider again the decision problem faced by Anton represented in Table 1. Anton initially has a high precise credence in A , say 0.9, and then he receives higher-order evidence which increases the imprecision of his credal state. This could be modeled by a kind of 'discounting' operation. This method has been used to model sensor or expert reliability, and is often employed as part of an aggregation

procedure where the information from various sources is combined into one belief state, taking account of the reliability of each source by discounting it appropriately (Moral & Del Sagrado, 1998; Mercier et al., 2008; Moral, 2018; Stewart & Quintana, 2018). Discounting involves making a convex combination of the original credal state (with weight $1 - \alpha$) and the fully vacuous state $[0,1]$ (with weight α). The result is a more imprecise credal state, with imprecision equal to α . It is like mixing the initial state with a state of complete ignorance. The reliability is given by the weight given to the original credal state $r = 1 - \alpha$.⁹ Suppose then that because Anton is told that doctors in his situation are only 60% reliable, he discounts himself as a source with a weight given by $r = 1 - \alpha = 0.6$. The result of doing this is an imprecise credal state $\{p(s_1) : 0.54 \leq p(s_1) \leq 0.94\}$.¹⁰

The consequence, according to the decision theory of Gärdenfors and Sahlin, is that Anton should now choose C, if we assume that all the probabilities in the set $\{p(s_1) : 0.54 \leq p(s_1) \leq 0.94\}$ count as serious possibilities. That is because option C has a better worst case than either A or B – the minimal expected utility for C is zero, whereas for A it is -34, and for B it is -11.02. Thus, according to this theory, the increase in imprecision of the credal state would lead to taking the option of giving the patient neither medicine. This can be thought of as taking the option of hesitating to act, or postponing the decision.

In a forced choice, where Anton has to choose immediately between options A and B (i.e. a case where C is not one of the options), the theory recommends that Anton choose the less risky option B.

Overall, although there can be differences in the specific verdicts of different IP decision theories,¹¹ they do tend to recommend the behavioural consequence that we argued should be expected as a result of receiving higher-order evidence, namely a tendency to avoid taking immediate action or opting for the less risky option.

5 | RELATION TO OTHER VIEWS

In this paper, I have argued that an IP-framework provides a natural setting to tackle the issues raised by cases of higher-order evidence. One might still wonder however whether we really need all the extra complexity of an imprecise framework. Is there not some alternative way that a precise Bayesian could handle the problem?

There is of course a large debate in the background here. Many people are not convinced in general of the need or desirability of imprecise probability, and argue that precise Bayesianism has various advantages which recommend it as a superior normative theory (Carr, 2019; White, 2010). I cannot get into this general debate here.¹² What I have hoped to show is that some of the things that seem puzzling about higher-order evidence might be more routine if viewed from an IP framework. In particular, this framework allows for an explicit way of representing conviction in one's first-order attitudes and how it can be lost. Such loss of conviction seems to be precisely what is involved in higher-order evidence cases, and representing such a change in doxastic state has long been alleged to present a difficulty for precise Bayesianism.

In a recent paper, Steglich-Petersen has presented a view which attempts to accommodate some of these difficulties, but without leaving the precise Bayesian framework. He appeals to an already existing distinction in the Bayesian framework between the level of a credence and the 'resilience' of the credence. The notion of resilience has been used to account for weight of evidence (Skyrms, 1977). Earlier we considered the difference between a credence of 0.5 that a coin will land heads when one has no information about the coin, and a credence of 0.5 when one has the results of many tosses, showing a symmetrical distribution of heads and tails. It has been proposed that

the difference here be understood in terms of the ‘resilience’ of the credence. This measures how much the level of credence should change in the face of additional data. In the first case, the credence has low resilience, because throwing a sequence of heads in a row will produce a large shift in the credence level, whereas in the second case, it would have less impact, because the credence will stick closer to the prior.

Steglich-Peterson proposes that in lopsided cases, higher-order evidence provokes revision of both the level and the resilience of credences. In non-lopsided cases, higher-order evidence may undermine the resilience of a credence, but without shifting the agent to a different level of credence. This gives us a way to understand cases like Holmes. Here Holmes’ testimony lowers the resilience of Watson’s credence that the butler did it, but without requiring him to shift to a different level of credence in the butler’s guilt. Steglich-Petersen argues that ‘it is not necessarily irrational to hold some level of credence while being doubtful that that level is rational’ (p. 224). The resilience view then captures the intuition that in non-lopsided cases, the evidence that serves as higher-order evidence does not appear to give reason to revise the first-order attitude in any particular direction.

Steglich-Peterson also argues that his view avoids commitment to the rationality of epistemic akrasia in a form such as Watson believes that the butler did it, while also believing that the evidence does not support that first-order belief. To make this case, he relies on a certain view of what full belief involves: namely, full belief in a proposition p requires not just a high level of credence in p but also a certain level of resilience.¹³ By making Watson’s credence lose resilience, Holmes’ testimony can have the effect that Watson no longer believes that the butler did it, even though he has a high credence for that. This means that Watson’s beliefs are not akratic in the above sense, because he actually does lose his first-order belief in the butler’s guilt.

The resilience view thus arguably has some of the advantages I have claimed also attach to my view, and it has some intuitive appeal. However, the view I have proposed has an advantage in terms of dealing appropriately with behavioural effects. I have argued that higher-order evidence should have an effect on action – namely moving one to making more conservative decisions or to not acting at all, if that is an option. The resilience view, on the other hand, would have the same consequences for action before and after the higher-order evidence was received. This is because the credence level is the same, and so the same decisions get made.

It is true that the higher-order evidence is supposed to have behavioural effects on the resilience account. In particular, a lowering of resilience is supposed to make one more responsive to new evidence. That means that one may, given new evidence, revise one’s view and thus be prepared to undertake a different action. However, a change in behaviour seems reasonable not just after obtaining new evidence, but in the situation before new evidence is received. But here the resilience view would imply no effect.

There may be certain ways that a proponent of the resilience view could find to address this difficulty. One might, for example, try to combine the view with a different kind of decision theory which takes resilience into account. This would need to be worked out. Or perhaps the idea is that it is full belief that guides action. This would be an unusual view, since most decision theories make use of probabilities. Nonetheless one could perhaps make use of the stability theory of belief Leitgeb (2017). Or one might appeal to considerations about what kind of policy with respect to betting will have maximise expected value over the long-run (Coates (2012), p. 121). At the very least, more needs to be said about the implications of the resilience view for decision-making. The advantage of the view that I have presented is that it reflects the impact of higher-order evidence in a framework which gives appropriate behavioural implications according to the usual decision-theoretic principles associated with IP models.

It is also worth noting that there have been some sophisticated attempts to handle higher-order evidence by combining a Bayesian approach with epistemic logic, building on work by Timothy Williamson (Dorst, 2019a, 2019b; Lasonen-Aarnio, 2015; Williamson, 2019, 2000). It is beyond the scope of this paper to assess the relation of this proposal to those efforts, though the comparison would be an interesting one.

In this paper, I have argued that imprecise probability provides a suitable framework for developing the kind of two-dimensional representation of doxastic attitude which helps to fully understand cases involving higher-order evidence. I do not mean to suggest that this is the only possible way to develop the view. There are other modelling frameworks which also have a two-dimensional character. For example, some authors have developed theories involving higher-order probabilities (Gaifman, 1988), and there are also decision theories which specifically take account of ‘confidence’ in a probability distribution (Bradley, 2017; Hill, 2013). I will not attempt to adjudicate the relative merits of these approaches here.¹⁴ What I hope to have done is to make a case for the usefulness of two-dimensional representations of doxastic attitudes in analysing cases involving higher-order evidence, using imprecise probability models as an illustration.

6 | CONCLUSION

There appear to be two kinds of effects which evidence in ‘higher-order evidence’ cases can have. Such evidence can call for a substantive change in our first-order attitude – for example, if it tells us that one first-order proposition is more likely to be true than another. But in some cases, the effect of the evidence is more to undermine the conviction with which we hold our first-order attitude. Which of these effects is more dominant depends on the case. Getting a clear picture of the impact of higher-order evidence requires then a representation of doxastic attitudes which captures both the substantive dimension and the conviction dimension. Representing doxastic attitudes using imprecise probabilities has this feature, making it a suitable framework for investigation of the phenomena identified in the higher-order evidence debate. I have argued that making use of such a representation allows us to avoid some of the difficulties associated with both Level-Splitting and Revision views. Consideration of cases of higher-order evidence suggests that such evidence should have implications for the risks we are willing to take in decision-making. Those behavioural consequences can be captured in an IP-framework, by appealing to existing decision principles.

ACKNOWLEDGMENTS

I acknowledge support from a Veni grant (275-20-058) from the Dutch Organisation for Scientific Research (NWO). Thanks to members of the Groningen Epistemology Reading Group and to Teddy Seidenfeld for useful discussion and to Alexander Gebharder and Remco Heesen for helpful comments on earlier drafts of this work.

ORCID

Leah Henderson  <https://orcid.org/0000-0002-8709-9765>

ENDNOTES

¹ My characterisation is similar, though not identical, to that of Kappel (2019).

² This is a more appropriate way of talking than saying that evidence supports a proposition itself. Good reasons for this are given by Worsnip (2018), p. 10.

- ³ Coates puts this as follows: ‘it is a mistake to think of Holmes’s assessment as counterevidence [to Watson’s belief that the butler did it]. Recall that Holmes may criticize Watson’s belief as irrational even though he thinks it is true, if he thinks Watson arrived at it in an irrational manner. Since Watson knows this, he cannot treat Holmes’s assessment as evidence favoring the butler’ (Coates (2012), p. 115).
- ⁴ An example of the contrast here is given by the characters Doubtful Ava and Doubtful Brayden in Christensen (2010b), p. 121.
- ⁵ Note that epistemic akrasia is often phrased as having high confidence that ‘ p , but my evidence doesn’t support p ’, eg. Horowitz (2014). However, as explained above, I take evidence to support a doxastic attitude towards a proposition, rather than supporting that proposition itself.
- ⁶ Proponents of IP often prefer not to use the term ‘uncertainty’ here to emphasise that it is not the same kind of uncertainty as is expressed by the credence itself. Instead they prefer to say that imprecision in the credal state represents ‘indeterminacy’ (Levi, 1974; Walley, 1991).
- ⁷ This is a situation that Popper called the ‘paradox of ideal evidence’ (Popper, 1959). Popper regarded it as ‘a little startling’ that our degree of rational belief in a is completely unaffected by all the extra evidence contained in e and that ‘the absence of any statistical evidence concerning [the coin] justifies precisely the same ‘degree of rational belief’ as the weighty evidence of millions of observations which, *prima facie*, support or confirm or strengthen our belief’ (Popper (1959), p. 426).
- ⁸ An important alternative is the theory of Isaac Levi (Levi, 1980). See Seidenfeld (1988, 2004) for discussion of differences between the theories. Discussion of these and other decision theories for imprecise probability can be found in Troffaes (2007); Williams (2014); Augustin et al. (2014).
- ⁹ A convex combination of two imprecise probabilities $\underline{p}_1(A)$ and $\underline{p}_2(A)$ is given by $\underline{p}(A) = (1 - \alpha)\underline{p}_1(A) + \alpha\underline{p}_2(A)$, and $\overline{p}(A) = (1 - \alpha)\overline{p}_1(A) + \alpha\overline{p}_2(A)$. For a convex combination of a precise probability $p_1(A) = \underline{p}_1(A) = \overline{p}_1(A)$, and the vacuous state $\underline{p}_2(A) = 0$, $\overline{p}_2(A) = 1$, this yields $\underline{p}(A) = (1 - \alpha)p_1(A)$, $\overline{p}(A) = (1 - \alpha)p_1(A) + \alpha$, which has imprecision $\overline{p}(A) - \underline{p}(A) = \alpha$.
- ¹⁰ Anton’s initial credence is precise, so $\underline{p}_1(A) = \overline{p}_1(A) = 0.9$. We mix this with the fully vacuous state with $\underline{p}_2(A) = 0$, $\overline{p}_2(A) = 1$. The resulting lower and upper probabilities are then respectively $\underline{p}(A) = (1 - \alpha)\underline{p}_1(A) + \alpha\underline{p}_2(A) = 0.6 \cdot 0.9 + 0.4 \cdot 0 = 0.54$, and $\overline{p}(A) = (1 - \alpha)\overline{p}_1(A) + \alpha\overline{p}_2(A) = 0.6 \cdot 0.9 + 0.4 \cdot 1 = 0.94$. The imprecision of the resulting state is $\overline{p}(A) - \underline{p}(A) = 0.4$.
- ¹¹ For example, in the choice between the three options, Levi’s theory would recommend that Anton choose the less risky option B, rather than C. The two theories agree on the forced choice case.
- ¹² For defences of IP, see Levi (1974, 2009); Joyce (2010); Schoenfield (2012).
- ¹³ Steglich-Peterson here appeals to the ‘stability theory of belief’ (Leitgeb, 2017)
- ¹⁴ There are of course various points of discussion. For example, some authors would resist the use of probability at the higher-order level (Levi, 1974, 2009; Savage, 1954).

REFERENCES

- Augustin, T., Coolen, F. P. A., de Cooman, G., & Troffaes, M. C. M. (Eds.). (2014). *Introduction to imprecise probabilities*. UK: Wiley.
- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4), 325–370.
- Carr, J. R. (2019). Imprecise evidence without imprecise credences. *Philosophical Studies*, 24(1), 1–24.
- Christensen, D. (2010a). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Christensen, D. (2010b). Rational reflection. *Philosophical Perspectives*, 24(1), 121–140.
- Christensen, D. (2016). Disagreement, drugs, etc.: from accuracy to akrasia. *Episteme*, 13(4), 397–422.
- Coates, A. (2012). Rational epistemic akrasia. *American Philosophical Quarterly*, 49(2), 113–124.
- De Finetti, B. (1931). Probabilism: a critical essay on the theory of probability and on the value of science. *Erkenntnis*, 31(2-3), 169–223.
- Dorst, K. (2019a). Evidence: a guide for the uncertain. *Philosophy and Phenomenological Research*.

- Dorst, K. (2019b). Higher-order uncertainty. In M. Skipper & A. Steglich-Peterson (Eds.), *Higher-order evidence: new essays*. Oxford University Press.
- Elkin, L., & Wheeler, G. (2018). Resolving peer disagreements through imprecise probabilities. *Noûs*, 52(2), 260–278.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers without gods: meditations on atheism and the secular life* (pp. 194–214). Oxford University Press.
- Feldman, R. (2009). Evidentialism, higher-order evidence, and disagreement. *Episteme*, 6(3), 294–312.
- Gaifman, H. (1988). A theory of higher order probabilities. In *Causation, chance and credence* (pp. 191–219). Dordrecht: Springer.
- Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking and decision making. *Synthese*, 361–386.
- Greco, D. (2014). A puzzle about epistemic akrasia. *Philosophical Studies*, 167(2), 201–219.
- Hazlett, A. (2012). Higher-order epistemic attitudes and intellectual humility. *Episteme*, 9(3), 205–223.
- Hill, B. (2013). Confidence and decision. *Games and economic behavior*, 82, 675–692.
- Horowitz, S. (2014). Epistemic akrasia. *Noûs*, 48(4), 718–744.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24, 281–323.
- Kappel, K. (2019). Escaping the akratic trilemma. In M. Skipper & A. Steglich-Peterson (Eds.), *Higher-order evidence: new essays*. Oxford University Press.
- Keynes, J. M. (1921). *A treatise on probability*. London: Macmillan.
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston: Houghton Mifflin.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88(2), 314–345.
- Lasonen-Aarnio, M. (2015). New rational reflection and internalism about rationality. In *Oxford Studies in Epistemology*, volume 5. Oxford University Press.
- Leitgeb, H. (2017). *The stability of belief: how rational belief coheres with probability*. Oxford University Press.
- Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71(13), 391–418.
- Levi, I. (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press.
- Levi, I. (2009). Why indeterminate probability is rational. *Journal of Applied Logic*, 7, 364–376.
- Mercier, D., Quost, B., & Denooux, T. (2008). Refined modeling of sensor reliability in the belief function framework using contextual discounting. *Information Fusion*, 9, 246–258.
- Moral, S. (2018). Discounting imprecise probabilities. *The mathematics of the uncertain*, 685–697.
- Moral, S., & Del Sagrado, J. (1998). Aggregation of imprecise probabilities. In B. Bouchon-Meunier (Ed.), *Aggregation and fusion of imperfect information*. Physica-Verlag.
- Peirce, C. S. (1932). *Collected papers*. Belknap Press.
- Pollock, J. L. (1986). *Contemporary Theories of Knowledge*. Roman and Littlefield.
- Popper, K. (1959). *The Logic of Scientific Discovery*. Routledge.
- Roush, S. (2009). Second-guessing: a self-help manual. *Episteme*, 6(3), 251–268.
- Roush, S. (2017). Epistemic self-doubt.
- Sahlin, N.-E., & Persson, J. (1994). Epistemic risk: the significance of knowing what one does not know. In B. Brehmer & N.-E. Sahlin (Eds.), *Future Risks and risk management*. Springer.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Dover Publications, Inc.
- Schoenfield, M. (2012). Chilling out on epistemic rationality: a defense of imprecise credences (and other imprecise doxastic attitudes). *Philosophical Studies*, 158, 197–219.
- Schoenfield, M. (2015). A dilemma for calibrationism. *Philosophy and Phenomenological Research*, 91(2), 425–455.
- Schoenfield, M. (2018). An accuracy based approach to higher order evidence. *Philosophy and Phenomenological Research*, 96(3), 690–715.
- Seidenfeld, T. (1988). Decision theory without "independence" or without "ordering": what is the difference? *Economics and Philosophy*, 4, 267–290.
- Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: Gamma-maximin versus e-admissibility. *Synthese*, 140, 69–88.
- Skyrms, B. (1977). Resiliency, propensities and causal necessity. *Journal of Philosophy*, 74(11), 704–713.
- Sliwa, P., & Horowitz, S. (2015). Respecting all the evidence. *Philosophical Studies*, 172(11), 2835–2858.

- Smithies, D. (2012). Moore's paradox and the accessibility of justification. *Philosophy and Phenomenological Research*, LXXXV(2), 273–300.
- Steglich-Peterson, A. (2019). Higher-order defeat and doxastic resilience. In *Higher-order evidence: new essays* (pp. 209–225). Oxford University Press.
- Stewart, R. T., & Quintana, I. O. (2018). Probabilistic opinion pooling with imprecise probabilities. *Journal of Philosophical Logic*, 47(1), 17–45.
- Titelbaum, M. G. (2015). Rationality's fixed point (or: in defense of right reason). In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, volume 5. Oxford University Press.
- Troffaes, M. C. M. (2007). Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45, 17–29.
- Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. Chapman and Hall.
- Weatherston, B. (ms). Do judgments screen evidence?
- Wedgwood, R. (2012). Justified inference. *Synthese*, 189(2), 273–295.
- White, R. (2009). On treating oneself and others as thermometers. *Episteme*, 6(3), 233–250.
- White, R. (2010). Evidential symmetry and mushy credence. In *Oxford Studies in Epistemology*, volume 3 (pp. 161–186). Oxford University Press.
- Williams, J. R. G. (2014). Decision-making under indeterminacy. *Philosophers' Imprint*, 14(4), 1–34.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Williamson, T. (2011). Improbable knowing. In *Evidentialism and its discontents*. Oxford University Press.
- Williamson, T. (2019). Evidence of evidence in epistemic logic. In M. Skipper & A. Steglich-Peterson (Eds.), *Higher-order evidence: new essays*. Oxford University Press.
- Worsnip, A. (2018). The conflict of evidence and coherence. *Philosophy and Phenomenological Research*, 96(1), 3–44.

How to cite this article: Henderson L. Higher-order evidence and losing one's conviction. *Noûs*. 2021;1–17. <https://doi.org/10.1111/nous.12367>