



# The Practical Import of Higher-Order Defeat: Resilience vs. Imprecise Credences

Jakob Donskov<sup>1</sup> · Asbjørn Steglich-Petersen<sup>2</sup>

Received: 13 June 2024 / Accepted: 2 February 2025 / Published online: 6 March 2025  
© The Author(s) 2025, corrected publication 2025

## Abstract

In some cases of higher-order defeat, you rationally doubt whether your credence in  $p$  is rational without having evidence of how to improve your credence in  $p$ . According to the resilience framework proposed by Steglich-Petersen (Higher-order defeat and Doxastic Resilience), such cases require loss of doxastic resilience: retain your credence level but become more disposed to change your mind given future evidence. Henderson (Higher-Order Evidence and Losing One's Conviction) responds that this allows for irrational decision-making and that we are better off understanding higher-order defeat in terms of imprecise probabilities. We argue first that Henderson's imprecise probability framework models the wrong kind of thing. Credal imprecision is neither necessary nor sufficient for higher-order doubt. Second, we offer two ways of understanding the practical import of higher-order defeat given loss of doxastic resilience.

## 1 Introduction

How should you deal with evidence of your own irrationality? Recent work has framed this question in terms of *theoretical* rationality: what is the epistemically rational response to higher-order evidence? As higher-order evidence bears on your own rationality, it can induce doubts sufficient to make it irrational for you to maintain your beliefs. In other words, higher-order evidence can defeat first-order beliefs. For example, a stressed doctor who is informed that she has not slept for 24 h should

---

✉ Jakob Donskov  
j.g.donskov@leeds.ac.uk

✉ Asbjørn Steglich-Petersen  
filasp@cas.au.dk

<sup>1</sup> School of Philosophy, Religion and History of Science, University of Leeds, Leeds, UK

<sup>2</sup> Department of Philosophy and History of Ideas, Aarhus University, Aarhus, Denmark

doubt the diagnosis she just formed. The information about sleep deprivation is higher-order defeating as it undermines the doctor's justification for believing her diagnosis by challenging her reliability. But this is not the only effect such higher-order evidence should have. It should also influence which *decisions* the doctor should make, that is, such evidence also bears on what it is *practically* rational for her to do. At first glance, it seems obvious that since the doctor's beliefs are defeated by the higher-order evidence, this ought to call for more cautious decisions regarding the patient's care.

However, one of the reasons higher-order evidence has received much attention in recent years and why it seems puzzling from the viewpoint of theoretical rationality, is due to the fact that higher-order evidence can defeat beliefs without indicating that those beliefs are false. Evidence of sleep deprivation is not evidence that the diagnosis is incorrect. This problem is only highlighted when framed in terms of credences. Evidence for sleep deprivation should challenge the doctor's high confidence in her diagnosis. That is, the higher-order evidence should lead to higher-order doubt. But it also seems that such doubt should be reflected in the credences themselves. Yet it is not clear how this should be done. Why think that the rational response is for her to *lower* her credence? The higher-order evidence says nothing about the probability of the diagnosis, and so it seems arbitrary for the doctor to respond by judging that the probability of the diagnosis is lower than previously estimated.

This has led Steglich-Petersen (2019) to propose that higher-order defeat can work by undermining the *resilience* of one's credences rather than their *level*. The resilience of your credence in  $p$  is a measure of how disposed you are to change your mind about  $p$  given new evidence. When you receive higher-order evidence of your own rational failure, this influences how likely you can rationally believe your credences to reflect your evidence. In light of this, you should become more disposed to changing your credence when new evidence comes in. As such, both the level of a credence and the resilience of that credence are measures of uncertainty. Both dimensions of your credences reflect your uncertainty, albeit about different matters. While your credence level reflects your uncertainty about the target proposition, your degree of resilience reflects your uncertainty that your credence level is rational. In principle, then, higher-order evidence can require a loss of resilience of a credence without requiring a change of the credence level.

This suggests a way in which higher-order defeat can be dealt with from the viewpoint of theoretical rationality. However, this theoretical solution seems to cause trouble in the practical domain. If the credence level sometimes remains intact and decisions are informed by credences, how can higher-order evidence influence decisions in those situations? This question has recently been posed as a challenge to the resilience framework by Henderson (2022). She agrees that standard talk of credence levels alone cannot capture the sense in which higher-order evidence induces uncertainty in every case of higher-order defeat. More is needed if we are to capture the sense in which sleep deprivation, for example, defeats the doctor's confidence.

Where Henderson disagrees is in the solution. She argues that the resilience framework sits poorly with decision theory. More precisely, she argues that in cases where the resilience framework recommends that people maintain their credence levels in the face of higher-order defeat, such defeat will have no effect on which decisions

they should make. Instead, Henderson argues that the uncertainty induced by higher-order evidence should be modeled as the degree of *precision* one's credence has, and therefore proposes to analyze higher-order defeat in terms of *imprecise probabilities*.

In this paper we offer two cases which challenge the adequacy of modeling rational higher-order doubt using imprecise probabilities. We then offer two distinct ways of thinking about how doxastic states relate to rational decision-making such that loss of credal resilience does indeed influence rational action. The first appeals to the recently popular view that either knowledge or rational belief is the norm of practical reasoning. We argue that since rational belief requires resilient credence, higher-order evidence can influence rational action by defeating resilience. Second, we argue that low resilience can influence rational decision-making by rendering credences inadmissible in the agent's practical deliberation.

The paper proceeds as follows. In §2, we introduce the idea that higher-order evidence should sometimes affect the resilience of one's credences rather than their level. In §3, we describe Henderson's objection that this view fails to capture how higher-order defeat should affect practical decisions and discuss her own proposal of understanding higher-order doubt in terms of imprecise probabilities. In §4, we outline two different ways in which the resilience of a credence can affect rational decision-making, thus defusing Henderson's challenge. §5 is a brief summary and conclusion.

## 2 Higher-order Evidence Can Defeat Credal Resilience

Higher-order evidence defeats by inducing doubt that the credence you have in some proposition is rational given your epistemic situation. In many cases of higher-order defeat, the rational response seems to require revising one's credence. For example, learning that 93% of people in your demographic tend to think they are better than average drivers should make you lower your credence that *you* are a better than average driver.<sup>1</sup> Here, the higher-order evidence indicates that some of your credences deviate from what is rational in some particular manner; that you are *too* sure of being a good driver. In such cases, higher-order evidence calls for lowering one's credence.

However, there are also cases where it seems rationally permissible for you to maintain the level of credence you have in a proposition, even in light of rational doubt that this is the right level of credence. Here are two cases discussed in the literature:

**Doubtful Ava:** “[...] Ava is considering the possibility that (D) the next U.S. President will be a Democrat. She gives D some particular credence, say, 7; this reflects a great deal of her general knowledge, her feel for public opinion, her knowledge of possible candidates, etc. But given the possibility that her credence is affected by wishful thinking, protective pessimism, or just failure to focus on and perfectly integrate an unruly mass of evidence, Ava very much

<sup>1</sup> One classic study found just that: Svenson (1981).

doubts that her credence is exactly what her evidence supports.” (Christensen, 2010b: 121).

**Doubtful Bob:** “A taxicab has been involved in a hit and run accident at night, and Bob has been called to serve on a jury in court, where a driver from Blue Cabs stands accused. There are just two taxi companies in town, Green Cabs being the dominant one. Only 5% of the taxis in town are blue, 95% are green. A witness saw the incident, and identified the cab as blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colours 80% of the time, and failed 20% of the time. Bob is carefully considering how likely it is that the cab was blue rather than green in light of the available evidence, and ends up giving it credence.<sup>17</sup> However, Bob always found probabilities tricky, and his grasp of Bayes’ Rule is pretty hazy. He thus severely doubts that his credence is right given the evidence. In fact, he would not be very surprised if he is quite far off the mark.” (Steglich-Petersen, 2019: 210).

In these cases, it seems rational for Ava and Bob to doubt that their credences are rational. But this does not in and of itself seem to undermine the rationality of maintaining the levels of their credences. Of course, these levels of credence might be irrational for other reasons: perhaps they have misinterpreted the available evidence, or perhaps their credences are inconsistent with some of their other beliefs. But the fact that they harbor rational doubts that their credences are rational in light of their evidence should not in itself move them to revise their credences. Even if Ava and Bob should doubt that their credences are rational, they may well be the most rational credences for them to maintain.

Why do some cases of rational higher-order doubt call for revision of one’s credence level while other cases do not?<sup>2</sup> The explanation has to do with the nature of higher-order evidence. Different definitions have been offered. Higher-order evidence has been defined as evidence about the existence, merit, or significance of one’s evidence (Feldman, 2009); as evidence which indicates that the agent has evaluated their available evidence in an irrational way (Christensen, 2010a); and as evidence indicating that one’s doxastic state is the result of a flawed process (Lasonen-Aarnio, 2013). But common between these characterizations is that higher-order evidence

<sup>2</sup> A reviewer asks whether it has to do with when the higher-order evidence is acquired. If you first form a credence and then learn it is irrational, revision is required. If you form a credence while realizing that you might be irrational, revision might not be required. That is an interesting suggestion, but it cannot be the full story. In particular, it does not explain *how* one should revise. If we suppose that Ava has already heeded her worries about wishful thinking, *we*, the theorists, might realize that her 0.7 credence is rational. Still, it seems that *she* should doubt whether her credence is in fact rational. The same goes if we suppose that Ava first forms her 0.7 credence and only then learns of the possibility of wishful thinking. It seems that in either variation of the case, Ava should exhibit some doubt that her credence is rational even though it is unclear to her how she should revise the credence so as to improve it. The same holds true of Bob. Suppose he first does the calculations, arrives at his answer 0.17%, and only then learns that he has nearly flunked his statistics course. Even then it remains unclear how Bob can revise his 0.17 credence to accommodate his doubt.

bears only *indirectly* on the truth-value of the proposition which one has a first-order belief in or gives credence.

For example, evidence that Bob is a poor Bayesian reasoner is higher-order evidence which seems to challenge the rationality of his 0.17 credence that the cab was blue. But the proposition that *Bob is a poor Bayesian reasoner* does not by itself affect the probability that the cab was blue. Indeed, one possibility you cannot rule out given this type of evidence is that your credence is exactly right given your first-order evidence, but that you have arrived at that credence through an irrational process; that you somehow did not base your credence on the evidence in a rational way. Still, the proposition seems to bear indirectly on whether the cab was blue; it is clearly relevant information for Bob's thinking through the matter.<sup>3</sup>

Importantly, there are two ways in which higher-order evidence can bear indirectly on a proposition. Christensen (2010b) draws attention to a distinction between *lopsided* and *non-lopsided* defeat. Evidence that you have *overestimated* your driving capabilities should be accommodated by *lowering* your credence. That is, the higher-order evidence about overconfidence already contains information indicating how your credence has deviated from the ideal level. Another example comes from cases of revealed peer-disagreement: two parties learn what credence their peer has given some proposition. Here the evidence lopsidedly indicates how your credence should be revised: you accommodate this information by revising your credence in the direction of your peer (Christensen, 2007).

Non-lopsided defeat, by contrast, tells you nothing about how you are falling short of rationality. Learning of sleep deprivation, intoxication, and other reasoning distorting impairments are often evidentially relevant but need not tell you in what way your credence is irrational.<sup>4</sup> For example, evidence that some of your credences are influenced by wishful thinking or based on invalid reasoning is relevant evidence, yet it does not indicate how you should revise your credences.

For this reason, it seems unmotivated to change your credence when faced with non-lopsided defeat. Consider the doctor case mentioned above in more detail:

**The Sleep Deprived Doctor:** After thorough examination, a doctor's credence that her patient suffers from coeliac disease is 0.9. Now a colleague pops in. Surprised that the doctor is still at work, he informs her that she has not slept for 24 h and so is in no condition to be diagnosing patients. The doctor trusts her colleague and although she is surprised to hear this at first, she immediately recognizes how exhausted she is.

Suppose the doctor merely responds by revising her credence down to, say, 0.3 that her patient suffers from the disease. This cannot be the rational response. For one,

<sup>3</sup> We might say that a piece of evidence is higher-order not in virtue of its content but in virtue of which propositions it is taken to be evidence about. So, we need not assume that first- and higher-order evidential support relations are mutually exclusive. A piece of evidence might both indicate the truth-value of one proposition and indicate what credence in some other proposition reflects the evidence (Skipper, 2020, Hedden & Dorst 2022).

<sup>4</sup> See Elga (ms) for the typical hypoxia case, and Christensen (2010a) for several other cognitive impairment cases.

there is nothing which tells her that the probability is 30%. But the problem is not just that adopting a specific, lower credence will be unfounded by the evidence as her evidence points to no specific lower credence. One might think that all the doctor can be required to do given such unspecific evidence of sleep deprivation is to adopt any credence “well below” 0.9. Such a suggestion still leaves the main problem untouched: she doesn’t even know that her initial credence was too high. For all she knows, the probability could be higher than 90% but her sleep deprivation made her ignore symptoms which strongly indicate coeliac disease.

More generally, merely changing one’s credence in the face of non-lopsided defeat seems problematic as it does not capture the relevant kind of uncertainty such evidence ought to induce. To see this, notice that we do not fix every aspect of your uncertainty about a given proposition merely by fixing a particular credence level in that proposition. Stipulating that the doctor holds a 0.9 credence that her patient has coeliac disease does not fix what mental state the doctor is in. Which mental state is represented depends on the entire probability distribution which makes up the doctor’s epistemic space. In particular, it seems plausible that which mental state is represented in a 0.9 credence in  $p$  depends on probability assignments to various propositions on which  $p$  depends.

This is a natural thought: how confident you are that  $p$  depends in part on how disposed you are to change your mind about whether  $p$ . Your 0.9 credence in  $p$  might be subject to a significant change if you learn that  $q$ —that happens if your credence in  $p$  conditional on  $q$  is significantly different from 0.9. Alternatively, if your various conditional credences in  $p$  given  $q$  and  $p$  given  $r$  and  $p$  given  $z$ , etc. are all close to 0.9, then your unconditional credence in 0.9 is *resilient*. It is a credence that will remain relatively fixed as you gain information. Steglich-Petersen’s (2019) proposal is that higher-order evidence can influence rational belief by constraining the resilience of one’s credences. That is, when you doubt that your credences reflect your evidence, you should be disposed to change your mind more readily given new evidence.

The concept “resilience” was originally formulated by Skyrms (1977, 1980). The motivation for introducing it goes back to Peirce (1931) and Keynes (1921), both of whom complain that a single probabilistic measure does not reflect evidential strength and so propose to add a second, wholly distinct value to credal states. Skyrms points out that the conditional probabilities already included in the Bayesian probability distribution could do the work. When you receive new evidence, you update the entire probability distribution. This includes updating your conditional probabilities. Once we pay attention to the relationship between the conditional and unconditional credences, it becomes apparent that evidential strength is already reflected in the probability distribution as a whole.<sup>5</sup>

<sup>5</sup> Skyrms (1977: 707–708) defines resilience as probabilistic invariance over some “sublanguage” which specifies a set of propositions relevant to one’s epistemic context. If you are resilient, learning of these propositions will not make you change your mind much. Skyrms himself is skeptical that a unique answer can be given as to what exactly the relevant sublanguage is. Leitgeb (2015) proposes that a given context specifies a set of “serious possibilities” for an agent which can play this role. We return to this in Sect. 4.1.

It is worth noting that an alternative route is possible. Following Jeffrey (1983: § 12.5), resilience might be thought of as a relation between your current and previous doxastic state. Learning some proposition  $q$  would have changed your mind about  $p$  significantly yesterday. But today, learning that  $q$  does not change

For example, a meteorologist might receive initial reports which suggest that the probability of rain tomorrow is 10% and therefore rationally form a credence of 0.1 that it will rain. Suppose the meteorologist receives more evidence throughout her day, all of which suggests that rain is indeed 10% likely. Here the probability of rain remains the same, even as she amasses more and more evidence bearing on the hypothesis. Through this process, her level of credence ought to remain constant, but it seems that something else should change in her credal state to reflect the more substantial body of evidence. The idea is that it is the resilience of her 0.1 credence which should increase. Her conditional credences that it will rain given some new evidence  $q$  should come closer to 10%. In turn, this reflects her increased higher-order confidence that her credence in rain should be 0.1.

Thus, how we should interpret a rational 0.1 credence that it will rain tomorrow or a rational 0.9 credence that the patient has coeliac disease depends on how resilient those credences are, which, in turn, depends on how strong the available evidence is that rain is 10% likely or that the probability of coeliac disease is 90%. More generally, how resilient you ought to be is a function of the strength of your evidence. This allows us to link evidential strength with higher-order confidence via resilience: stronger evidence means more resilience which permits more higher-order confidence.<sup>6</sup>

Steglich-Petersen's (2019) suggestion, then, is that non-lopsided higher-order evidence can defeat your credence by making you doubt that you have rationally evaluated what the evidence supports or that you have rationally gauged the strength of that evidence, thus calling for you to lower the resilience of your credence. Initially, the doctor formed her credence after having examined the patient, where the examination gave rise to, let us assume, strong evidence that the patient is quite likely to suffer from the disease. When she learns of her sleep deprivation, this suggests that the examination provided unreliable evidence and hence provides significantly weaker evidence than first anticipated. As such, the rational response for the doctor is to lower the resilience of her credence in the diagnosis. Yet, the new evidence says nothing about the probability of coeliac disease, so she has no grounds for changing her probability estimate and hence she should retain her credence level.

---

your mind about  $p$  much. For example, yesterday you knew nothing about the bias of the coin, so observing 10 heads in a row would have significantly increased your credence that the coin is biased. But today, after observing an approximately even number of heads and tails distributed randomly over 1000 flips, observing 10 heads will *not* significantly increase your credence that the coin is biased. Here the increased resilience of your later credence is specified by the difference in your earlier and later conditional credence in bias given observing 10 heads. So, even if there are going to be certain propositions which, if learnt, will drastically change your mind both before and after observing the 1000 flips, the later credence will be more resilient relative to the earlier.

Steglich-Petersen's (2019) idea which is relevant for present purposes is that just as evidence can serve to increase your resilience, so evidence can serve to decrease it. In particular, evidence of your own rational failure can require you to lose resilience in a proposition you would otherwise have been more resilient towards.

<sup>6</sup> Leitgeb (2014, 2017) has further developed the idea, showing that by paying attention to how the conditional probabilities evolve, we can account for rational stability and thereby link rational credence with rational belief. We will return to this below.

**Table 1** Desirability matrix for the sleep deprived doctor

	S1	S2
Administer drug A	20	-80
Administer drug B	-13	20
Postpone treatment	0	0

### 3 Decision-making and Henderson's Alternative Account

We are now in a position to lay out Henderson's (2022) charge that the resilience framework cannot account for the behavioral changes that higher-order defeat intuitively ought to have. The objection is straightforward: higher-order defeat ought to have an effect on one's behavior. In cases of non-lopsided higher-order defeat, the resilience view allows for maintaining one's credence. Since rational decisions are guided by credences, the view gets things wrong in at least those cases. Thus, the resilience view permits rationally dubious behavior as it effectively allows one to disregard the higher-order evidence in one's practical reasoning. A plausible theory would, by contrast, promote decisions associated with more uncertainty.

To bolster this argument, Henderson considers and fills in the details of the sleep deprived doctor case. The doctor is trying to determine whether her patient suffers from disease S1 or S2, and thereby determine which course of treatment to recommend. Henderson provides a desirability matrix which specifies the expected utility associated with each course of action relative to the possible states of nature. We replicate this matrix in Table 1.

Given these desirability assignments and the assumption that the decision table exhausts the possibilities, if the doctor seeks to maximize expected utility, she should administer drug A given a credence in S1 equal to or above 0.8, drug B if her credence in S1 is equal to or below 0.606, and should otherwise postpone treatment.

The doctor's initial credence that the patient suffers from S1 was rationally high given her observations, and so she rationally decides to administer drug A. When she later receives evidence of being sleep deprived, this ought to induce doubts severe enough to make it irrational for her to stand by her decision. She ought to recommend the less risky option that is drug B or postpone treatment if possible. As Henderson points out, however, if the doctor's credence level in S1 remains unchanged, the same course of action maximizes expected utility. Thus, according to Henderson the resilience view recommends administering drug A even after learning of sleep deprivation.

#### 3.1 Henderson's Proposal: Higher-order Defeat Calls for Imprecision

Before we respond to Henderson's challenge, we will consider her own preferred view. She claims that invoking imprecise probabilities allows for a view which makes sense of higher-order defeat in both the theoretical and practical domain.

The central claim of imprecise probability theories is that we should represent doxastic states by a cluster of probability functions, rather than a single probability function as standard Bayesianism would have it (Levi, 1980; Joyce, 2012; Bradley, 2017). Imprecise probabilities can thereby map imprecise doxastic states as credences which take a range of values. Suppose, for example, that your evidence regarding  $p$

is compatible with every probability function in the interval  $[0.4, 0.7]$  and so, the thought goes, your credence ought to match this imprecise probability (Joyce, 2010). The doxastic states of real human beings might very well be imprecise and hence modeled in many cases along the suggested lines. At least, we will not dispute that assumption here. And so, we will not dispute that the difference between the upper and lower bounds of the interval measures how imprecise that doxastic state is.

However, Henderson's key claim, the claim we reject below, is that the degree of imprecision itself represents a degree of uncertainty of a certain kind. Greater imprecision means greater uncertainty. If your credence in  $p$  is  $[0.4, 0.7]$  and mine is  $[0.5, 0.7]$ , I am more confident in  $p$  even though our maximal assignment is identical. To make this point, Henderson distinguishes between two types of uncertainty. First, there is (regular) uncertainty about whether a particular state of affairs obtains which is represented by the level of credence one has in that state of affairs obtaining. Importantly, this uncertainty pertains only to the contents of your credences. The uncertainty is therefore represented by the level of your credences.

The second type of uncertainty is uncertainty about the rationality of one's own credences. Henderson calls this "conviction". She says, "the degree of imprecision in an [imprecise] credal set can be used to represent greater 'uncertainty' about the probability" (p. 521). Your degree of conviction is then the degree to which you are sure your credences are rational. The idea is that as your confidence in the rationality of your credences grows—perhaps because you gather more evidence which supports your initial (median) credence level—your credal set narrows. You become more sure your (median) credence is rational and so rule out the outlying probabilities from your credal set. By contrast, defeating higher-order evidence defeats conviction. That is, such defeat calls for broadening the credal set. Roughly, if you learn that your credence is irrational, you should broaden the interval to include more of the outlying probabilities.

How, then, are decisions influenced according to the imprecise probabilities model? Henderson illustrates how an imprecise credal set can impact decisions using Gärdenfors and Sahlin's (1982) theory: once you have broadened your credal set, you calculate the expected utility of each action according to every probability assignment included in your imprecise state. Having done that, you "should then choose from the options the one which has the largest minimal expected utility, given the admissible probability measures" (Henderson, 2022: 523). As Gärdenfors and Sahlin (1982: 371) put it: "The alternative with the largest minimal expected utility ought to be chosen."

For concreteness, assume that the doctor starts out having a precise 90% credence that S1 obtains and that the higher-order evidence about sleep deprivation indicates that the doctor is only 60% reliable. Henderson's suggestion is to perform a convex combination of the initial credal state (0.9) and the fully vacuous state ( $[0, 1]$ , i.e. the maximally imprecise credal state) with the reliability measure acting as weights on the initial and fully vacuous states. This gives an imprecise probability of  $[0.54, 0.94]$  that S1 obtains.<sup>7</sup>

<sup>7</sup>If you are 60% reliable, then there is a 40% chance you are not. So, to get the lower bound, you multiply the reliability with the credence and add the lower bound of the fully vacuous state multiplied by those

The doctor can then calculate the expected utilities of each decision given every probability assignment included in the imprecise credence. This leaves the doctor with an enormous set of “expected” utilities for each act. For example, the expected value of administering drug A will be -26 according to the lower bound (0.54) of the doctor’s credal set. According to the almost-lower bound (0.55), the expected utility of drug A is -25, i.e. slightly better. However, according to the upper bound (0.94), the doctor expects 14 utilities for administering drug A. Given this array of expected utilities for each action, how is the doctor to decide according to Henderson? She cannot simply pick the act which maximizes expected utility simpliciter, since the various probability assignments in her imprecise credal state yield different expected utilities. So, Henderson recommends picking the act that, given all of the admissible probability measures, has the largest minimal expected utility. (Henderson, 2022: 523). This gives the reasonable result that the doctor should postpone treatment if possible, and should otherwise administer drug B.

We can thus summarize Henderson’s view as follows. The effect of non-lopsided higher-order defeating evidence is to broaden one’s credal set. More probability distributions become compatible with one’s evidence and so each different probability distribution should be included in one’s doxastic state. This gives rise to an imprecise credence. Since each probability in the credal set gives different expected utilities for the different decisions, Henderson recommends that the doctor chooses the act which has the best “worst-case” scenario.

### 3.2 Two Counterexamples to the Imprecision Model of Higher-order Uncertainty

While the proposal recommends the right course of action in the doctor-case, we think that Henderson’s imprecise probabilities view is unattractive as a framework for representing how rational agents respond to higher-order defeat. In this section we will argue the view is vulnerable to counterexamples.

It seems that imprecise probabilities are neither necessary nor sufficient for higher-order uncertainty. Consider first a case illustrating that imprecise probabilities are not necessary for higher-order uncertainty.

**Statistics Exercises:** Adam studies genetics and is trying to figure out whether a given woman carries a certain X-linked recessive gene given that she has two healthy male siblings. After thinking carefully, he comes to the conclusion that the probability of the woman carrying the gene is 10%. However, Adam is well aware that these types of exercises are not his strong suit. In fact, he is quite uncertain that the probability is 10%.

How should we represent Adam’s doxastic state? He is quite uncertain in his own reasoning. Why should this lack of confidence be represented by a credal set clustered around 10%? Adam would not be too surprised to learn that his result is mistaken. But he has no reason to think that values close to 10% are more likely to be rational

---

40%  $(0.6(0.9)+0.4(0)=0.54)$ . To get the upper bound you multiply with the upper bound of the fully vacuous state instead  $(0.6(0.9)+0.4(1)=0.94)$ . See Henderson (2022: n10) for the details.

than values quite far from 10%. In fact, it seems that he should not be any more certain that the rational value is close to 10% than that it is far from 10%. Friends of the imprecision model might suggest that Adam should adopt a very broad credal set, perhaps spanning the entire interval. But that does not seem right either. Given his calculations, Adam has most reason to think that the probability is 10%, and no reason to think that it is any other particular probability. So a very broad credal set does not seem to capture Adam's doxastic state.

Attributing a precise credence to Adam is, of course, the orthodox Bayesian view. And for good reason: Adam assigning a 0.1 credence to the proposition that the woman carries the gene reflects that he thinks the probability is 10%. Despite Adam's precise credence, he should be uncertain whether that credence is rational. The question, then, is how Adam's uncertainty is reflected in his doxastic state. As we outlined in Sect. 2, the uncertainty can be reflected in the low resiliency of his 0.1 credence. Adam should be highly disposed to change his 0.1 credence that the woman carries the gene. If this is right, it is possible for Adam to assign a precise credence while being uncertain about his own rationality.

Consider next a counterexample to the sufficiency of imprecise probabilities for higher-order uncertainty.<sup>8</sup>

**The Supercomputer:** Beth the meteorologist is trying to figure out the probability that it will rain tomorrow. Luckily, the institute she works at has just acquired a new supercomputer Deep Thought which has revolutionized her field. Just give Deep Thought your data, and the computer will give you an extremely reliable estimate of the probability of rain. Beth plugs her data in and is given the result that the probability of rain is within the interval [0.619, 0.773]. Having experienced the previous reliability of Deep Thought, Beth trusts the estimate.

Beth has good reason to be highly confident that the imprecise probability  $P[0.619, 0.773]$  represents the probability of rain given her evidence. Her credence should arguably be imprecise, yet she can rationally be highly confident in her own credal state. It thus seems that imprecise probabilities do not always indicate higher-order uncertainty.

The resilience framework can capture such cases by saying that Beth should adopt the interval of probabilities recommended by Deep Thought with a high degree of resilience. It should take a lot of evidence to move Beth away from this credal set. Not so for Adam. His 0.1 credence should be unresilient: it should be relatively easy for evidence to change his mind about whether the woman carries the recessive gene.

Our general diagnosis of these cases is that imprecise probabilities map the wrong thing. Credal imprecision does not reflect the kind of doubt which higher-order defeat should induce. The problem is that we cannot capture the relevant kind of uncertainty

<sup>8</sup> A reviewer points out that Henderson does not explicitly claim that imprecision always indicates a lack of conviction. While she does say that "the degree of precision of the credal set represents our 'conviction' about our first-order attitudes" (p. 521) she might—as we think she should—reject that imprecision is sufficient for higher-order doubt. But then one might wonder what the relevant difference is between a higher-order confident and higher-order uncertain imprecise credence.

merely as uncertainty about whether *the woman carries the recessive gene* or whether *it will rain tomorrow*. We need to encode uncertainty whether 10% is the rational answer and we need to encode confidence that the probability of rain tomorrow is in the given interval.

Notice that the imprecision model creates a tight link between what probability one assigns to a proposition and one's confidence in that assignment. Since Adam thinks the answer to the problem is 10%, his uncertainty about this judgment should be modeled around 10%. These cases create trouble for the imprecision framework by showing that such judgments can come apart from one's higher-order confidence.

Defenders of the imprecision model have a natural response here which is worth considering, even though we ultimately do not think it presents a way out for them. They might argue that cases like these really involve giving credence to propositions with a probabilistic content. Then we can encode the degree of conviction as a regular (imprecise) credence in a probabilistic proposition irrespective of the probability figuring as the content of the relevant (imprecise) credence. Take Adam. When reflecting on his own mathematical shortcomings, he begins to doubt that his answer is rational. What is the object of his belief? A natural candidate is the proposition that *the probability of the woman carrying the recessive gene is 10%*. If so, one might respond that this is the proposition Adam should assign an imprecise credence to. His evidence is compatible with various probabilities in that probabilistic proposition and so he should assign a relatively imprecise probability to it, say  $[0.25, 0.65]$ . The difference between the upper and lower bounds being relatively large (namely 0.4), represents Adam's lack of conviction. Notice, then, that the imprecise probability Adam should assign the probabilistic proposition does *not* have to be centered around 10%. This suggests that by introducing probabilistic propositions, the imprecision model can avoid the tight link between Adam's 10% judgment and his higher-order doubt.

Even so, introducing a probabilistic proposition does not make the non-probabilistic proposition *the woman carries the recessive gene* go away. This is still a proposition Adam can give credence. After all, it is the proposition he was entertaining while doing the exercise. Further, the credence Adam gives to the non-probabilistic proposition should surely somehow be constrained by the credence he gives the probabilistic counterpart. In particular, it seems that the degree of conviction (that is, the credal span of 0.4) Adam assigns to the probabilistic proposition should constrain the degree of conviction he gives the non-probabilistic proposition. But what imprecise probability with a 0.4 span should Adam assign to *the woman carries the recessive gene*, such that it reflects that Adam takes "10%" to be his best guess? It was exactly the fact there is no plausible answer to that question which motivated appealing to probabilistic propositions in the first place. So for this strategy to get off the ground, we need a more general account of when and why ordinary non-probabilistic propositions take a back seat in the doxastic lives of people that find themselves in Adam's rather commonplace predicament. Opting to introduce probabilistic contents into one's models has to be motivated by some fairly general principles guiding when they are the appropriate objects of thought and how they constrain or altogether replace the relevant non-probabilistic propositions.

For this reason, we are not optimistic that Henderson can escape the counterexamples in any straightforward manner by appealing to probabilistic propositions. More

generally, it is not without difficulty to introduce probabilistic propositions in the hope that they can do any theoretical work which credences in regular propositions cannot do.

## 4 Resilience and Decision-making

While we think that Henderson's own view faces counterexamples, we also think she draws attention to a crucial issue. If the resilience view is to be viable, it needs to address how practical decisions should be influenced by higher-order defeat. So far, we have argued that higher-order defeat, when non-lopsided, can require one to lower one's credal resilience while permitting one to maintain one's credal level. How and why are rational decisions influenced by this loss of resilience? We see two different ways in which the resilience-level of a credence can affect rational decision-making.

### 4.1 Resilience and Categorical Belief

When Bayesians theorize about rational decision-making, they do so in terms of credences or degrees of belief. The folk, however, tend to think about rational decision-making in terms of *categorical belief* or *knowledge*, which, on most accounts, requires categorical belief, and there is a rich philosophical tradition premised on this approach (Gerken, 2017). For example, a recently popular view holds that it is permissible to treat  $p$  as a practical reason (and thus act as if  $p$  is true) if and only if one *knows* that  $p$ , which implies in turn that it is permissible to treat  $p$  as a practical reason only if one *rationally believes* that  $p$ , insofar as rational belief is necessary for knowledge (Stanley & Hawthorne, 2008).<sup>9</sup>

What is the relationship between the credal and categorical ways of thinking about rational decision-making? That depends on how one should understand the relationship between credence and categorical belief. The issue here is complex, and as yet unresolved, but at least one promising way of understanding the relationship is helpful for our present purpose. This is the "stability theory" of belief, developed by Leitgeb (2014, 2015, 2017), according to which a high credence in  $p$  is necessary but not sufficient for categorical belief that  $p$ . To qualify as a belief, the high credence must also be sufficiently "stable." Stability is what enables belief to "play its characteristic functional role in decision-making, reasoning, and asserting [...] in the course of processes such as perception, supposition and communication" (Leitgeb, 2015: 146). He shows that a certain kind of stability is a necessary condition for rational categorical belief (see Theorem 1 in Leitgeb, 2014: 140).

On Leitgeb's account, stability is a matter of the credence being preserved under the supposition of new evidence or information (2015: 156). In particular, a credence in  $p$  is sufficiently stable to warrant categorical belief that  $p$  only if the credence in  $p$  is what he calls P-stable: the credence in  $p$  remains greater than 0.5 under the sup-

<sup>9</sup> For other representatives of the knowledge-based approach, see e.g. Levi (1980); Williamson (2000); Stanley & Hawthorne (2008); and Weisberg (2020). For arguments that categorical belief plays a key role in decision-making, see e.g. Ross & Schroeder (2014); Holton (2014); and Jackson (2019, 2020).

position of any proposition  $q$  which is consistent with what you categorically believe. These propositions count as serious possibilities in the given context of reasoning (Leitgeb, 2014: 138–139).<sup>10</sup>

We can then identify a specific threshold for when loss of resilience entails loss of categorical belief. Suppose you start with a high credence in  $p$  which amounts to categorical belief that  $p$ . Then that credence will be P-stable. In order for a loss of resilience to entail loss of categorical belief in  $p$ , you should lose resilience such that your conditional credence in  $p$  given at least one serious possibility  $q$  is less than 0.5.

This theory allows that in cases of non-lopsided higher-order defeat, it can be rationally permitted to maintain one's high level of credence in  $p$ , while abandoning belief that  $p$ . This is because in such cases, one can be required to lower the resilience of one's credence to a level lower than that required for categorical belief. This, we take it, is a plausible upshot.

Consider, for example, the case of doubtful Ava. It seems plausible that Ava may maintain a credence of 0.7 that the next president will be a Democrat, despite her strong doubt that this is indeed the rational level of credence. But even if this level of credence satisfies the required level for categorical belief that the next president will be a Democrat, it seems clear that, in light of her doubts, Ava should *not* believe that the next president will be a Democrat. Rather she should suspend judgment on the matter.<sup>11</sup> The stability theory, coupled with the resilience account of higher-order defeat, can explain this intuition: while Ava's credence level is consistent with belief, its degree of resilience is not. This is so since there is some proposition in Ava's context of reasoning which by Ava's lights counts as a serious possibility, such that Ava thinks it is more likely than not that the next president will not be a Democrat conditional on that proposition. For example, in light of her doubt, it seems that Ava should be significantly swayed by counterevidence like being told by a trusted friend that the evidence points towards a Republican president. Said differently, conditional on such testimony, Ava should give it less than 0.5 credence that the next president will be a Democrat. This means that her high credence that the next president will be a Democrat fails to be P-stable and, in turn, that it fails to count as a categorical belief.

Return now to the question of rational decision-making. If rational belief that  $p$  is required for the permissibility of acting as if  $p$ , and a suitably high degree of resilience is required for a credence to qualify as belief, it seems that a suitably high degree of resilience is also required for the permissibility of acting as if  $p$ . If that is the case, higher-order defeat of a credence that  $p$  can undermine the permissibility of acting as if  $p$  by defeating the resilience of that credence.

## 4.2 Resilience and Decision-making Under Ignorance

Not everyone will be attracted to the above solution, since it rests on controversial assumptions about the relationship between belief and credence, and the role of categorical belief in rational decision-making. However, there is another way in which

<sup>10</sup> See also Leitgeb (2015: § 3) for further details on how to determine the set of serious possibilities.

<sup>11</sup> Carr (2020) likewise argues that higher-order evidence leads to normative uncertainty which in turn requires suspension of judgment.

resilience can be relevant to rational decision-making, which does not rely on such assumptions.

Decision theorists standardly distinguish between decisions under *certainty*, which are made when one knows what the outcome of each alternative act will be; decisions under *risk*, which are made when one does not know the outcome of each act but can assign a probability to each outcome; and decisions under *ignorance*, which are made when one is unable to assign reliable probabilities to the outcomes.

It seems quite plausible that higher-order defeat could create a situation where one should make decisions as if under ignorance. In particular, this is plausible in cases of non-lopsided higher-order defeat, where the higher-order evidence undermines the rationality of one's level of credence without indicating that another credence is more rational. In such situations, the higher-order defeat should make one less sure that one's credence is rational, without letting one surmise where one should move it to. One is rationally stuck with an uncertain credence that cannot be relied upon.<sup>12</sup> The resilience theory understands this as a readiness to change one's credence in response to new evidence. And if one expects one's credence to change in the face of new evidence, one should not rely on it for making decisions.

So given sufficient higher-order defeat, one should make decisions as if under ignorance. Various principles have been proposed that can be relied upon when making decisions under ignorance. Least controversially, when one act leads to better outcomes than all alternative acts no matter what state of nature turns out to be true, and so *dominates* all alternatives, that action ought to be chosen. If no act or sets of acts dominates the alternative, such as in Henderson's case of the sleep deprived doctor, some additional criterion is needed. The most cautious decision principle to rely on such situations is the *maximin* rule (Milnor, 1954: 50):

**Maximin rule:** Choose an act where the minimum possible value is maximized.

According to this rule, you should choose an act with the best possible worst-case scenario. Applying the rule to the decision represented in Table 1 would yield postponing treatment. If postponing is not an option, that is, if the doctor is forced to choose between A and B, then the rule still offers the sensible recommendation of administering drug B, since the worst possible outcome given B is better than the worst possible outcome given A. The Maximin rule thus yields the decisions recommended by Henderson.

Despite the dissimilarity between Henderson's decision procedure and the one we advocate here, a similarity is worth pointing out. Recall, according to the imprecision

<sup>12</sup> Carr (2020) has a parallel suggestion. She argues that if your evidence is imprecise, your doxastic state ought to exhibit normative uncertainty. And this uncertainty consists partly in uncertainty about how to resolve the uncertainty. So there's no straightforward way of revising one's credences to accommodate the normative uncertainty. Rather, she suggests, you need to resort to some form of principle which can help guide your actions. The only principle she considers is "the most conservative one", namely the principle which says to ignore your normative uncertainty when making decisions. This move amounts to biting the bullet of Henderson's objection: the doctor should proceed with administering the drug she thought was best prior to becoming higher-order uncertain. However, as we'll now argue, we are not forced to opt for the most "conservative" principle. There are plausible principles available to us from the decision theory of how to act under ignorance.

model, higher-order doubt serves to broaden your credal set. During practical deliberation, your credal set will now yield an expected utility for every probability included in your credal set. The imprecise model then needs to appeal to some principle to determine which expected utility should guide decision-making. As we saw, Henderson recommended letting the minimum expectation guide action: perform the action with the least-bad expected outcome. This last step in the decision procedure is the impreciser's counterpart of Maximin. Indeed, Gärdenfors and Sahlin (1982: 371) explicitly call this step of their procedure "The maximin criterion for expected utilities."

We do not want to overstate the similarity between the two proposals. After all, Henderson's proposal still suggests maximizing certain expected utilities—namely the minimum expected utilities—while ours makes no such suggestion, since we think you should not rely on your own expectation given sufficient higher-order defeat. While we both appeal to minimum-maximizing principles, the principles are not identical.

However, there is still a commonality in how we arrive at our respective principles. Neither follow from utility maximization, for example. And so we do not have a principled way of picking out the correct principle amongst the various decision principles available in the literature. It does not follow from anything built into the imprecise framework that one should place special emphasis on avoiding the worst-case scenario. Indeed, Henderson would presumably allow that in some cases of credal imprecision, you are not required to let the minimum expectation guide action. For example, it seems that there will be cases where the best case scenario is sufficiently good to warrant a less risk-averse decision rule.<sup>13</sup> Nor is there anything built into the resilience framework which leads to Maximin.

However, when we judge that the doctor ought to postpone—and if that is not possible, that the doctor ought to administer drug B—it is plausible that our intuitive judgment implicitly relies on a principle like Maximin. While we cannot prove that Maximin is the correct decision principle in a given case, we can point to the fact that the principle is compatible with the general framework and that the principle explains our intuitive judgments. In other cases, Maximin will seem overly cautious. Further principles for decision-making under ignorance have been suggested to handle such cases (see Peterson (2013: Ch. 3) for a helpful overview). For example, if there is not much to lose but a lot to gain between choices where we can assign no probabilities, we tend to opt for optimistic decision principles. And rationally so. Better, then, to not build a principle like Maximin into the theory directly, but rather leave it context dependent which principle yields the correct result in a specific case. This is not the place, however, for a detailed review of the various decision principles on offer. The point is that rational decision-making is possible, and our present claim is that in cases of sufficiently detrimental non-lopsided higher-order defeat, one ought not to rely on the relevant credences when deciding what to do, but rather disregard them and act as if under ignorance.

<sup>13</sup> If the best-case scenario is sufficiently good and the worst-case is not so bad, it seems that rather than letting the minimum expectation guide action, one should let the maximum expectation influence one's decision. Decision principles like  $\Gamma$ -maximax (which uses the maximum expectation) and Hurwicz's rule (which combines the minimum and maximum expectation) have been proposed to allow for less conservative decision-making for imprecise probabilities. See Huntley et al. (2014) for a helpful overview of various decision principles available to the imprecise probability framework.

The present proposal leaves us with the question of when higher-order defeat is so detrimental that you should act as if under ignorance. Should you disregard your credences and adopt a decision principle like Maximin if you learn there is just a 1% chance that your reasoning is impaired by exhaustion? Presumably not. So one way or another, our proposal must be supplemented with some considerations about the degree of resilience required in order for it to be permissible to rely on your credences in decision-making. And indeed, *any* decision theory operating with a difference between cases where one's credences are reliable enough to guide decisions, and cases where they are not, are faced with the task of specifying some threshold of reliability. Although important, we will therefore leave this question for another occasion.<sup>14</sup>

## 5 Conclusion

We have sought to do two things in this paper. We first argued that Henderson's imprecise probabilities framework proves unsatisfactory when it comes to representing a certain kind of higher-order defeat. We then described two ways that loss of doxastic resilience can affect rational decision-making, thus answering Henderson's objection to the resilience framework for understanding higher-order defeat. First, if rational categorical belief that  $p$  is necessary for acting as if  $p$ , a high degree of resilience seems necessary for acting as if  $p$  as well, since, on a plausible way of linking credences and categorical belief, a high degree of resilience is necessary for categorical belief. Second, even in the absence of such a link between credences and categorical beliefs, sufficiently detrimental loss of resilience should lead you to act as if under ignorance. Fortunately, we have principles which can guide our decisions even in such cases, yielding the results recommended by Henderson.

**Acknowledgements** We wish to thank Darren Bradley, Robbie Williams, Dan Greco, and two anonymous reviewers for valuable feedback on earlier drafts of this paper.

## Declarations

**Conflicts of Interest** The authors report no conflicts of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

---

<sup>14</sup> We suspect that the right answer must appeal to the practical stakes, and to the cost and prospects of further inquiry. For discussion of the latter consideration, see e.g. Good (1967) and Dorst et al. (2021).

## References

- Bradley, R. (2017). *Decision theory with a Human Face*. Cambridge University Press.
- Carr, J. R. (2020). Imprecise evidence without Imprecise Credences. *Philosophical Studies*, 2735–2758. <https://doi.org/10.1007/s11098-019-01336-7>. 177.9.
- Christensen, D. (2007). Epistemology of disagreement: The Good News. *Philosophical Review*, 116(2), 187–217. <https://doi.org/10.1215/00318108-2006-035>
- Christensen, D. (2010a). Higher Order Evidence. *Philosophy and Phenomenological Research*, 81.1, pp. 185–215. <https://doi.org/10.1111/j.1933-1592.2010.00366.x>
- Christensen, D. (2010b). ‘Rational Reflection’. *Philosophical Perspectives*, 24.1, pp. 121–140. <https://doi.org/10.1111/j.1520-8583.2010.00187.x>
- Dorst, K., Levinstein, B., Salow, B., Husic, B., & Fitelson, B. (2021). Deference Done Better Philosophical Perspectives, 35(1): 99–150.
- Elga, A. (manuscript) ‘Lucky to Be Rational’. <https://philpapers.org/rec/ELGLTB>
- Feldman, R. (2009). ‘Evidentialism, Higher-Order Evidence, and Disagreement’. *Episteme*, 6.3, pp. 294–312. <https://doi.org/10.3366/e1742360009000720>
- Gärdenfors, P. and Nils-Eric Sahlin (1982) ‘Unreliable Probabilities, Risk Taking, and Decision Making’. *Synthese*, 53.3, pp. 361–386. <https://doi.org/10.1007/bf00486156>
- Gerken, M. (2017). *On Folk Epistemology: How we think and talk about knowledge*. Oxford University Press.
- Good, I. J. (1967). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17, 319–322.
- Hawthorne, John and Jason Stanley. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590. <https://doi.org/10.5840/jphil20081051022>
- Hedden, Brian and Kevin Dorst. (2022). (Almost) All Evidence is Higher-Order Evidence. *Analysis*, 82.3, pp. 417–425. <https://doi.org/10.1093/analys/anab081>
- Henderson, L. (2022). Higher-Order Evidence and Losing One’s Conviction. *Noûs*, 56.3, pp. 513–529. <https://doi.org/10.1111/nous.12367>
- Holton, R. (2014). Intention as a model for belief. *Rational and Social Agency: Essays on the philosophy of Michael Bratman*. Ed. By Manuel Vargas and Gideon Yaffe. Oxford University Press.
- Huntley, N., Hable, R., & Troffaes, M. C. M. (2014). ‘Decision making’. In: *Introduction to Imprecise Probabilities*. Ed. by T. Augustin, F. P. A. Coolen, G. de Cooman, M. C. M. Troffaes. Wiley, pp. 190–206.
- Jackson, E. (2019). Belief and Credence: Why the Attitude-Type Matters. *Philosophical Studies*, 176.9, pp. 2477–2496. <https://doi.org/10.1007/s11098-018-1136-1>
- Jackson, E. (2020). The Relationship Between Belief and Credence. *Philosophy Compass*, 15.6, pp. 1–13. <https://doi.org/10.1111/phc3.12668>
- Jeffrey, R. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.
- Joyce, J. M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24.1, pp. 281–323. doi: 10.1111 / j. 1520–8583.2010.00194.x.
- Joyce, J. M. (2012). Regret and Instability in Causal Decision Theory. *Synthese*, 187.1, pp. 123–145. <https://doi.org/10.1007/s11229-011-0022-6>
- Keynes, J. M. (1921). *A treatise on Probability*. Dover Publications.
- Lasonen-Aarnio, M. (2013). Disagreement and evidential attenuation. *Noûs*, 767–794. <https://doi.org/10.1111/nous.12050>. 47.4.
- Leitgeb, H. (2014). The Stability Theory of Belief. *Philosophical Review*, 123(2), 131–171. <https://doi.org/10.1215/00318108-2400575>
- Leitgeb, H. (2015). The Humean Thesis on Belief. *Aristotelian Society Supplementary Volume*, 89.1, pp. 143–185. <https://doi.org/10.1111/j.1467-8349.2015.00248.x>
- Leitgeb, H. (2017). *The Stability of Belief: How rational belief coheres with probability*. Oxford University Press.
- Levi, I. (1980). *The enterprise of knowledge: An essay on knowledge, Credal Probability, and Chance*. MIT Press.
- Milnor, J. W. (1954). ‘Games Against Nature’. In: *Decision Processes*. Ed. by R.M. Thrall, C.H. Coombs, and R.L. Davis. New York: Wiley, pp. 49–60.
- Peirce, C. S. (1931). *Collected papers*. Harvard University Press.
- Peterson, M. (2009). *An introduction to decision theory*. Cambridge University Press.

- Ross, J., & Mark Schroeder (2014). Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88.2, pp. 259–288. <https://doi.org/10.1111/j.1933-1592.2011.00552.x>
- Skipper, M. (2020). Does Rationality Demand Higher-Order Certainty?. *Synthese*, 198.12, pp. 11561–11585. <https://doi.org/10.1007/s11229-020-02814-w>
- Skyrms, B. (1977). Resiliency, propensities, and Causal Necessity. *Journal of Philosophy*, 74(11), 704–713. <https://doi.org/10.2307/2025774>
- Skyrms, B. (1980). *Causal necessity: A pragmatic investigation of the necessity of laws*. Yale University Press.
- Steglich-Petersen, A. (2019). Higher-order defeat and Doxastic Resilience. *Higher-order evidence: New essays*. Ed. By Mattias Skipper and Asbjørn Steglich-Petersen. Oxford University Press.
- Svenson, O. (1981). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica*, 47(2), 143–148.
- Weisberg, J. (2020). Belief in psyontology. *Philosophers' Imprint*, 20.11.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.