



How to act on what you know

Roman Heil¹ 

Received: 13 October 2023 / Accepted: 6 May 2024
© The Author(s) 2024

Abstract

That we may rely on our knowledge seems like a platitude. Yet, the view that knowledge is sufficient for permissible reliance faces a major challenge: when much hangs on whether we know, relying on our knowledge seems to license irrational action. Unfortunately, extant proposals to meet this challenge (Hawthorne & Stanley, 2008; Williamson, 2005a; Schulz, 2017, 2021b) either fail to make the correct predictions about high-stakes cases or, as I will argue, face a substantial objection. In this paper, I will offer two novel proposals for defending the sufficiency of knowledge for permissible reliance. My favoured proposal, which I call Flexible Fallibilism, is based on the idea that stakes can affect how we rely on our knowledge and how our knowledge determines what is rational for us to do. Besides making the correct predictions about high-stakes cases, I will argue that Flexible Fallibilism also provides us with a knowledge-based account of when to simplify our practical reasoning. Afterwards, I will present another proposal, which I call Dual Infallibilism, that combines two recent claims by Moss (2013, 2018) and Jackson (2019a) about credal knowledge and the metaphysics of doxastic states. While the offered proposals share various virtues, I will argue that we ultimately should prefer Flexible Fallibilism.

Keywords Knowledge norms · Practical reasoning · Decision theory · Reliance · Fallibilism

1 Introduction

That we may rely on our knowledge in action seems like a platitude. According to a prominent view, this platitude explains why knowledge is distinctively valuable: knowledge can guide action in a way that mere true belief cannot (Meno 97, a–c; Williamson, 2000, pp. 78–79, 101–102; Hawthorne & Stanley, 2008, p. 590; Hyman, 2010).¹ Knowledge talk also plays a prominent role in our assessment of practical

¹ For a comprehensive overview of the value of knowledge debate, see Pritchard and Carter (2018).

✉ Roman Heil
r.heil@em.uni-frankfurt.de

¹ Institute of Philosophy, Goethe University Frankfurt, Norbert-Wollheim-Platz 1, 60629 Frankfurt am Main, Germany

reasoning: we criticize those who don't rely on knowledge in practical reasoning, and can defend our own practical reasoning by pointing out that we only relied on knowledge. Various authors have argued that this is best explained by our practical reasoning being governed by a knowledge norm:²

Knowledge norm for practical reasoning (KPR) One may rely on a proposition p in practical reasoning iff one knows that p .³

Given the platitudinousness of the right-to-left direction of KPR, it comes as a surprise that it faces a substantial challenge from counterexamples: when much hangs on whether we know, relying on our knowledge seems to license irrational action.⁴ Consider the following case:

Jellybean You are participating in a psychological study intended to measure the effect of stress on memory. The researcher asks you questions about Roman history—a subject with which you are well acquainted. For every correct answer you give, the researcher will reward you with a jelly bean; for every incorrect answer, you are punished by an extremely painful electric shock. There is neither reward nor punishment for failing to give an answer. The first question is: when was Julius Caesar born? You are confident, though not absolutely certain, that the answer is 100 BC. You also know that, given that Caesar was born in 100 BC, the best thing to do is to provide this answer (i.e. this course of action will have the best consequences—you will be one jelly bean richer!). (Reed, 2010, pp. 228–229)

Jellybean and similar cases have been proposed as counterexamples to the right-to-left direction of KPR. Reed tells us that in Jellybean, you confidently know that Julius Caesar was born in 100 BC. However, the reward for getting the answer right is meager and the punishment for getting it wrong severe. Intuitively, risking a shock for your confident but not certain knowledge does not seem like the rational course of action. Generally speaking, if stakes are high, it is sometimes not rational to perform some salient action when we rely on what we know.⁵ This shows, critics claim, that we may not always rely on our knowledge in practical reasoning and hence that KPR must be false.

One response to counterexamples like this is to embrace Impurism about knowledge and deny that one knows the relevant proposition in Jellybean. According to

² See Hawthorne (2004), Williamson (2005b), Stanley (2005), Hawthorne and Stanley (2008), Fantl and McGrath (2009), Smithies (2012), Mehta (2016), Williamson (2017) Simion (2021) and Schulz (2021b). Not all of these authors believe that knowledge is both necessary and sufficient for permissible reliance.

³ Proposals differ on whether knowledge sets the standard for when one may treat a proposition as a *premise in one's practical reasoning* or as a *reason for action*. I will stay neutral on these matters and follow others in capturing the core commitment that is common to these proposals by employing the notion of *reliance* (Fantl & McGrath, 2019, p. 259; Fritz, 2022, p. 1193). Whatever the propositions we rely on in practical reasoning ultimately are, knowledge is sufficient for permissibly relying on them.

⁴ In what follows, I will only be concerned with the sufficiency of knowledge for permissible reliance.

⁵ I will use *high stakes* as a placeholder characterization of what drives our intuitions in cases like Jellybean. Anderson and Hawthorne (2019) have recently argued that it is hard to find a plausible characterization of high stakes that conclusively characterizes “our patterns of worry and concern” (2019, p. 255) that are distinctive of these cases. For the purpose of this paper, I will simply proceed with the uncontroversial assumption that there is such a pattern of worry and concern in cases like Jellybean, whatever a proper characterization of this pattern might look like in the end.

Impurism, whether a subject's true belief qualifies as knowledge partly depends on the subject's practical interests (Roerber, 2018). However, many reject the idea that whether one knows depends on what is at stake. Those who want to stick to epistemological orthodoxy hence need to opt for another response to high-stakes cases.

The following observation suggests a way of moving forward: High-stakes cases, in which one knows but in which it is intuitively not rational to act in some salient way, do not constitute counterexamples to KPR by themselves.⁶ The reason for this is that KPR only tells us when we may *rely on* propositions in practical reasoning, not how to *act* in the light of the propositions we may rely on. To assess whether high-stakes cases like Jellybean constitute a challenge to KPR, there needs to be a view in place to determine which actions are rational for us to perform in the light of the propositions we may rely on. In other words, proponents of KPR need a supplementary *knowledge-based decision theory* to tell us what is rational for us to do, given what we know. If, for instance, such a view were to tell us that answering that Caesar was born in 100 BC is the rational thing to do in Jellybean, then cases like Jellybean would indeed be (indirect) counterexamples to KPR. Hence, the prospects of defending KPR very much depends on there being an independently plausible knowledge-based decision theory that delivers the correct verdict about high-stakes cases like Jellybean.

One could of course deny that there is any relation between what we may rely on in practical reasoning and what is rational for us to do. While Jellybean would be immediately taken care of, I think this response is unattractive. First, it is puzzling: why think there is a norm for practical reasoning in the first place if compliance with this norm has no downstream effect on rational action? This response denies the plausible view that the rationality of practical reasoning is linked to the rationality of action. As Schulz (2017, p. 463) aptly puts it, “[i]f decisions should be based on what one knows, knowledge should play a fundamental role in decision theory.” Second, this response undermines any prospects of providing a unified knowledge-first view of practical rationality, a project many proponents of KPR are attracted to (Goldschmidt forthcoming, 4). Indeed, it is natural to think of KPR and knowledge-based decision theory as complementary constraints on ex post rational action. Knowledge-based decision theory is concerned with the question of what is ex ante rational for us to do. Plausibly, an action is ex post rational only if it was ex ante rational for one to perform. KPR is concerned with the rationality of practical reasoning. Plausibly, an action is ex post rational only if it is based on rational practical reasoning.⁷ Putting this together, we get the following, knowledge-based view: An action is ex post rational only if it was ex ante rational for one to perform, given what one knows (knowledge-based decision theory), and performed by relying on one's knowledge (KPR). While I think that the resulting unified view is theoretically attractive and worth pursuing, I want to note that knowledge-based decision theory has also been proposed and motivated independently of KPR (Moss, 2018; Dutant, forthcoming; Goldschmidt, forthcoming). For those not wanting to commit to KPR, my paper can be read more narrowly as advancing two

⁶ See Neta (2009), Ichikawa (2012), Jackson (2012), Weisberg (2013), Locke (2014), Roerber (2018) and Beddor (2020).

⁷ Compliance with KPR is only a necessary condition for ex post rational action. For discussion about what's missing for sufficiency, see, e.g., Lord (2018, ch. 5).

novel knowledge-based decision theories that, in contrast to extant proposals, provide us with the correct verdict about what is rational to do in high stakes cases.⁸

In this paper, I will first argue that extant proposals either face a substantial objection or result in high-stakes cases being counterexamples to KPR. I will then propose and compare two novel knowledge-based decision theories that avoid these problems, get the verdict right about cases like Jellybean, provide us with a view about when to simplify our practical reasoning and ultimately vindicate the platitude that we may rely on what we know.

The paper is structured as follows. In Sect. 1, I will reject the kind of knowledge-based decision theory that is commonly assumed on behalf of KPR. According to this view, which I will call *Basic Infallibilism*, an action is rational if it maximizes expected utility, *conditional on the totality of what one knows*. Basic Infallibilism assigns probability 1 to our knowledge and thus gives no probabilistic weight to the dire outcome in Jellybean. Hence, it tells us that answering maximizes expected utility in Jellybean, which is clearly at odds with our intuitions. In Sect. 2, I will discuss what I call *Sophisticated Infallibilism* (Williamson, 2005a; Schulz, 2017, 2021b), the view according to which, if stakes are high, we should rely and conditionalize only on an epistemically particularly safe subset of our knowledge. While Sophisticated Infallibilism makes the right predictions about Jellybean, I will show that it faces a substantial objection.

After having made my negative case against extant proposals, I will propose two novel knowledge-based decision theories. In Sect. 3, I will lay the groundwork for the view I ultimately favour. I will argue that our ordinary notion of reliance is much more fallibilist than has hitherto been acknowledged, an insight which has important ramifications for how knowledge connects to rational action. I will then employ a generalized form of conditionalization to develop a knowledge-based decision theory that implements these lessons. The resulting proposal, which I call Flexible Fallibilism, not only makes the right predictions in Jellybean but, as I will show in Sect. 4, also provides us with a novel knowledge-based view about how to simplify our practical reasoning. In Sect. 5, I will compare Flexible Fallibilism to yet another proposal, which I call Dual Infallibilism, that combines two recent claims by Jackson (2019a) and Moss (2013, 2018). While Dual Infallibilism also makes the right predictions about Jellybean and offers a theory of how to simplify our practical reasoning, I will argue that Flexible Fallibilism has various advantages over it and is thus preferable.

⁸ Fassio and Gao (2021) have recently raised various objections against knowledge-based decision theory (and perspectivism more generally). While I will solely be concerned with the challenge from high stakes in this paper, I briefly want to comment on one of the worries they raise: Gao and Fassio conceive of KPR and knowledge-based decision theory as competing views that are differently motivated and offer conflicting recommendations. Regarding the first point, they argue, for instance, that while KPR can be motivated by appealing to data about praise and criticism of practical reasoning and (ex post) rational action (Hawthorne & Stanley, 2008) knowledge-based decision theory in isolation cannot, with the worry being that the latter is thus ill-motivated (but see Goldschmidt forthcoming). They also consider and reject a pluralist proposal, according to which the two views serve different purposes and address different decision problems (Fassio & Gao 2021, p. 7055). However, they do not consider the view sketched above, namely that KPR and knowledge-based decision theory constitute tightly connected, serial constraints on ex post rational action. Although I have to leave addressing their objections in detail for another occasion, one can already see how treating KPR and knowledge-based decision theory as integrative components of a unified, knowledge-based theory of rational action deals with the mentioned motivational problems.

2 Basic infallibilism

To defend KPR, a decision theory has to be offered that is knowledge-based. A decision-theory is knowledge-based if one’s knowledge plays a crucial role in determining what is rational for one to do. In what follows, I will focus on the epistemic role knowledge might play in this regard.

The knowledge-based decision theory that is sympathetically discussed by Hawthorne and Stanley (2008) and commonly assumed in the literature says that an action is rational for one to perform if it *maximizes expected utility, conditional on what one knows*.⁹ On this picture, knowledge plays a role in determining what is rational for one to do because the probability function that we use to calculate expected utilities is conditionalized on one’s knowledge. Since known propositions receive probability 1 if conditionalized on the totality of one’s knowledge K (i.e., if $p \subseteq K$, then $P(p|K) = 1$), this view is most naturally classified as a form of infallibilism about knowledge (Brown, 2011, p. 159). Let’s call this version of knowledge-based decision theory *Basic Infallibilism*.¹⁰

The problem is that Basic Infallibilism leads right into trouble for KPR. Consider the following decision matrix that represents our choice in Jellybean between answering (A) and not answering ($\sim A$), where p is the proposition that Caesar was born in 100 BC:

	p	$\sim p$
A	Win a jellybean (1)	Electric shock (−1000)
$\sim A$	0	0

In case Caesar was not born in 100 BC ($\sim p$), answering leads to an outcome with substantial disutility, an extremely painful electric shock (−1000). Yet, according to Basic Infallibilism, its disutility should be ignored in determining what is rational for us to do. After all, if one knows that p , $\sim p$ receives a conditional probability of 0 (i.e. $P(\sim p|K) = 0$), so outcomes associated with $\sim p$ do not receive any probabilistic weight at all in the expected utility calculation. According to Basic Infallibilism, answering would maximize expected utility and thus be the rational thing to do.¹¹ Yet, we clearly have the intuition in Jellybean that risking an extremely painful electric shock is not something that is rational for us to do. If stakes are high, we shouldn’t ignore the possibility that what seems like an ordinary piece of knowledge could be false after all.¹²

The argument now proceeds as follows: If we are committed to Basic Infallibilism, but this leads to the just noted implausible consequences, then we have ample reason to

⁹ Goldschmidt (forthcoming) has recently provided a Savage-style representation theorem for this view. The view’s implausible consequence discussed below is due to his K7 axiom (ibid., 18–19).

¹⁰ Weatherson’s (2012) proposal is also naturally classified as a form of Basic Fallibilism, but adds on top that what one knows determines what possibilities should be included in the decision table, namely those one doesn’t know not to obtain (ibid., 82; Fassio & Gao, 2021, fn. 4).

¹¹ $EU(A) = P(p|K) \times 1 + P(\sim p|K) \times (-1000) = 1 \times 1 + 0 \times (-1000) = 1$; $EU(\sim A) = 0$; $EU(A) > EU(\sim A)$

¹² For further discussion, see, e.g., Williamson (2005c), Greco (2013) and Schulz (2017).

reject KPR and to give up on the platitude that knowledge is sufficient for permissible reliance. To avoid this, we should reject Basic Infallibilism and look elsewhere for a supplementary knowledge-based decision theory.

Before discussing more promising proposals, it is instructive to investigate our intuitions about high-stakes cases like Jellybean more closely. Williamson (2005c, pp. 480–481) discusses a high-stakes case where you get to decide whether to bet on a complex logical falsehood $\sim q$, standing to win a carrot if q is false and to be horribly tortured to death if q is true. He writes that

few reasonable humans would accept the bet, even if they had worked out the truth-table. The penalty for a small computational error is just too high. Reasonable humans have *cognitive habits for managing their own fallibility* which the probability calculus makes no attempt to reflect. (my emphasis)

I think Williamson is exactly right in tracing our intuitions about high-stakes cases to our fallibilist habits.¹³ Humans often err, and yet we have to act. We are painfully aware of our own fallibility and have developed cognitive habits that allow us to *price in* our fallibility when making decisions. These cognitive habits are plausibly the source of our intuitions in high-stakes cases. When a lot hangs on whether we know, relying on what we take ourselves to know with an infallibilist spirit is rightfully not something we are comfortable doing, at least not without double-checking, second thoughts or further assurances.

I think our fallibilist intuitions affect (and should affect) how we think about what is rational for us to do. Whether an action is rational should reflect whether much can be lost if the agent is incorrect about the things she takes herself to know. A desideratum for any knowledge-based decision theory then is that it needs to take human fallibility into account.¹⁴

3 Sophisticated infallibilism

In what follows, I will discuss a more refined infallibilist proposal, which I will call *Sophisticated Infallibilism*, that performs better when it comes to satisfying the desideratum from the last section.¹⁵ I will first introduce the view in general and then note two variants afterwards.

Generally put, Sophisticated Infallibilism says that in high-stakes cases, we should not rely and conditionalize on the totality of what we know, but only on a certain *epistemically privileged subset* of our knowledge. If stakes are high (such as in Jellybean), we should rely and conditionalize on the subset of *epistemically particularly secure* knowledge to protect the goods at stake. Depending on one's favourite theory of epistemic justification, only knowledge that is particularly well warranted, well supported by one's evidence, reliable, sensitive or safe—surpassing what is needed to qualify as

¹³ For similar considerations, see Ross and Schroeder (2014, pp. 266–267). For a critical discussion, see Fassio (2018).

¹⁴ See also Schulz, (2017, p. 468) on this.

¹⁵ Proponents of Sophisticated Infallibilism are Williamson (2005a) and Schulz (2017, 2021b).

mere knowledge—is particularly secure in this sense.¹⁶ If stakes are low, by contrast, our ordinary, moderately secure knowledge is fit to be relied and conditionalized on.

Let's suppose that we can broadly divide what is at stake into different levels (Schulz, 2017, p. 469). So, for instance, if a carrot is at stake, we are at the lowest level of stakes. If one's house, job, physical well-being or the life of a loved one is at stake (ibid.), we typically will be at a high level of stakes. Depending on one's utility function, the different levels in between can be populated accordingly. We are now ready to formulate a version of KPR that captures the commitments of Sophisticated Infallibilism:

KPR_{SI} One may rely on a proposition p in practical reasoning involving n th-level stakes iff one knows that p with n th-level of epistemic security.

Here's how Sophisticated Infallibilism deals with Jellybean: Since Jellybean is a high-stakes decision situation, we may only rely and conditionalize on the subset of knowledge that comes with a correspondingly high level of epistemic security. Sophisticated infallibilists assume that particularly secure knowledge will only be rather sparsely available. Much of our ordinary knowledge will not be so secure and thus not part of the privileged subset of knowledge that is conditionalized on in high stakes cases. Hence, conditional on what we *do* know with the needed level of security, the proposition that Caesar was born in 100 BC will plausibly receive a probability of less than 1. But now, conversely, the outcomes associated with its falsity receive probabilistic weight in one's expected utility calculation. Since one of these outcomes, the outcome of answering wrongly if Caesar was not born in 100 BC, is decidedly terrible, giving it some probabilistic weight would be enough to make *not answering* the action that maximizes expected utility and thus the rational thing to do. Since Sophisticated Infallibilism gives probabilistic weight to the possibility that the proposition one takes oneself to know could be false, it takes into account our fallibilist intuitions about Jellybean. Relying and conditionalizing only on particularly secure knowledge is what those with "cognitive habits for managing their own fallibility" can be expected to do if stakes are high.

With Sophisticated Infallibilism in general introduced, let's look at some more specific proposals to fill in the blanks. First, what subset of knowledge is epistemically particularly secure? According to Williamson (2005a, p. 232) and Schulz (2017, pp. 467–468), the knowledge that is particularly secure is our *higher-order knowledge*. Why should we think that higher-order knowledge is epistemically particularly secure? One answer is that higher-order knowledge is *safer* than first-order knowledge (Schulz, 2017, p. 467).¹⁷ In his more recent proposal, Schulz (2021b, p. 8081) directly cashes out his candidate for particularly secure knowledge—*knowledge of higher strength*—in terms of safety: one piece of knowledge is stronger than another piece of knowledge in this sense if it is safer. On both proposals, the subset of knowledge sophisticated infallibilists recommend relying and conditionalizing on in high-stakes cases is epistemically particularly secure.

¹⁶ I will use the notion of *epistemic security* as a placeholder to stay neutral on the different, more specific proposals sophisticated infallibilists have put forward. I will elaborate these proposals in more detail below.

¹⁷ Influential proponents of the safety condition include Sainsbury (1997), Sosa (1999), Williamson (2000) and Pritchard (2005).

Which version of Sophisticated Infallibilism should we prefer? There have been many objections to cashing out Sophisticated Infallibilism in terms of higher-order knowledge.¹⁸ To note just one worry, higher-order Sophisticated Infallibilism tells us that we have to engage in higher-order reasoning to acquire higher-order knowledge (Gerken, 2011, p. 539; Gao, 2019a, pp. 99–101). But it is quite puzzling why this is something we ought to do. Higher-order reasoning does not increase first-order safety and if a belief is first-order safe enough for higher-order knowledge, why not rely on it directly instead of investing cognitive resources to acquire the corresponding higher-order knowledge? These and other worries have led Schulz (2021b) to cash out Sophisticated Infallibilism directly in terms of knowledge of higher strength.

In what follows, I will argue that there is a structural objection to Sophisticated Infallibilism that can't be avoided by opting for knowledge of higher strength. The objection shows that we should reject Sophisticated Infallibilism, despite its correct verdicts in high-stakes cases.

4 The indiscriminate security objection

The problem is that Sophisticated Infallibilism's demand for epistemic security is too indiscriminate, resulting in intuitively irrational choices and implausible predictions about how stakes matter for rational action. Consider the following case.

Camping Trip Ronda and her young daughter Daphne are going camping. Their backpacks were already packed by Ronda last night. Since she was quite tired after work, she just about knows that she packed all the things needed for a great trip. By contrast, after a full night's sleep, she confidently knows that the cat feeder is turned on, which takes care of their beloved house cat by dispensing water and food while they are gone. Ronda has to decide between two campings spots: the first camping spot C1 is surrounded by beautiful scenery and involves a few hours hike to get there. The second camping spot C2 is much less beautiful, but only 20min away from their home. Ronda goes through her mental checklist: "I packed some ground coffee (l_1), Daphne's plush-deer (l_2), ..., (l_{100}). The cat feeder is also turned on (h_1)." She then contemplates her choice: The scenery is much more beautiful at C1 than at C2 and the hike is going to be fun. Going to C2 would allow her to go to her yearly dental appointment and, on her way back, check on the cat and pick up missing things. However, she confidently knows that the feeder is turned on and the appointment can easily be postponed. Although she less confidently knows that she packed all things yesterday, missing some of them would not affect their trip significantly. Ronda concludes her reasoning by deciding to go to C1.

Let's note a few things about the case: First, stakes intuitively are high since their cat's life is at stake in Ronda's decision: if they go to C1 and the cat feeder is not turned on, then their cat will die as a result. Second, note that not all propositions contribute equally to what is at stake: while being wrong about h_1 would have a significant negative impact on the outcomes of her choice, being wrong about some of

¹⁸ See in particular Gao (2019a). For a helpful summary, see Schulz (2021b, pp. 8077–8079).

the propositions l_1, l_2, \dots, l_{100} would only have a slight negative impact. For instance, if Ronda was mistaken about knowing l_1 or l_2 , then she might be tired without coffee in the morning and Daphne would be cranky without her plush-deer. Yet she would at some point be awake even without coffee and Daphne's mood would surely improve given the chance of seeing a real deer. Third, Ronda's decision to go to C1 is intuitively clearly rational, given the beautiful scenery and the fun hike to get there. Furthermore, she confidently knows that (h_1) the cat is well taken care of. While she less confidently knows that l_1, l_2, \dots, l_{100} , she can reasonably expect to know most of l_1, l_2, \dots, l_{100} . Perhaps she is in fact mistaken about knowing to have packed one thing or another, but that does not significantly affect their trip.

To see what Sophisticated Infallibilism would say about the case, the utility assignments have to be specified. Let's suppose that in C1, the falsity of h_1 —resulting in the death of their cat—contributes -5000 utility to what is at stake. In C2, the falsity of h_1 contributes 0 utility: after her dental appointment, Ronda will stop by their home anyway and check on the cat. In both C1 or C2, h_1 being true upholds the status quo and thus contributes 0 utility. Let's suppose further that in C1, the truth of each of l_1, \dots, l_{100} contributes 15 utility, whereas the falsity of each of l_1, \dots, l_{100} contributes -20 utility to the overall outcome (adding a mild annoyance on top of the thing missing). In C2, the truth of each of l_1, \dots, l_{100} only contributes 10 utility, pricing in the lack of beautiful scenery (after all, most things are more enjoyable with a nice view!). However, they also only contribute -5 utility if false: Ronda can bring them along while stopping by home, with the remaining downside being that she has to find them first.

To specify the epistemic characteristics of the case, let's suppose that Ronda has only ordinary, level-1 knowledge that l_1, l_2, \dots, l_{100} , reflecting that she was fairly tired when she packed, but level-2 knowledge that h_1 , reflecting that she turned on the cat feeder with fresh eyes in the morning.

What does Sophisticated Infallibilism tell us about Ronda's decision in Camping Trip? Since her decision involves goods—their life of her cat—that demand a high level of epistemic security (Schulz, 2017, pp. 469–470), Sophisticated Infallibilism tells us that she should make her decision from the perspective of a correspondingly secure level of knowledge. Let's suppose that this is her level-3 knowledge.¹⁹

Let's assume that given her level-3 knowledge, the probability of l_1, \dots, l_{100} is .7 each and the probability of h_1 is .99. While not much hangs on this, the probability distribution strikes me as plausible: Ronda was tired when acquiring knowledge of l_1, \dots, l_{100} , so taking on the perspective of a more secure level of knowledge will not include propositions that just about meet the standard for level-1 knowledge. However, they will still be assessed as moderately probable from this more secure perspective. By contrast, she acquired her knowledge that h_1 after a full night's sleep, which in the absence of defeating conditions is a very reliable (but not infallible) belief-forming method that produces fairly secure knowledge. That h_1 is still highly likely, given the more demanding perspective of level 3 knowledge, is thus natural.

¹⁹ This aligns with Schulz (2017, pp. 471–475) very elaborate proposal on how to determine the needed level of knowledge, the details of which can be omitted.

The problem now is that Sophisticated Infallibilism fails to vindicate our intuitions about Camping Trip. As we noted earlier, it is clearly intuitively rational for Ronda to opt for C1. Sophisticated Infallibilism makes the opposite prediction. The view tells us that, given Ronda's level-3 knowledge, it is not rational for her to go to C1:

$$\begin{aligned} EU(C1)_3 &= 100 \times (.7 \times 15 + .3 \times (-20)) + .99 \times 0 + .01 \times (-5000) = 400 \\ EU(C2)_3 &= 100 \times (.7 \times 10 + .3 \times (-5)) + .7 \times 0 + .3 \times (0) = 550 \\ EU(C2)_3 &> EU(C1)_3 \end{aligned}$$

According to Sophisticated Infallibilism, it is rational for her to go to the camping spot C2 close to their home, which is intuitively implausible.

Note that we should not take our intuition that it is rational for Ronda to go to C1 mislead us into thinking that she perhaps does in fact have level-3 knowledge that h_1 , conditional on which $\sim h_1$ receives probability 0. Rather, it is intuitively rational for her to "tolerate a very small probability of a fairly big loss for a moderate gain" (Schulz, 2017, p. 478). We regularly do so, too, when we go about our day driving cars, crossing streets and eating food prepared by others. This does not mean that we have secure knowledge that doing so can't go wrong, but rather that it is sufficiently likely that it won't.

We can trace the source of this problem to Sophisticated Infallibilism's demand for *indiscriminate high security* in high-stakes cases for *all* decision-relevant propositions, no matter whether this security is needed for all of them. Sophisticated Infallibilism says that one's *overall epistemic perspective* on all decision-relevant propositions should take on a level of secure knowledge that matches what is at stake. However, this demand neglects important differences between these propositions: not all propositions that matter for one's decision need this level of epistemic security. If we look at Camping Trip, for instance, each of l_1, \dots, l_{100} contribute very little to the overall potential outcome of the decision, so it is doubtful that Ronda needs to treat them individually with the same level of epistemic security that is rightfully afforded to h_1 . After all, the latter is much more closely tied to the potential outcome that makes Camping Trip a high-stakes case.

The problem is that Sophisticated Infallibilism cannot recognize these differences. Instead, if stakes are high, it advises us to indiscriminately assess all propositions from a highly secure perspective, leading to intuitively irrational choices. Put succinctly, it's like moving to a new apartment and packing one's entire belongings in bubble wrap because a single vase needs to be protected like this.

The problem of indiscriminate security affects Sophisticated Infallibilism on a structural level. Hence, proponents of KPR should be wary to turn to any variant of Sophisticated Infallibilism for a defence of their view.

5 How to rely on what you know

In what follows, I will argue that there is a reason the proposals discussed so far have failed to serve as plausible knowledge-based decision theories: there is a commonly held, unquestioned assumption in the literature about the central notion of KPR, the notion of reliance, that is inadequate and has downstream effects on how we think

about knowledge determining rational action. The assumption is that once we may rely on a proposition in practical reasoning, we may rely on it with an *infallibilist stance*. In what follows, I will argue that our ordinary notion of reliance is much more fallibilist than commonly assumed. This, in turn, has important ramifications for the kind of knowledge-based decision theory we should adopt.

It is a commonly held assumption about the notion of reliance that once we may rely on a proposition in practical reasoning, we may rely on it with an infallibilist stance. Let's call this way of relying on things *infallible reliance*. Illustrative for this way of thinking about reliance is Fantl and McGrath's statement that once you know a proposition, "you can *take it for granted, assume it's true, count on it, take it to the bank, and book it*" (2012, p. 441, their emphasis). That many accept this assumption about the notion of reliance can most easily be seen by considering what is commonly accepted to *follow* from permissible reliance. In the debate about norms of practical reasoning, many hold that if one may infallibly rely on a proposition p in practical reasoning, one ought to choose the action that maximizes expected utility, *conditional on p* (see, e.g., Fantl & McGrath, 2002, pp. 76–78, 2007, p. 559; Hawthorne & Stanley, 2008, p. 580f; Ross & Schroeder, 2014, p. 261; Locke, 2014, p. 86; Beddor 2021, p. 194).²⁰ Recall that by conditionalizing a known proposition on what we know, we assign probability 1 to said proposition and probability 0 to its negation. If relying on p is spelled out by conditionalization, then relying on p means that we don't assign any probabilistic weight—and thus ignore—the possibility, however remote, that p might be false. Hence, the infallibilist stance.

For Basic Infallibilism, the problem with infallible reliance arises directly. For Sophisticated Infallibilism, infallible reliance leads to the demand for indiscriminate security for *each and every* proposition we rely on. As Camping Trip shows, this is not tenable either.

In what follows, I will show that we should challenge the assumption that our notion of reliance is so rigidly infallibilist. Indeed, I will argue that our ordinary notion of reliance is much more fallibilist than has hitherto been assumed.

Let's start out by making some observations about our ordinary notion of reliance.²¹ There is something that proponents of infallible reliance get right: in many ordinary decision situations, we do infallibly rely on things. Consider the following case that exemplifies this:

Watering Paul enjoys having plants on his balcony and relies on a moisture gauge to take proper care of them. The gauge is an instrument with no backup system. Since Paul is quite bad at estimating by other means how much water the plants need, he solely relies on the gauge's reading for watering the plants. Giving them the correct amount of water is not a high stakes situation (for him, anyway), so he's fine with treating the propositions acquired by reading the gauge as certain.

The example illustrates our way of relying when stakes are low. Whether we rely on restaurant opening times we read online, the weather report we saw on TV, the promise a usually reliable friend gave us or the reading of the moisture gauge of our plants,

²⁰ If two or more actions maximize expected utility, it is rationally permissible to choose among them.

²¹ For some general discussion of the notion of reliance, see Holton (1994) and Alonso (2014, 2016).

we rely on them with infallibilist conviction—if not much hangs on it. Importantly, doing so seems entirely appropriate.

My claim now is that those who accept an infallibilist conception of reliance are mistaken about how we rely on things when stakes are high. To see this, consider the following case:

Low Visibility Christine is trying to land her old Cessna on a remote runway during a foggy night. Since visibility is low, she has to rely on various instruments to successfully land her plane. For instance, she constantly has an eye on her altimeter which measures the altitude of the plane. Yet, she is aware that instruments can and do fail, especially in old airplanes, so she is alert to any signs of malfunctioning. As a backup, she has a backup altimeter on her phone that she can consult if the main altimeter breaks down.

In Low Visibility, Christine relies on the built-in altimeter to land her plane. Importantly, however, she does not rely on the built-in altimeter the way Paul relies on his moisture gauge. In contrast to Paul, Christine relies on the built-in altimeter in a more fallibilist-minded way: in doing so, she stays alert to the possibility that the altimeter might malfunction and that she might have to consult the backup altimeter on her phone. Let's call the way Christine relies on her built-in altimeter *fallible reliance*. Generally put, when we fallibly rely on things, we rely on them in a more fallibilist-minded way, i.e., by remaining sensitive to the possibility that the things we rely on might be undependable. Again, importantly, that Christine relies on the primary altimeter only in this way seems entirely appropriate.

Here are two quick objections to this way of describing the case. First, couldn't one say instead that Christine *infallibly* relies on the *combined* system of altimeters? This description is inadequate. To see this, suppose the built-in altimeter indeed malfunctions and Christine has to use the backup altimeter on her phone. In that case, we would say that Christine *stopped* relying on the built-in altimeter and *started* relying on her backup altimeter to assess the altitude. Naturally, one only relies on things in φ -ing if one let's oneself be guided by these things in φ -ing. This is clearly true of the built-in altimeter, as Christine lets herself be guided by it while landing the plane. The back-up altimeter, by contrast, is true to its name: something that *could* guide her if she had to rely on it in a case of malfunction.

However, one might now object, if Christine starts to rely on her back-up altimeter, doesn't she just stop infallibly relying on the built-in altimeter? Again, I think this description is inadequate, with the main worry being that her relying seems importantly different from typical cases of infallible reliance: when we infallibly rely on p , we have psychologically closed the matter of whether p and would be blindsided if it turned out that $\sim p$. For instance, someone who is diagnosed with cancer will be caught off guard in this sense, realizing that they can't take their health for granted any more. Christine, by contrast, shows low need for closure (Nagel, 2008) and the awareness that it is in principle possible that altimeters in old planes can break, indicating fallible reliance. If her altimeters would break, she would not be caught off guard. She would initially be concerned but pull out her phone with the back-up altimeter right after.

There are plenty of cases like Low Visibility. A young athlete fallibly relies on her knowledge of having a clean health record in signing a disability insurance. An elderly

man fallibly relies on his knowledge that he can walk the round trip to the park while still carrying an umbrella to possibly lean on to. An engaged couple fallibly relies on their knowledge of their unshakable love for each other in signing a prenuptial agreement. In all of these cases, we rely on our knowledge in a way that is sensitive to the possibility that we might fail to know after all.

Some might initially feel somewhat uneasy about the notion of fallible reliance. Doesn't it sound odd for Christine to say, "I rely on the reading of the altimeter, but it is possible that the reading is false"? While I think there is a tension to be felt, we shouldn't worry about it. First, the felt tension is a familiar one: fallibilism about knowledge licenses similarly odd-sounding concessive knowledge attributions (CKAs), such as "I know that p but possibly $\sim p$ ". Yet, many accept fallibilism, because they think these oddities can plausibly be explained away in one way or other (Stanley, 2005; Dougherty & Rysiew, 2009, 2011; Worsnip, 2015)²² or because they prefer it over worse alternatives (Lewis, 1996, p. 550). Hence, those who accept fallibilism about knowledge should not hesitate to embrace fallibilism about reliance and related notions. A second point worth making is that certain concessive reliance attributions sound just fine. Consider, for instance, "We can rely on prices being stable for deciding on the budget for the next quarter, even though it is *of course* possible that oil prices may fall dramatically". Or consider: "We have presented considerable evidence indicating the existence of quarks. While we acknowledge that there is always a chance that these results will be overturned, we will, in the design of further experiments, rely on there being quarks."²³ If some concessive reliance attributions sound fine, we might wonder how dependable a feeling of uneasiness about the notion of fallible reliance really is.

I argued that our fallibilist habits of managing our own fallibility affect *how* we rely on our knowledge. If stakes are high, we rely on our knowledge but remain sensitive to the possibility that we might fail to know after all. If stakes are low, we can make our life easier by relying on our knowledge with an infallibilist stance and treat it as certain, because it is certain enough for our purposes. If this is right, then our notion of reliance is not a rigid, infallibilist notion, as it has hitherto been assumed, but a flexible one: depending on what is at stake, reliance will manifest differently.²⁴

Let's call the resulting view about how can we rely on our knowledge *Flexible Fallibilism*. A version of KPR in the spirit of Flexible Fallibilism goes as follows:

KPR_{FF} If one knows that p one may (1) infallibly rely on p in low-stakes practical reasoning and (2) fallibly rely on p in high-stakes practical reasoning.

KPR_{FF} is a precisification of KPR. It captures the idea that our fallibilist habits affect both how we rely on what we know and when different ways of relying are appropriate. Naturally, this should have ramifications for the role knowledge plays in determining

²² I think similar explanations could be provided for the notion of fallible reliance, although showing this would take us too far off track.

²³ Cf. Dougherty and Rysiew (2009, p. 130).

²⁴ More work might be done by the notion of reliance. For instance, Schulz (2021a) has recently argued that to deal with decision-making under uncertainty, a notion of *partial* reliance, tied to normatively proper credences, allows us to capture practical reasoning involving propositions that are *not* known.

what is rational for us to do. In what follows, I will offer a novel knowledge-based decision theory that implements this idea.

6 How to act on what you know

My proposal employs a generalized form of conditionalization, so-called *Jeffrey conditionalization* (Jeffrey, 1983, ch. 11). Jeffrey conditionalization is typically conceived as a diachronic rule that we should use to learn from uncertain evidence. In what follows, this standard conception will be used to formally introduce Jeffrey conditionalization. I will then reinterpret the latter for my purposes.

Jeffrey conditionalization allows us to conditionalize on a weighted partition of mutually exclusive and jointly exhaustive propositions $E = \{(E_1, \mu_1), (E_2, \mu_2), \dots, (E_n, \mu_n)\}$. The weights μ_1, \dots, μ_n sum up to 1 and codify one’s uncertainty about the propositions in E . Formally:

Jeffrey Conditionalization
$$P_J(\cdot) = \sum_{i=1}^n \mu_i P(\cdot | E_i).$$

In what follows, I will not use the standard conception of Jeffrey conditionalization as a diachronic learning rule. Rather, the key idea will be to employ Jeffrey conditionalization as a *synchronic* rule that measures how *our way of relying contributes to determining what is rational for us to do*.²⁵ Having specified this, I can now explicitly state my proposal for a flexibly fallibilist knowledge-based decision theory.

KDT_{FF} An action A is rational for a subject S to perform iff A maximizes expected utility, Jeffrey-conditional on S ’s knowledge.

To see how KDT_{FF} works, let’s suppose we rely on our knowledge that p . If stakes are low, I suggested that we can agree with infallibilists in that we may infallibly rely on p and ignore the possibility that $\sim p$. To determine what is rational to do, we have to consider the partition that consists of a piece of knowledge p and its negation: $\{p, \sim p\}$. To capture the idea that we can ignore $\sim p$ in decision-making, we should choose the weights accordingly: $\mu_p = 1$ and $\mu_{\sim p} = 0$, giving us the weighted partition $L = \{(p, 1), (\sim p, 0)\}$. Note, however, that it is a well-known feature of Jeffrey conditionalization that it yields standard conditionalization if we assign weights like this.²⁶ Hence, Flexible Fallibilism agrees with Basic Infallibilism that if stakes are low, we should simply conditionalize on our knowledge.

Where Flexible Fallibilism disagrees with Basic Infallibilists is when stakes are high. I argued that when much is at stake, we only fallibly rely on our knowledge, i.e.,

²⁵ Bayesians have been hesitant to employ Jeffrey conditionalization because its applicability depends on an invariance condition being met (see, e.g., Levi (1967), Jeffrey (1970), Jeffrey (2004), Williamson (2000) and Schwan and Stern (2017)) which is not met in all cases (for an example in which it isn’t, see Osherson and Zhao (2010, pp. 2533–2535)). To update by Jeffrey conditionalization, the conditional probabilities before updating cannot change in the course of updating due to the uncertain evidence one updates one’s credences with. However, this is only a worry if Jeffrey conditionalization is employed as a *diachronic* learning rule. In my account, it is employed as a *synchronic* rule. No learning takes place in the decision situation, so the needed invariance condition is trivially satisfied.

²⁶ With $L = \{(p, 1), (\sim p, 0)\}$, we get: $P_L(\cdot) = \mu_p \times P(\cdot|p) + \mu_{\sim p} \times P(\cdot|\sim p) = 1 \times P(\cdot|p) + 0 \times P(\cdot|\sim p) = P(\cdot|p)$.

we rely on it while being sensitive to the possibility that we might fail to know after all. When it comes to determining what is rational to do, we should again choose weights that corresponds to this stance towards our knowledge. One way to do so, generally speaking, is to include non-negative parameter Δ in our weighting that reflects our fallibilist stance towards our knowledge: $\mu_p = 1 - \Delta$ and $\mu_{\sim p} = 0 + \Delta$, resulting in the following weighted partition: $H = \{(p, 1 - \Delta), (\sim p, 0 + \Delta)\}$.

How should we assign the weights when stakes are high? One *salient* answer is to look at the epistemic probability of the proposition *that one knows that p*. This probability will typically be less than 1, reflecting the fallibilist idea that knowing is compatible with an epistemic chance that p (and thus Kp) is false.²⁷ Weighing our knowledge that p with the epistemic probability of Kp naturally fits the idea that if stakes are high, our fallibilist habits raise awareness of our own fallibility and incorporate our assessment of our fallibility into how strongly we should let p determine what is rational for us to do.

Before elaborating the conception of epistemic probabilities at play here, let me first illustrate how my proposal works. Let's assume that the epistemic probability of the proposition (Kp) that you know Caesar was born in 100 BC is .9, reflecting that it is epistemically highly likely that in Jellybean you know that p . Correspondingly, the parameter Δ is set at .1, capturing your remaining sensitivity to the possibility that you might after all not know that p . The resulting weights are as follows: $\mu_p = .9$; $\mu_{\sim p} = .1$. Choosing the weights like this is a natural option for the case at hand, since the impact of one's knowledge on what is rational for one to do is just as high as is warranted by one's epistemic position towards one's knowing that p . It corresponds to the notion that just because we know, we should not ignore that we only fallibly do so when we determine what is rational for us to do.

With the resulting weighted partition being $J = \{(p, .9), (\sim p, .1)\}$, a quick calculation reveals that not answering ($\sim A$) maximizes expected utility, Jeffrey-conditional on one's knowledge p , and is hence the rational thing to do.

$$\begin{aligned} EU_{FF}(A) &= (\mu_p \times P(p|p) + \mu_{\sim p} \times P(p|\sim p)) \times 1 + (\mu_p \times P(\sim p|p) + \\ &\mu_{\sim p} \times P(\sim p|\sim p)) \times (-1000) = -99.1 \\ EU_{FF}(\sim A) &= 0 \\ EU_{FF}(\sim A) &> EU_{FF}(A). \end{aligned}$$

This is the intuitive verdict. Furthermore, one can (fallibly) rely on one's knowledge in deciding that not answering is rational. Thus, according to Flexible Fallibilism, Jellybean is not a counterexample to the view that knowledge is sufficient for permissible reliance. Since the respective subjects are in a high-stakes case and fallibly rely on their knowledge, they simply keep in mind the small epistemic chance that they might fail to know after all and act accordingly.

Having illustrated how my proposal works, let me now specify the notion of epistemic probability at work. While there are many conceptions of epistemic probability available,²⁸ the employed conception of epistemic probability needs to be suitably

²⁷ See, e.g., Reed (2002), Fantl and McGrath (2009), Dougherty and Rysiew (2009) and Dougherty (2011) for why we should conceive of fallibilism in probabilistic terms.

²⁸ For an overview, see Moon and Jackson (2020, p. 656).

knowledge-based. If knowledge did not play a role in determining epistemic probabilities, then one might worry that the resulting view might not be properly classified as knowledge-based decision theory.²⁹ Fortunately, however, a fitting and independently plausible knowledge-based conception of epistemic probability is available. On this conception, we arrive at the epistemic probability of a proposition by conditionalizing it on a selected set of known propositions.³⁰ This set is typically taken to *include* certain basic types of knowledge (e.g., perceptual, introspective and evidential knowledge, etc.) and background knowledge (e.g., about the reliability of our cognitive faculties or instruments). The set is taken to *exclude* inductive or inferential knowledge that “goes beyond our evidence” (Goodman & Salow, 2018, p. 90) and that is plausibly particularly susceptible to being scrutinized by our fallibilist habits.³¹

Let’s apply this conception to Low Visibility and Jellybean. Suppose that, in the former case, Christine sees that the altimeter displays an altitude of y and gains perceptual knowledge that the altimeter displays an altitude of y . How epistemically likely should she consider it that she has inductive knowledge that the altitude *is* y ? Conditional on her perceptual knowledge and her background knowledge that the altimeters in her plane are fairly reliable, she considers it epistemically highly likely, though not certain, that she knows that the altitude is y . By fallibly relying on her knowledge, she rationally stays alert and keeps her phone altimeter close by. Likewise, in Jellybean, you are confident, but not certain, that Caesar was born in 100 BC and your background knowledge includes that you very occasionally have misremembered specific historic dates. Conditionalizing on this knowledge, it is epistemically highly likely that you in fact know that Caesar was born in 100 BC. Weighing your knowledge that p accordingly in your practical reasoning, you rationally decide to not give an answer, given the meager rewards and risk of an extremely painful shock. This knowledge-based conception of epistemic probabilities thus provides us with a natural and intuitively plausible description of these cases.

The proposed account also delivers an answer to the question concerning what is at stake in rational decision-making.³² As Schulz (2017, pp. 470–471) rightly notes, outcomes need to be in one’s epistemic ken to contribute to what’s at stake. They are in one’s epistemic ken on my view iff they are an epistemically possible consequence of an available action, i.e., iff they receive a probability of >0 , conditional on the selected set of knowledge characterized above and the action being performed.

The resulting decision theory is knowledge-based through and through. If stakes are low, we give the known propositions a weight of 1, resulting in conditionalizing on them to determine what is rational for one to do. If stakes are high, I argued that we should weigh the known propositions by the epistemic probability that one knows

²⁹ I am grateful to two anonymous referees for pressing me on this.

³⁰ I am grateful to an anonymous referee who encouraged me to consider this option.

³¹ While offering a full-fledged account of epistemic probability along these lines is beyond the scope of this paper, proposals from the literature include those from Ichikawa (2017) (who calls the restricted set in question “basic knowledge”), Littlejohn (2011) (“non-inferential knowledge”), Comesaña (2018) (“basing evidence”) and Goodman and Salow (2023) (“evidential knowledge”). As it is argued by some of these authors, this conception independently has advantages over a Williamsonian conception.

³² I am grateful to an anonymous referee for inviting me to elaborate here.

them, which is in turn determined by the selected set of knowledge characterized above.

Now I am also in a position to give a fuller characterization of fallible reliance. By fallibly relying on p , we are sensitive in practical reasoning to the question how epistemically likely it is that we know that p . This, in turn, lessens the impact of p in our practical reasoning, corresponding to its reduced impact on determining what is rational for us to do, as specified by KDT_{FF} . We can further clarify this by putting it in terms of reasons. Suppose p is a reason for φ -ing. If S fallibly relies on p in her practical reasoning, then S should take p to speak slightly less in favor of φ -ing than if she had infallibly relied on p . Depending on one's decision situation, these general characteristics of fallible reliance can manifest in various ways. In Low Visibility, fallible reliance will manifest as an increased alertness by Christine towards signs that the altimeter reading she relies on is false, making her practically reason that she should keep a backup altimeter at hand. In Jellybean, fallible reliance will manifest as a hesitation to answer, giving expression to one's registering of the lessened (and ultimately insufficient) support that knowing that Caesar was born in 100 BC provides for this possible course of action. By being sensitive in one's practical reasoning to the possibility that one may fail to know after all, one weighs one's knowledge accordingly and concludes that one should better not answer, arriving at what is rational for one to do.

I argued that one salient way to set the weights in Jeffrey-conditionalizing on known propositions is to use the epistemic probabilities that one knows the respective propositions. In the next section, I will give a general account of how to weigh one's knowledge. I will identify two factors relevant in weight-setting, which I call the *total stakes factor* and the *stake proportionality factor*. The resulting account not only makes the correct predictions in the cases of interest, but has the independent virtue of providing us with a knowledge-based explanation of when to simplify our decision-making.

7 Weight setting

First, as already suggested, one important factor is the *totality of what is at stake*. Assuming that we can divide what is at stake into different levels (Schulz, 2017, p. 469), our setting of the relevant weights should plausibly be responsive to what level of stakes we are dealing with in decision-making. If stakes are low, assigning $\mu_K = 1$ to what one knows and $\mu_{\sim K} = 0$ to what one does not know is adequate. It is a harmless way of simplifying one's decision-making that saves cognitive resources and allows one to get on with one's life. If we face middling stakes, we might think that we should not disregard the possibility that what we know is false, but still simplify somewhat. Suppose that the epistemic probability of your decision-relevant knowledge that p is .87963 and the probability of $\sim p$ is .12037 accordingly. One way to simplify your reasoning without neglecting the possibility that what you take yourself to know, might turn out to be false would be to use a $\mu_K = .88$ for p and a $\mu_{\sim K} = .12$ for $\sim p$. Finally, if stakes are high, we might consider using the exact relevant epistemic probabilities of that one knows the relevant propositions as weights to accurately adjust its impact

on rational action to its epistemic standing. Thus, depending on the level of the stakes, we should choose weights that allow as much simplification in decision-making as is permissible for what is at stake. According to one plausible view, then, the range of weights for one's knowledge goes from 1 to the epistemic probability of whatever one knows: if stakes are low, we may fully simplify. If stakes are high, we should match the impact of our knowledge on what is rational to do with its epistemic probability. Intermediate levels of stakes, in turn, allow for various grades of simplification in between. Flexible Fallibilism thus yields a natural view about the simplification of practical reasoning: simplify just as much as is warranted by what is at stake.

The second factor that should affect the weights is *stake proportionality*. How we should set the weights also depends on how much the truth-value of a decision-relevant proposition contributes to a given outcome. Recall how in *Camping Trip*, the truth-values of the propositions l_1, \dots, l_{100} respectively have only little impact on what is at stake. Intuitively, Ronda may respectively rely on these propositions infallibly. By contrast, Ronda should only fallibly rely on her knowledge of h_1 , the proposition on which proportionally a lot hangs when it comes to what is at stake. It is now easy to see that the indiscriminate security problem does not arise for Flexible Fallibilism: Ronda should Jeffrey-conditionalize on l_1, \dots, l_{100} with $\mu_K = 1$ and $\mu_{\sim K} = 0$ weights, yielding standard conditionalization, but Jeffrey-conditionalize on h_1 with a weight that matches the epistemic probability of her knowledge of h_1 , reflecting the proposition's proportional contribution to what is at stake.³³

How do these two factors interact? For high stakes decisions that concern only two ways the world might be—such as *Jellybean*—the stake proportionality factor simply collapses into the total stakes factor. When it comes to high stakes decisions that involve more than two ways the world might be—such as in *Camping Trip*—the stake proportionality factor acts as a *corrective* to the total stakes factor. The total stakes factor can be understood as reflecting our general fallibilist habits that initially manifest if stakes are high, whereas the stake proportionality factor fine-adjusts the manifestation of these habits by retracting the fallibilist stance towards those propositions that do not contribute significantly to what is at stake.

It is worth pointing out that the picture just presented suggests that pragmatic factors can affect the weights needed for Jeffrey-conditionalizing on our knowledge. The result is a mild, non-threatening form of pragmatic encroachment when it comes to the epistemic input that contributes to determining what is rational for us to do.³⁴

³³ Why shouldn't one instead directly use one's knowledge-based epistemic probabilities and simplify accordingly? Call this the *mere-epistemic-probability* view (MEP). Let me point out three advantages of my more elaborate proposal: First, knowledge marks propositions concerning which we should simplify. On the MEP view, there is no guidance concerning which propositions we can treat as certain in low-stakes scenarios, carrying with it the risk of over-/undersimplification. Second, my proposal—but not the MEP view—vindicates the intuition that Ronda does in fact rely on l_1, \dots, l_{100}, h_1 in her practical reasoning. It does so while at the same time still giving weight to our fallibilist inclinations. Finally, my proposal allows us to defend the view that knowledge is sufficient for reliance in practical reasoning, which, as noted above, many think is an intuitively and theoretically attractive view. Needless to say, not everyone shares these theoretical commitments and intuitions. However, if one is sympathetic to a knowledge-based view of reliance and rational decision-making, then my more elaborate proposal is preferable over the MEP view. I am grateful to an anonymous referee for pressing me on this.

³⁴ One might object here by claiming that subjects who let pragmatic considerations affect this input are susceptible to diachronic Dutch books (see Greco, 2013; Rubin, 2015; Schroeder, 2018 for such a concern

Importantly, the epistemic domain remains entirely unaffected. On my view, pragmatic considerations affect only the *relation* between knowledge and action, but not knowledge itself.

We thus arrive at a full characterization of Flexible Fallibilism, a knowledge-based view of practical reasoning and decision-making. It encompasses two complementary components:

KPR_{FF} If one knows that p , one may (1) infallibly rely on p in low-stakes practical reasoning and (2) fallibly rely on p in high-stakes practical reasoning.

KDT_{FF} An action A is rational for a subject S to perform iff A maximizes expected utility, Jeffrey-conditional on S 's knowledge.

KPR_{FF} is a precisification of KPR, telling us how we should rely on what we know, given what's at stake. KDT_{FF} is the supplementary knowledge-based decision theory that determines what is rational for us to do, given what we may rely on and how we may rely on it. Taken together, we get a view about the simplification of practical reasoning: how we may rely on our knowledge (mirroring the impact our knowledge has on determining what is rational for us to do) corresponds to how much is at stake, as specified by the total stakes factor and the stake proportionality factor, yielding various stages of simplification.

In the next section, I will compare Flexible Fallibilism to another proposal, inspired by views recently advanced by Jackson (2019a) and Moss (2013, 2018). I will show that while this further proposal gives us the right results in high-stakes cases and tells us when to simplify our practical reasoning, Flexible Fallibilism has various advantages over it and is thus preferable.

8 Can we do (even) better with credal knowledge?

In the final section, I will briefly compare Flexible Fallibilism to a further proposal for defending the view that knowledge is sufficient for permissible reliance in practical reasoning. This further proposal combines two claims from Jackson (2019a) and Moss (2013, 2018). Jackson's (2019a) claim is that when stakes are high, we do not rely on our beliefs, but on our credences. This, on her view, explains our intuitions in high-stakes cases and also yields a view of when we simplify our reasoning. Moss's (2013), (2018) claim is that credences can constitute knowledge.

The backdrop of the first component of the proposal is a dualist metaphysics of beliefs and credences ('dualism', for short). One main motivating idea for dualism is that beliefs and credences seem to play different roles in practical reasoning (Jackson, 2019a, pp. 516–517).³⁵ By relying on our beliefs, we take the believed propositions

about pragmatic encroachment on credences). For instance, while the subject would choose a weight of 1 for some piece of knowledge p in a low-stakes context, the subject might choose a weight < 1 for p in a high-stakes context. In response, we could plausibly claim that a subject considers diachronic factors when simplifying and thus won't be susceptible to a Diachronic Dutch book, given that expediency and resourcefulness are characteristically forward-looking motivations for simplification. For other convincing responses that are easily adaptable to Flexible Fallibilism, see Gao (2019b) and Beddor (2021).

³⁵ See also Lance (1995), Smithies (2012), Wedgwood (2012), Ross and Schroeder (2014), Tang (2015), Staffel (2019), Weisberg (2020).

for granted and ignore small error possibilities, thereby simplifying our reasoning and reducing our cognitive load.³⁶ By contrast, when we rely on our credences, we don't ignore these small possibilities any more, resulting in higher accuracy but increased cognitive costs. Jackson now argues that when stakes are low, we want to be cognitively efficient and thus simplify our reasoning by relying on our beliefs. However, if stakes are high, such as in Jellybean, we should expend some cognitive resources and rely on our credences instead, giving some weight to the possibility that Caesar being born in 100 BC might be false after all (Jackson, 2019a, p. 521).

The second component is due to Moss (2013, 2018), who has recently argued that our credences can have similar epistemic qualities as our beliefs, and thus can likewise constitute knowledge (I will call this kind of knowledge 'credal knowledge', in contrast to ordinary 'outright' knowledge).

A version of KPR that combines these components now goes as follows:

KPR_{DI} (1) One may rely on a proposition p in low-stakes practical reasoning if one knows that p ; (2) One may rely on one's credence in p in high-stakes practical reasoning if one's credence in p constitutes credal knowledge.

Since the view says that we may infallibly rely on the two kinds of knowledge, depending on stakes, I will call it *Dual Infallibilism*.

Again, a corresponding knowledge-based decision theory can be developed. According to this knowledge-based decision theory, different kinds of knowledge are relevant for determining expected utility, depending on what is at stake: If stakes are low, then an action is rational for one to perform if it maximizes expected utility, conditional on the totality of what one knows. However, if stakes are high, we use the credences that qualify as credal knowledge as input to our expected utility calculation. Let's call a credence function that qualifies as knowledge " C_K " and the result of an expected utility calculation that uses C_K as epistemic input one's C_K -expected utility. If stakes are high, Dual Infallibilism tells us that an action is rational for one to perform if it maximizes C_K -expected utility.

It is easy to see that Dual Infallibilism gives us the right results in Jellybean. Let's assume that we have a .9 credence in p , corresponding to the description that we are confident, but not certain, that Caesar was born in 100 BC. You are well acquainted with Roman History, so this credence plausibly constitutes credal knowledge (i.e., $C_K(p) = .9$). Since stakes are high in Jellybean, Dual Infallibilism tells us that we may infallibly rely on our credal knowledge about p . Accordingly, to now determine what is rational for us to do, we need to calculate our C_K -expected utilities. Since $C_K(\sim p) = .1$, the possibility that p is false, and the negative outcome associated with it, receive minor probabilistic weight. Hence, not answering maximizes C_K -expected utility, the intuitively correct prediction.³⁷

A further selling point of Dual Infallibilism is that it brings along a theory of when to simplify practical reasoning. In these regards, Dual Infallibilism is on a par with Flexible Fallibilism. However, Dual Infallibilism faces various problems. As I will argue in what follows, Dual Infallibilism is both less flexible and epistemically more

³⁶ But see Dinges (2021) who claims that it is not our beliefs, but our credences that simplify our reasoning.

³⁷ $EU(A) = C_K(p) \times 1 + C_K(\sim p) \times (-1000) = -99.1$; $EU(\sim A) = 0$; $EU(\sim A) > EU(A)$.

demanding than Flexible Fallibilism. While these are not fatal problems, they suggest that Flexible Fallibilism is the preferable choice.

The first problem concerns the account of simplification provided by Dual Infallibilism. Dual Infallibilism recognizes only two kinds of stakes: when stakes are low, it tells us to simplify by relying on your outright knowledge; when stakes are high, it tells us to rely on our credal knowledge. But what about intermediary stakes, where we may plausibly simplify a bit? The worry is that Dual Infallibilism does not allow *continuous simplification*, providing no plausible treatment of intermediary stakes. By contrast, Flexible Fallibilism allows for continuous simplification: depending on what is at stake, the weights can be set in a way that yields the appropriate degree of simplification.³⁸

As an anonymous referee has noted, Dual Infallibilist could argue that one's credences could likewise be subject to continuous simplification. However, those committed to a dualist metaphysics of belief should be inclined against this move. After all, one motivating thought for the dualist picture is to argue that credences and beliefs play different roles in our cognition. If dualists were to allow the continuous simplification of credences, perhaps sometimes even treating propositions as certain in which one has a credence of just above .5, the line between the functional roles of credences and beliefs would be significantly blurred. This, however, would undermine a key motivation for adopting a dualist picture in the first place (Jackson, 2019b). Furthermore, there are plausibly principled barriers to credences being the right kind of attitude for simplification (Palmira, 2023).

The second problem concerns the epistemic demandingness of Dual Infallibilism. To cover both low- and high-stakes decision-making, both outright and credal knowledge have to be available. However, as I will show below, there are cases in which this availability is non-trivial. I will argue that Dual Infallibilism thus either makes it harder to rely on one's knowledge or makes the wrong predictions about high-stakes decision-making.

Consider the following case to illustrate this worry.

Gettierized Jellybean Let everything be as above in Jellybean. Suppose that in addition to the original researcher (let's call her "Hi"), a second researcher (let's call him "Lo") comes in and also asks you questions about Roman history. When it comes to Lo's line of questioning, the conditions are slightly different: you win a Jellybean if you answer correctly, but you receive a barely noticeable electric shock if you answer incorrectly. While sipping on a cup of coffee, you contemplate what is rational for you to do. Given that you think that Dual Infallibilism is true, you decide to rely on your credal knowledge to determine whether it is rational to answer Hi, and to rely on your outright knowledge to determine whether it is rational to answer Lo. Since you take yourself to have confident, but not certain knowledge that p , you infer that your corresponding high confidence in p constitutes credal knowledge. However, unbeknownst to you, Hi has mixed a substance into your coffee that affects this kind of introspective inference: in most cases, inferring your

³⁸ Jackson (2019a, pp. 516–517) is aware that intermediate stakes are problematic on her account. However, she is only concerned with the worry that if stakes are intermediate, we have to decide whether we should rely on our belief or credences, which involves additional cognitive effort. She does not address the inadequacy of either option, given the intermediary stakes.

confidence from your knowledge has the side effect that your confidence is slightly lowered in this process. Yet, you have been lucky: although your confidence could easily have been lowered by the inference, you are in one of the rare cases where this doesn't happen.

Does your credence constitute credal knowledge? No, since it could have easily been the case that your credence had failed to correspond to how likely p is.³⁹ Hence, your credence is Gettierized and thus fails to constitute credal knowledge. However, you still know outright that Caesar was born in 100 BC, since a counterfactually slightly lowered confidence doesn't affect your outright knowledge. As Gettierized Jellybean shows, knowing that p does not guarantee credal knowledge about p .

Since you lack credal knowledge about p , how should you now determine what is rational for you to do, according to Dual Infallibilism? Naturally, you should consider what you *do* know. Thus, the next best fallback position is your outright knowledge. However, if you (infallibly) rely on your outright knowledge that p and conditionalize on it to determine what is rational for you to do, we again get the unintuitive result that it is rational to answer Hi's question. As an anonymous referee has pointed out, another response would be to simply acknowledge that there might be unfortunate cases in which you lack the needed credal knowledge. Of course, cases like this can arise for Flexible Fallibilism, too, e.g., when one's outright beliefs fail to constitute knowledge due to being false or Gettierized. However, Flexible Fallibilism, or so I argued, is less susceptible to scenarios of this kind, given that it involves only outright knowledge, and not outright and credal knowledge.

By contrast, Flexible Fallibilism does well when it comes to Gettierized Jellybean. Since you know that p , you rely on p to determine both whether to answer Hi and whether to answer Lo. When it comes to determining whether it is rational to answer Lo, you infallibly rely on what you know and set your weights accordingly ($\mu_p = 1$; $\mu_{\sim p} = 0$). When it comes to determining whether it is rational to answer Hi, you fallibly rely on what you know and likewise set your weights more carefully, say, by considering the epistemic probability of p ($\mu_p = .9$; $\mu_{\sim p} = .1$). For both decisions, knowledge that p is sufficient, no credal knowledge is needed.

So, the second problem of Dual Infallibilism is that it is more demanding. More things have to go epistemically well in order for you to be able to rely on knowledge in both high- and low-stakes decision-making. Furthermore, if things don't go well, like in Gettierized Jellybean, then the most plausible fallback position fails to vindicate our fallibilist intuitions.⁴⁰

³⁹ As Moss notes, credal knowledge, like outright knowledge, also has to satisfy modal conditions meant to capture this, such as safety and sensitivity (Moss, 2013, pp. 18–20, 2018, ch 5), and is susceptible to Gettier cases (Moss, 2013, pp. 9–10, 2018, Sect. 5.6).

⁴⁰ Two worries: First, couldn't one imagine that one's estimation of the epistemic probability is similarly Gettierized? Second, what if it is slightly lower or higher? Weighing seems to be affected in both cases, which might threaten the advantage of Flexible Fallibilism over Dual Infallibilism. My response to the first worry is that this isn't as much of a problem for Flexible Fallibilism, since the weighing also involves non-epistemic factors and therefore plausibly does not presuppose non-Gettierized inputs in the first place. My tentative response to the second worry is that it less strongly affects my proposal: even if we don't get the weighing quite right, we are left with an approximately correct weighing. Dual Infallibilists, by contrast, are left with the unattractive fallback position of having to infallibly rely on their outright knowledge. Hence,

While Dual Infallibilism and Flexible Fallibilism are on a par when it comes to getting the verdict right about cases like Jellybean, I argued that Flexible Fallibilism has advantages over Dual Infallibilism and is thus preferable: first, its theory of simplification accommodates continuous differences in stakes, and second, it is less demanding when it comes to the availability of one's knowledge.

Finally, zooming out, one might wonder how Flexible Fallibilism compares to a Bayesian view, according to which our (rational) credences, not our knowledge, play the central epistemic role in decision theory. While some have argued that knowledge-based decision theory has some advantages over Bayesian decision theory (see, e.g., Greco, 2013) others have been skeptical (see, e.g., Schiffer, 2007, and Fassio & Gao, 2021). I think this is a worthwhile topic for future research (Heil, MS) but beyond the scope of this paper. Here, I argued for the conditional that if knowledge plays a role for our practical reasoning and decision-making, then Flexible Fallibilism is the most promising way to capture this role.

9 Conclusion

In this paper, I defended the view that knowledge is sufficient for permissible reliance against the challenge from high stakes. After having pointed out the need for a knowledge-based decision theory, I argued that extant proposals face the indiscriminate security problem. I then explored two novel proposals, Flexible Fallibilism and Dual Infallibilism, and argued that we should opt for the former: Flexible Fallibilism vindicates the platitude that we may rely on what we know, but reminds us to be attentive to our fallibility when we do.⁴¹

Funding Open Access funding enabled and organized by Projekt DEAL. Funding was provided by Deutsche Forschungsgemeinschaft (Grant no. SCHU 3080/3-1).

Declarations

Conflict of interest The authors declare that they have no conflicts of interest.

Research involving human participants and/or animals This article does not contain any studies involving animals performed by any of the authors.

Informed consent This article does not contain any studies involving human participants performed by any of the authors.

Footnote 40 continued

my view still fares better in the relevant range of cases. I am grateful to two anonymous referees for pressing me on this.

⁴¹ I would like to thank in particular Moritz Schulz for invaluable discussions, and Paul Silva Jr. and three anonymous referees for their constructive comments on earlier drafts of this article. I also would like to thank Dominik Kauss, Jakob Koscholke, Thomas Krödel, Andreas Mueller, Patricia Rich, Sergiu Spătan, Alexandra Zinke and audiences in Dresden, Frankfurt, Glasgow, Hamburg and Salzburg for their helpful feedback. This work was supported by Grant SCHU 3080/3-1 to Moritz Schulz from the DFG as part of the Emmy Noether Group *Knowledge and Decision*.

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/>.

References

- Alonso, F. M. (2014). What is reliance? *Canadian Journal of Philosophy*, 44(2), 163–183.
- Alonso, F. M. (2016). Reasons for reliance. *Ethics*, 126(2), 311–338.
- Anderson, C., & Hawthorne, J. (2019). Knowledge, practical adequacy, and stakes. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 6, pp. 234–257). Oxford University Press.
- Beddor, B. (2020). Certainty in action. *Philosophical Quarterly*, 70(281), 711–737.
- Beddor, B. (2021). Shifty evidence and shifty books. *Analysis*, 81(2), 193–198.
- Brown, J. (2011). Fallibilism and the knowledge norm for assertion and practical reasoning. In J. Brown & H. Cappelen (Eds.), *Assertion: New philosophical essays*. Oxford University Press.
- Comesaña, J. (2018). A note on knowledge-first decision theory and practical adequacy. In B. Kim & M. McGrath (Eds.), *Pragmatic encroachment in epistemology*. Routledge.
- Dinges, A. (2021). Beliefs don't simplify our reasoning, credences do. *Analysis*, 81(2), 199–207.
- Dougherty, T. (2011). Fallibilism. In D. Pritchard & S. Bernecker (Eds.), *The Routledge companion to epistemology*. Routledge.
- Dougherty, T., & Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1), 123–132.
- Dougherty, T., & Rysiew, P. (2011). Clarity about concessive knowledge attributions: Reply to Dodd. *Synthese*, 181(3), 395–403.
- Dutant, J. (forthcoming). Knowledge-first evidentialism about rationality. In F. Dorsch, & J. Dutant (Eds.), *The new evil demon problem*. Oxford University Press.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.
- Fantl, J., & McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3), 558–589.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Fantl, J., & McGrath, M. (2012). Précis of Knowledge in an Uncertain World. *Philosophy and Phenomenological Research*, 85(2), 441–446.
- Fantl, J., & McGrath, M. (2019). Clarifying pragmatic encroachment: A reply to Charity Anderson and John Hawthorne on knowledge, practical adequacy, and stakes. *Oxford Studies in Epistemology*, 6, 258–266.
- Fassio, D. (2018). Moderate skeptical invariantism. *Erkenntnis*, 85(4), 841–870.
- Fassio, D., & Gao, J. (2021). Do we really need a knowledge-based decision theory? *Synthese* 1–29.
- Fritz, J. (2022). Knowledge and the many norms on action. *Erkenntnis*, 87(3), 1191–1210.
- Gao, J. (2019a). Against the iterated knowledge account of high-stakes cases. *Episteme*, 16(1), 92–107.
- Gao, J. (2019b). Credal pragmatism. *Philosophical Studies*, 176(6), 1595–1617.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178(3), 529–547.
- Goldschmidt, Z. (forthcoming). Foundations for knowledge-based decision theories. *Australasian Journal of Philosophy*.
- Goodman, J., & Salow, B. (2018). Taking a chance on kk. *Philosophical Studies*, 175(1), 183–196.
- Goodman, J., & Salow, B. (2023). Epistemology normalized. *Philosophical Review*, 132(1), 89–145.
- Greco, D. (2013). Probability and prodigality. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 82–105).
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Heil, R. (Unpublished). Knowledge and rational action. Manuscript in preparation.

- Holton, R. (1994). Deciding to trust, coming to believe. *Australasian Journal of Philosophy*, 72(1), 63–76.
- Hyman, J. (2010). The road to larissa. *Ratio*, 23(4), 393–414.
- Ichikawa, J. J. (2012). Knowledge norms and acting well. *Thought: A Journal of Philosophy*, 1(1), 49–55.
- Ichikawa, J. J. (2017). Basic knowledge first. *Episteme*, 14(3), 343–361.
- Jackson, A. (2012). Two ways to put knowledge first. *Australasian Journal of Philosophy*, 90(2), 353–369.
- Jackson, E. (2019a). How belief-credence dualism explains away pragmatic encroachment. *Philosophical Quarterly*, 69(276), 511–533.
- Jackson, E. G. (2019b). Belief and credence: Why the attitude-type matters. *Philosophical Studies*, 176(9), 2477–2496.
- Jeffrey, R. (1970). Dracula meets Wolfman: Acceptance vs. partial belief. In M. Swain (Eds.), *Induction, acceptance, and rational belief* (pp. 157–185). D. Reidel.
- Jeffrey, R. (2004). *Subjective probability: The real thing*. Cambridge University Press.
- Jeffrey, R. C. (1983). *The logic of decision* (2nd ed.). University of Chicago Press.
- Lance, M. N. (1995). Subjective probability and acceptance. *Philosophical Studies*, 77(1), 147–179.
- Levi, I. (1967). Probability kinematics. *British Journal for the Philosophy of Science*, 18(3), 197–209.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Littlejohn, C. (2011). Evidence and knowledge. *Erkenntnis*, 74(2), 241–262.
- Locke, D. (2014). Knowledge norms and assessing them well. *Thought: A Journal of Philosophy*, 3(1), 80–89.
- Lord, E. (2018). *The importance of being rational*. Oxford University Press.
- Mehta, N. (2016). Knowledge and other norms for assertion, action, and belief: A teleological account. *Philosophy and Phenomenological Research*, 93(3), 681–705.
- Moon, A., & Jackson, E. (2020). Credence: A belief-first approach. *Canadian Journal of Philosophy*, 50(5), 652–669.
- Moss, S. (2013). Epistemology formalized. *Philosophical Review*, 122(1), 1–43.
- Moss, S. (2018). *Probabilistic knowledge*. Oxford University Press.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294.
- Neta, R. (2009). Treating something as a reason for action. *Noûs*, 43(4), 684–699.
- Osherson, D., & Zhao, J. (2010). Updating beliefs in light of uncertain evidence: Descriptive assessment of Jeffrey's rule. *Thinking and Reasoning*, 16(4), 288–307.
- Palmira, M. (2023). Reasoning simplifying attitudes. *Episteme*, 20(3), 722–735.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, Duncan, T. J., & Carter, J. A. (2018). The value of knowledge. In E. N. Zalta (Eds.), *The stanford encyclopedia of philosophy*, spring 2018 edition. Metaphysics Research Lab, Stanford University.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2), 143–157.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224–244.
- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs*, 52(1), 171–195.
- Ross, J., & Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2), 259–288.
- Rubin, K. (2015). Total pragmatic encroachment and epistemic permissiveness. *Pacific Philosophical Quarterly*, 96(1), 12–38.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–919.
- Schiffer, S. (2007). Interest-relative invariantism. *Philosophy and Phenomenological Research*, 75(1), 188–195.
- Schroeder, M. (2018). Rational stability under pragmatic encroachment. *Episteme*, 15(3), 297–312.
- Schulz, M. (2017). Decisions and higher-order knowledge. *Noûs*, 51(3), 463–483.
- Schulz, M. (2021a). Partial reliance. *Canadian Journal of Philosophy*, 51(6), 436–451.
- Schulz, M. (2021b). Practical reasoning and degrees of outright belief. *Synthese*, 199(3–4), 8069–8090.
- Schwan, B., & Stern, R. (2017). A causal understanding of when and when not to jeffrey conditionalize. *Philosophers' Imprint*, 17(8), 1–21.
- Simion, M. (2021). Knowledge and reasoning. *Synthese*, 199(3), 10371–10388.
- Smithies, D. (2012). The normative role of knowledge. *Noûs*, 46(2), 265–288.
- Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, 13(s13), 137–49.
- Staffel, J. (2019). How do beliefs simplify reasoning? *Noûs*, 53(4), 937–962.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.
- Tang, W. H. (2015). Belief and cognitive limitations. *Philosophical Studies*, 172(1), 249–260.

- Weatherson, B. (2012). Knowledge, bets and interests. In J. Brown & M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 75–103). Oxford University Press.
- Wedgwood, R. (2012). Outright belief. *Dialectica*, 66(3), 309–329.
- Weisberg, J. (2013). Knowledge in action. *Philosophers' Imprint*, 13(22), 1–23.
- Weisberg, J. (2020). Belief in psyontology. *Philosophers' Imprint*, 20(11), 1–27.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Williamson, T. (2005a). Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *Philosophical Quarterly*, 55(219), 213–235.
- Williamson, T. (2005b). Knowledge, context, and the agent's point of view. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 91–114). Oxford University Press.
- Williamson, T. (2005c). Replies to commentators. *Philosophy and Phenomenological Research*, 70(2), 468–491.
- Williamson, T. (2017). Acting on knowledge. In J. A. Carter, E. C. Gordon, & B. W. Jarvis (Eds.), *Knowledge first: Approaches in epistemology and mind* (pp. 163–181). Oxford University Press.
- Worsnip, A. (2015). Possibly false knowledge. *Journal of Philosophy*, 112(5), 225–246.

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