


ARTICLE

## Partial Reliance

Moritz Schulz 

Department of Philosophy, Technische Universität Dresden, Dresden, Germany  
Email: moritz.schulz@tu-dresden.de

### Abstract

According to a prominent thought, in one's practical reasoning one should rely only on what one knows (Fantl and McGrath 2002; Hawthorne and Stanley 2008; Williamson 2000, 2005a, 2017). Yet for many choices, the relevant information is uncertain. This has led Schiffer (2007, 189) to the following objection: oftentimes, we are fully rational in reasoning from uncertain premises which we do not know. For example, we may decide to take an umbrella based on a 0.4 credence that it will rain. There are various ways proponents of a knowledge norm for practical reasoning can respond. One option is to say that the right way of dealing with uncertain information requires knowledge of probabilities (Hawthorne and Stanley 2008, 581–85). Another option is to say that credences can be knowledge because they really are beliefs with an unusual kind of content that consists of a set of probability spaces (Moss 2018, chap. 9). Mixed accounts are possible as well (Weisberg 2013). On neither of the accounts in the literature, however, can reasoning from uncertain premises be taken as reasoning based on a graded attitude, a credence, toward an ordinary proposition. To make room for this possibility, I argue that reliance comes in degrees. The knowledge norm is only plausible when taken to be concerned with *full reliance*. *Partial reliance*, on the other hand, goes hand in hand with credence.

**Keywords:** Knowledge; practical reasoning; uncertainty; knowledge norm; credence

Over the last two decades, various philosophers have proposed a close connection between knowledge and action, including but not limited to Jeremy Fantl and Matthew McGrath (2002), John Hawthorne and Jason Stanley (2008), and Timothy Williamson (2000, 2005a, 2017). It should not come as a surprise that knowledge and action are intimately linked. After all, the human search for knowledge plausibly stems in large parts from the need to make good decisions.

One aspect of the connection between knowledge and action concerns practical reasoning. When should one rely on a proposition in one's reasoning about which action to perform? Williamson, as well as others, suggests that one needs to know for proper reliance:

There is an intuitively attractive link between knowledge and practical reasoning. If it is not true for me to say "I know that my door is locked," something is wrong with practical reasoning in which I rely on the premise that my door is locked, for in some sense I am not entitled to that premise, however much it seems to me that I am. (Williamson 2005a, 227)

Abstracting from the fact that Williamson is in this passage concerned with knowledge ascriptions rather than knowledge itself, we may capture his proposal in the following way:<sup>1</sup>

**KN:** It is appropriate to rely on  $p$  in practical reasoning iff one knows that  $p$ .

<sup>1</sup>As Williamson (2005a) defends a nonsensitive invariantism about knowledge ascriptions, descending from knowledge ascriptions to knowledge itself remains faithful to his account.

A knowledge norm of this kind answers the question of which epistemic status a proposition must enjoy in order to be appropriately relied upon in practical reasoning. According to **KN**, knowledge is the required epistemic status. To vary this thought, knowledge is the entry ticket which entitles a proposition to take part in the game of practical deliberation and provides the agent with a warrant to reason from this proposition.<sup>2</sup>

Note that ‘practical reasoning’ is used in a generic sense in **KN**. The right-to-left direction of **KN** merely states that if one knows  $p$ , one *can* appropriately rely on  $p$ , not that one always has to. Hence, **KN** does justice to the fact that, for most choices, only part of one’s knowledge will be relevant.<sup>3</sup>

The problem I wish to address in this paper concerns decision-making with uncertain propositions.<sup>4</sup> Right after the knowledge norm was proposed, there was an unease about how a norm like **KN** is supposed to relate to decision-making under uncertainty, in particular to Bayesian decision theory.<sup>5</sup> Stephen Schiffer (2007) expresses this worry very clearly (in the following quote ‘IRI’ refers to *interest-relative invariantism*):

I think that the [...] thesis that one should act only on what one knows isn’t plausible. The following sort of example is very common: you are completely justified in carrying an umbrella even though you don’t know that it will rain but merely believe to degree 0.4 that it will rain. Not only is this a prima facie counterexample to IRI, insofar as it appears to be an example in which one is justified in acting as one does on the basis not of knowledge but of a justified partial belief, but there is a familiar and widely-accepted Bayesian account of why one’s action is justified, even though one isn’t acting on knowledge: one is performing that action which has the greatest expected desirability. (Schiffer 2007, 189)

Schiffer describes what appears to be a fairly straightforward counterexample to the knowledge norm for practical reasoning (and its variants). I don’t know that it will rain, so the right-hand side of **KN**’s biconditional is false. On the other hand, it may seem that relying in my decision-making on a 0.4 credence that it will rain is appropriate, which would show the left-hand side of **KN** to be true.<sup>6</sup> I do not need to know that it will rain for deciding to take an umbrella. Carrying umbrellas is often a precautionary act by which we insure ourselves against rain: we pay a small price (carrying extra weight, etc.) which will pay off if it rains.

Schiffer also mentions what appears to be a well-entrenched contender to the knowledge norm: *Bayesian decision theory*. Expected utility maximization is well-equipped to explain the agent’s decision-making in Schiffer’s umbrella case. The agent decides to take an umbrella because it has a higher expected utility than not taking it.

<sup>2</sup>A number of variants of **KN** have been proposed. Most prominent is perhaps the proposal by Hawthorne and Stanley (2008) which connects knowledge to reasons for action. According to Hawthorne and Stanley, a proposition is only appropriately treated as a reason for action if it is known. Related proposals were made by Hyman (1999) and Unger (1975, 200–201). Hyman and Unger defend the position that something can only be a reason for an agent if it is known. Recent work on *reasons possession*—according to which one only possesses a reason if one knows it (or is in a position to know it)—varies this thought further (see Lord 2018 and Kiesewetter 2017). In this paper, I focus almost exclusively on **KN**. But I say something about how the results of this paper apply to alternative formulations of **KN** in my concluding remarks.

<sup>3</sup>See Hawthorne and Stanley (2008) for discussion. They explicitly restrict a norm similar to **KN** to *relevant propositions* because their norm targets reasons for *individual actions*.

<sup>4</sup>See Staffel (2013) for an account of reasoning with uncertain propositions (and a defense that this is indeed possible).

<sup>5</sup>Knowledge norms like **KN** also face other problems which I have to set aside for the purpose of this paper. Perhaps the most prominent kind of objection involves high-stakes decision situations. See Brown (2008), Gao (2017, 2019), Greco (2013), and Gerken (2011). For some salient responses, see Schulz (2017, 2021a) and Williamson (2005a). In order to circumvent the issue of high-stakes decision-making, I shall assume throughout that stakes are low.

<sup>6</sup>I sometimes speak loosely about ‘relying on a credence’ where I should strictly say ‘relying on the proposition which is the object of the credence.’ Similarly, I sometimes say that one ‘relies on a (outright) belief’ where I should strictly say that one ‘relies on the propositional content of the belief.’

We may extract from Schiffer's considerations the following argument against **KN**, the knowledge norm for practical reasoning. Call this the *problem of uncertainty* for **KN**:

- (1) It is sometimes appropriate to rely on uncertain propositions in practical reasoning.
- (2) At least in some such case, the uncertain propositions are not known.
- (3) Therefore, the knowledge norm is false.

The first premise establishes cases involving uncertain propositions in which the left-hand side of the knowledge norm is true, while the second premise expresses that in at least some such cases, the right-hand side of the knowledge norm is false. (In this reconstruction, a proposition is called 'uncertain' if the salient agent assigns it a nonextreme credence.)

Various responses to the problem of uncertainty can be sorted by asking whether they reject the first or the second premise of the argument above. Hawthorne and Stanley (2008, 581–85) reject the first premise. According to them, the agent in Schiffer's example, if fully rational, will act on her *knowledge of (epistemic) chances*, i.e., that the chances of rain are 40 percent.<sup>7</sup> According to this response, one should never rely on uncertain propositions. Cases in which it appears as if the agent appropriately relies on an unknown proposition are in fact cases in which the agent relies on a known proposition about chances.

In contrast, Sarah Moss (2018, chap. 9) makes a radically different proposal. On her theory of *probabilistic knowledge*, uncertain propositions like the one in Schiffer's example can be known. However, probabilistic knowledge has, on Moss's account, a novel kind of content: probabilistic knowledge is knowledge whose content is a set of *probability spaces*. Very roughly, knowledge that it is 0.4 likely that it will rain would be knowledge whose content is a set of probability spaces in which rain is assigned a probability of 0.4. The account by Moss puts one in a position to reject the second premise in the argument above: uncertain propositions can be known, yet they have a very different kind of content than one might have thought.

Jonathan Weisberg (2013) offers a two-part story which combines aspects of the previous two proposals. The first part is a proposal similar to Moss's, which would allow Weisberg to reject the second premise for a certain class of cases. As opposed to Moss, however, Weisberg does not go in for probabilistic contents. Rather, he suggests that credences can sometimes *constitute* beliefs about chances; and when they do, they can constitute knowledge in virtue of the corresponding belief being knowledge. This knowledge would be the same kind of knowledge—knowledge about epistemic chances—as proposed by Hawthorne and Stanley (2008). The second part of Weisberg's proposal is designed for cases in which one's credences do not constitute the relevant kind of knowledge. In such cases, Weisberg contends, the agent must know that it *may* rain. Depending on how likely the relevant possibility is, this knowledge about possibilities is then to be given different *weight* in one's reasoning. Thus, credences can also function as "weight-givers" in one's reasoning with possibilities.

Among the proposed solutions, a certain type of response to the problem of uncertainty for the knowledge norm seems to be missing. On all existing solutions, some aspects of decisions under uncertainty require an unconventional story. On no account can the agent's decision-making in a case like Schiffer's be taken at face value: (a) the agent relies on a partial belief, a credence, which is (b) a graded attitude toward an ordinary proposition. Hawthorne and Stanley (2008), and to some extent also Weisberg (2013), explain the (a) aspect away. It is not a graded attitude the agent should

<sup>7</sup>Hawthorne and Stanley (2008, 584–85) also consider a further possibility which requires knowledge of those propositions which ultimately ground a credence. For a response to this option, see Jackson (2012). Here it suffices to note that similar to the other accounts mentioned in the main text, such a proposal cannot take Schiffer's example at face value. One would have to say that what appears to be rational reliance on the content of a credence is in fact reliance on those aspects of one's knowledge which ground the credence. When in what follows I speak of "Hawthorne and Stanley's proposal," I refer to their proposal in terms of epistemic chances and not to the proposal in terms of grounding discussed in this footnote.

rely on, but rather an outright attitude toward a modal proposition, either about chances (Hawthorne and Stanley 2008) or about possibilities (Weisberg 2013). Moss (2018), on the other hand, works around the (b) aspect by taking good decision-making to be based on an outright attitude (knowledge) toward a novel kind of content (sets of probability spaces). The goal of this paper is to find a way of doing justice to both aspects of Schiffer's case without giving up on the knowledge norm for practical reasoning.

A view of the envisaged kind would have certain advantages over its rivals.<sup>8</sup> First, it can do justice to the *normative significance* of credences (in virtue of satisfying the [a] aspect). Consider a case in which an agent appears to act well as a result of relying on a partial belief. On the envisaged account, one can take such a case at face value. Yet on some of the alternative accounts, most notably the one by Hawthorne and Stanley (2008), relying on a credence can never be normatively proper. Either the case is not really one of acting well or the agent does not rely on a credence but rather on corresponding outright belief whose content is a modal proposition about chances. Second, sticking to the existing accounts seems to require an assumption of *conceptual availability*. For instance, Schiffer (2007, 190) remarks: "[...] Jane may be justified in acting on a partial belief even if she lacks the concept of the relevant epistemic probability and doesn't even believe, let alone know, that there's a 'non-negligible chance' that it will rain." If one does not possess the concept of probability, one cannot form beliefs whose content contains the concept of probability. But of course, one may have a certain credence without possessing the concept of probability.

The account by Moss may remain unaffected by the above considerations. But it has other problems. For instance, it is a vexed question whether there is a robust sense in which knowledge would remain *factive* (cf. Easwaran 2018; Moss 2018, ch. 6; Rich 2020). If the content of some of our knowledge is a set of probability spaces that is irreducible to an ordinary proposition, in what sense is what we know still true? There are, of course, ways of addressing this question (see Moss 2018, Ch. 6; and Rich 2020), but it may seem that from the perspective of theory choice, being able to do without a novel kind of content would be a simpler and in any case more conservative option. This makes it, in my mind, worthwhile to explore an alternative approach.

The proposal I would like to make is that reliance comes in degrees. Thus, in addition to fully relying on a proposition, there is also a sense in which one might only be partially relying on a proposition. Given these two senses, there is an interesting ambiguity in one of the knowledge norm's central phrases, namely in the phrase "to rely on *p* in one's practical reasoning." On one construal, the norm is concerned with "*fully* relying on *p* in one's practical reasoning." On another, broader construal, it is concerned with "*partially* relying on *p* in one's practical reasoning." I shall argue, based on an observation by Jackson (2012), that knowledge plausibly governs full reliance but not partial reliance. As Schiffer's example is arguably a case of partial reliance, it is not a counterexample to a knowledge norm for practical reasoning cast in terms of full reliance.

Saving a knowledge norm for practical reasoning by making explicit that it only governs full reliance raises a number of important questions. What does it mean to partially rely on a proposition? Is partial reliance governed by any kind of norm which would supplement the knowledge norm for full reliance? What would such a norm have to do with knowledge? In this paper, I hope to make progress on these questions, although I will not be able to answer all of them completely.

This is how we shall proceed. I start out in [section 1](#) by sketching one way of defending the knowledge norm. [Section 2](#) introduces partial reliance. [Section 3](#) draws a close parallel between outright belief and full reliance on the one hand and partial reliance and credence on the other. [Section 4](#) discusses the prospects of unifying the resulting picture. [Section 5](#) concludes.

<sup>8</sup>See Jackson (2012) and Müller and Ross (2017) for further discussion.

### 1. A limited norm

When considering Schiffer's objection to the knowledge norm for action, Alexander Jackson (2012) responds as follows (where 'AKP' is a variant of the knowledge norm **KN**):

The point is clear once we consider whether the case is a counterexample to AKP. Roughly, to treat  $p$  as a reason for action is to have a full belief that  $p$ . Jane's reason for carrying an umbrella is not that it will rain. It may well be that if Jane did take that proposition to be available as a reason for action, i.e., if she believed that it will rain, she would decide not to leave the house, rather than go out but take an umbrella [Hawthorne and Stanley 2008, 582] AKP gets things exactly right: Jane does not know that it will rain, and should not take it as a reason for action that it will rain. (Jackson 2012, 357)

AKP states "Treat the proposition that  $p$  as a reason for acting only if you know that  $p$ " (Jackson 2012, 354). Central to Jackson's point is the claim that treating a proposition as a reason for action is closely linked to having a full belief. The observation Jackson makes is now very straightforward. In Schiffer's case, the agent does not fully believe that it will rain. Rather, she merely has a 0.4 credence that it will rain. But if treating something as a reason for action requires fully believing the proposition in question, then the agent in Schiffer's case does not treat the proposition that it will rain as a reason for action. There is then no counterexample to the knowledge norm. A counterexample of the desired kind would require a case in which the agent appropriately treats a proposition as a reason for action without knowing it. What we would have instead is a case in which the agent does not know the proposition and does not treat it as a reason for action.

On Jackson's resolution, the knowledge norm for action is closely connected to, and possibly even derivable from, the knowledge norm for full belief (more on this in section 3). On this latter norm, one should fully believe a proposition  $p$  only if one knows  $p$ . It is instructive to note that there is no similar problem of uncertainty regarding the knowledge norm for full belief (Jackson 2012, 357). It is clear that credences (aka *partial beliefs*) are outside the scope of the knowledge norm for full belief. It would be a nonstarter to suggest that one should have a credence of 0.4 that it will rain only if one knows that it will rain. In a similar vein, one should hold, according to Jackson (2012), that acting on uncertain propositions is outside of the scope of the knowledge norm for practical reasoning.

Crucial to Jackson's point is the claim that treating a proposition as a reason for action requires full belief. But one may wonder what we are then doing when we decide to take an umbrella based on a 0.4 credence that it will rain. It still seems that our epistemic stance toward the proposition that it will rain is treated as counting in favor of taking an umbrella. So, at least *prima facie*, there seems to be a sense in which the proposition *is* treated as a reason for action.

I think there is an interesting ambiguity to be considered here. This ambiguity comes out very clearly in the variant of the knowledge norm for practical reasoning I wish to focus on in this paper. The phrase "relying on a proposition in one's practical reasoning" can be understood in a narrow sense where it is meant to express *fully* relying on a proposition. If we make this reading explicit, we arrive at the following norm:

**KNF:** It is appropriate to *fully rely* on  $p$  in one's practical reasoning iff one knows that  $p$ .

When one fully relies on a proposition, one treats the proposition as a fact. When treated as a fact (or as true), an agent evaluates her alternatives, as Ross and Schroeder (2014, 264) put it, "by the same procedure by which she would evaluate them conditional on  $p$ ." In other words, one does not reckon in one's practical reasoning with the possibility that the relevant proposition might be false.

In contrast, one can, it seems, merely partially rely on a proposition in one's reasoning. When one does so, one does not treat the proposition as a fact. In other words, one does reckon in one's practical reasoning with the possibility that it might be false.

We could, of course, ask whether partial reliance is also governed by knowledge:

**KNP:** It is appropriate to *partially rely* on  $p$  in one's practical reasoning iff one knows that  $p$ .

However, it seems clear that Schiffer's case is a counterexample to this norm. In Schiffer's case, it is plausible that the agent partially (but not fully) relies on the proposition that it will rain. And it is clear that the agent does not know that it will rain. So **KNP** is clearly false. Note, though, that this does not affect **KNF**, for in Schiffer's case the agent does not fully rely on the proposition in question.

In what follows, I shall try to develop an account of partial reliance. On this account, partial reliance is a *sui generis* mental act. Some existing solutions to the problem of uncertainty can be seen as trying to account for partial reliance in terms of full reliance on a "partial" proposition, that is a proposition about chances (Hawthorne and Stanley 2008) or possibilities (Weisberg 2013), or a proposition which itself consists of probability spaces (Moss 2018). Instead, the account I wish to develop qualifies the *attitude* rather than the *content*.

## 2. Partial reliance

In this section, I explain what an account of graded reliance is intended to deliver. I start with a brief explication of the notion of reliance. I will then describe the problem of coming up with a norm governing partial reliance.

According to Williamson (2000, 99), relying on a proposition  $p$  is to use  $p$  as a premise in practical reasoning. Following Williamson, I shall therefore construe reliance as a *mental act*. On this construal, when one relies on  $p$ , one assigns  $p$  a certain role in an episode of practical reasoning. An act-sense of reliance contrasts with a state-sense of reliance as proposed by Facundo Alonso (2014, 165). According to Alonso, reliance is a *mental state* similar in kind to a belief state. As I see it, these differences are largely terminological, for a state-sense of reliance seems to be derivable from an act-sense. One may say that one relies in the state-sense iff one is willing to rely in the act-sense under appropriate circumstances (see also section 3). On this way of setting things up, the state-sense of reliance relates to the act-sense similar to how a disposition relates to its manifestation. As a matter of fact, the state-sense of reliance may simply describe a disposition to rely in the act-sense of reliance under certain conditions.<sup>9</sup>

The philosopher's concept of relying on a proposition in one's reasoning does not seem to be deeply rooted in natural language. Nevertheless, sentences such as "I relied on the information provided by the *Guardian*" or "The court relied on the fact that the defendant had access to the building" seem to express something which comes at least close to it. To the extent that the everyday concept of reliance informs the philosophical notion, it is interesting to observe that the everyday concept is gradable. One can say things like "I relied more on the information provided by the *Guardian* than on the information provided by the *Sun*." In this case, what is expressed seems to be that the propositions communicated by the *Guardian* are treated by me as a better basis for further decision-making and reasoning than the propositions published in the *Sun*.

However, comparative talk about reliance can also trace the strength of reasons rather than their epistemic standing. For example, a statement to the effect that in a circumstantial case the court relied more on the fact that the defendant had blood on her hands than on the fact that she had motive can naturally be interpreted as saying that the court took the blood on her hands to be a

<sup>9</sup>On both the act-sense as well as the state-sense, reliance is a certain kind of *relation* (between [at least] an agent and a proposition). In this regard, the present terminology is compatible with the way Holton defines trust. There are differences concerning the assumed relata, though. For Holton trust is "a three-place relation holding between two people and an action (1994, 70)." See Alonso (2014, 179n5) for related discussion.

stronger reason for conviction than the fact that she had motive. This can be so even when the two facts have the same epistemic standing from the standpoint of the court.

In the light of this flexibility, developing a concept of partial reliance is in my eyes best seen as a project of *explication* in the sense of Rudolf Carnap ([1950] 1962, chap. 1): one takes a concept used in natural language but shapes it further based on theoretical considerations. With respect to the notion of partial reliance, I shall therefore assume that it has some basis in natural language without presupposing that there is a perfect match. In particular, I shall ignore that ‘reliance’ in natural language sometimes traces the strength of reasons rather than their epistemic standing.

With this in mind, what can we say about full and partial reliance? Let us take the somewhat easier part of this question first. To fully rely on a proposition  $p$  in one’s practical reasoning can plausibly be taken to ignore  $\neg p$ -possibilities for the purposes of solving a decision problem. For instance, if I rely fully on the proposition that my lottery ticket lost, then I might decide to throw it away (cf. Hawthorne and Stanley 2008, 572). I would not throw it away if I still reckoned with the possibility that it might have won. Full reliance thus requires a kind of *practical certainty* that the proposition in question is true.<sup>10</sup> Proponents of a knowledge norm think that knowledge provides the relevant standard for being practically certain. For instance, Brian Weatherson (2012) argues that a possibility can be legitimately excluded from one’s decision-making iff one knows that it is not actual.<sup>11</sup>

The somewhat harder task is to say what it means to partially rely on a proposition in one’s reasoning. In contrast to full reliance, partial reliance on a proposition  $p$  means to take  $p$  into account while solving one’s decision problem without ignoring  $\neg p$ -possibilities.<sup>12</sup> Since reliance is gradable, one can partially rely on  $p$  to a low, an intermediate, or a high degree. Relying on  $p$  to a high degree can plausibly be understood as putting a lot of weight on  $p$  in one’s decision-making while putting little weight on  $\neg p$ . Relying on  $p$  to an intermediate degree would mean putting roughly equal weight on  $p$  and on  $\neg p$ . And relying on  $p$  to a low degree would mean to putting little weight on  $p$  and a lot of weight on  $\neg p$ .<sup>13</sup>

To illustrate, consider again the lottery example. Given the high likelihood of one’s ticket having lost, one would give this proposition a lot of weight in one’s reasoning. Assuming, however, that one is not certain, one will also put some weight on the proposition’s negation, albeit very little. This latter fact is responsible for why one’s reliance on holding a losing ticket still counts as partial. Once one’s reliance is partial, one will typically not decide to throw the lottery ticket away, for this response typically requires that one can fully rely on not holding a winner.

The present idea moves the proposal close to one part of Weisberg’s account. He writes:

A subject with 0.4 credence in rain and another with 0.1 credence might share a reason for bringing an umbrella, namely that it may rain. But only the first (we may suppose) should end up bringing one. What cuts the difference is the weight each subject should give to that reason.

<sup>10</sup>Wedgwood (2012) distinguishes between practical and theoretical certainty and argues that decision-making requires one to be practically certain of various propositions.

<sup>11</sup>But see also Beddor (2020) for a critical examination of the knowledge norm in light of the idea that *certainty* plays an important role in one’s decision-making. Ultimately, Beddor argues that a context-dependent notion of certainty instead of knowledge is what should guide one’s practical reasoning. Locke (2015) takes a similar stand: it is rationally permissible to treat a proposition as true in one’s practical reasoning when the proposition is *practically certain*, where practical certainty results from epistemic certainty by ignoring practically irrelevant alternatives.

<sup>12</sup>As it is possible to partially rely to some degree on a proposition while simultaneously to partially rely on the proposition’s negation, partial reliance should not be construed as adopting the proposition as a premise in a deductive model of practical reasoning. Otherwise partial reliance would commit one to reason from an inconsistent set. Adopting a proposition as a premise should be reserved for full reliance only. Thanks to an anonymous reviewer for pointing this out.

<sup>13</sup>To guard against a possible misunderstanding, putting weight on a proposition in one’s reasoning does not require that one takes this proposition to strongly favor a certain response. I can put a lot of weight on the proposition that tomorrow will be a sunny day while taking this proposition only to weakly favor wearing my blue sweater.

For the first subject, the possibility of rain carries enough weight to overcome the reasons against carrying an umbrella. But not so for the 0.1 subject. So one role for credences in practical reasoning, I propose, is to determine how much weight we should give to certain reasons. (Weisberg 2013, 13)

Weisberg goes on to suggest that weighing reasons can be understood along the lines of expected value calculation. An expected value is, after all, a weighted average of one's utilities, where the weights are supplied by one's credences.

One difference between the present account and Weisberg's is that Weisberg takes the relevant reason to be a proposition about what is possible, i.e., that it *may* rain. Such modal reasons are then to be given different weight in one's reasoning. On the model I wish to develop, one partially relies on whatever is the embedded proposition. That is, one would partially rely on the proposition that it will rain and give this proposition different weight in one's reasoning. Of course, one could also (fully) rely on the proposition that it may rain, but that would not be an essential ingredient of the account I favor. Note that the analogy with expected value calculation is more direct on the present proposal. For what is given different weight in the calculation of expected value is not a modal proposition of the form 'may ( $p$ )' but rather the embedded proposition  $p$ , to which the corresponding credence attaches.

Before moving forward, let me pause and consider a possible objection.<sup>14</sup> Could it be that what I classify as degrees of reliance merely reflects a context-relativity of reliance? For instance, Alonso (2014, 177–78) suggests that one may be prepared to rely on a proposition in one context but not in another (see also Alonso 2016). The example Alonso presents is one in which an agent is prepared to rely on a rope holding her weight when deciding to climb a small rock but not when deciding to climb a big rock. In this example, the agent has some reasonable doubts about the rope's strength, which is why relying on it in climbing the big rock would mean risking her life, at least from the agent's own perspective.<sup>15</sup> Once such a context-relativity is granted, one could say that partial reliance occurs when one is disposed to rely on the proposition in question only in some but not all relevant contexts. One may go on to propose that one is disposed to rely on one proposition  $p$  more than on another proposition  $q$  iff for the set of contexts in which both propositions are relevant the set in which one is disposed to rely on  $q$  is a proper subset of the set of contexts in which one is prepared to rely on  $p$ .<sup>16</sup>

Although I agree that this way of construing reliance would give us some kind of gradability, it would still leave a certain type of gradability unaccounted for. The problem is that there can be *intra-contextual* differences concerning reliance. In other words, there can be single contexts with only one decision problem to solve in which I rely on one proposition more than another. Following up on the example above, I may only partially rely on the proposition that the rope will hold my weight when deciding whether to climb the big rock, but I may simultaneously fully rely on the proposition that the rock has at least a certain altitude (which grounds my concerns about what I stand to lose should I fall). These differences would show, for instance, by my assigning different probabilities to the respective propositions when calculating the expected value of potential actions. Thus, it seems there can be differences in reliance even in a single context which are not plausibly described in terms of context-relativity. (I continue discussing the rope example on in section 3.)

<sup>14</sup>I thank an anonymous reviewer for pressing me on this.

<sup>15</sup>In the present version of the example, the agent has a nonextreme credence in the relevant proposition and neither knows nor believes outright that it is true. We get a challenge of a different kind if we change this aspect: even when known, the proposition may under certain circumstances not be suitable for full reliance, namely when the stakes are extremely high (see also fn. 5). To deal with such cases, I elsewhere (Schulz 2021a) suggest that reliance may also be gradable along another dimension: just like there may be degrees of outright belief on top of credences, there may be degrees of full reliance, too.

<sup>16</sup>As two sets may overlap without standing in a subset relation, the present proposal would only define a partial order. I will not rest my case against this proposal on this consequence, though.



With a notion of partial reliance in place, we can ask whether it is norm governed in any way. Very roughly, finding a norm which governs partial reliance would amount to solving the following equation for ' $\phi$ ':

**PN?**: It is appropriate to *partially rely* on  $p$  in one's practical reasoning iff  $\phi(p)$ .

To give an example of what might solve this equation, let us go back to how Schiffer (2007, 189) introduces the problem of uncertainty for the knowledge norm. Next to his critical remarks, there is the positive suggestion that agents rely on a "justified partial belief" in decisions under uncertainty. This suggests to consider the following norm: it is appropriate to *partially rely* on  $p$  in one's practical reasoning iff one has a justified partial belief in  $p$ . There will be more to say on a norm for partial reliance. For the time being, it suffices to see how one may use **PN?** as a schema to search for a norm governing partial reliance. The next section discusses how one may simplify this search.

### 3. Reliance and belief

I think progress can be made regarding our main question by clarifying the relation between full reliance and full belief on the one hand and partial reliance and partial belief on the other.

There seems to be a tight connection between belief and practical reasoning: "Intuitively, one believes  $p$  outright when one is willing to use  $p$  as a premise in practical reasoning" (Williamson 2000, 99). If we want to be sensitive to the distinction between full and partial reliance, we may capture the proposed relation between full belief and reliance as follows:

**BP**: One fully believes  $p$  iff one is willing to fully rely on  $p$  in one's practical reasoning.

For present purposes, there is no need to decide whether this principle can serve as a definition of belief or whether it merely describes something closer to the functional role of belief for practical reasoning.<sup>17</sup>

With **BP** in place, it suggests itself to consider the mirror image of **BP** for partial belief:

**PP**: One partially believes that  $p$  iff one is willing to partially rely on  $p$  in one's practical reasoning.

On this proposal, to have a partial belief would mean to be willing to rely on  $p$  to some extent in one's practical reasoning. It would mean to put some weight on  $p$  in one's practical deliberation that is proportional to the degree to which one believes  $p$ . If we avail ourselves of a more fine-grained notion of credences and degrees of reliance, then we may put a more fine-grained principle in place of **PP**:

**PPx**: One has credence  $x$  in  $p$  iff one is willing to rely on  $p$  to degree  $x$  in one's practical reasoning.

To have a certain credence in  $p$ , on this view, is to be prepared to put a corresponding weight on  $p$  in one's practical reasoning.

On a reliance-based account of credence and belief as presented so far, there is a continuity between credences and full belief which mirrors a corresponding continuity between partial reliance and full reliance. As higher credences go with higher degrees of reliance and full belief goes with full reliance, full belief is a kind of practical certainty which implies credence 1. It bears emphasizing

<sup>17</sup>See Weatherston (2005) and Ross and Schroeder (2014) for further discussion. For instance, Ross and Schroeder (2014, 267) hold that "[...] at least part of the functional role of belief is that believing that  $p$  defeasibly disposes the believer to treat  $p$  as true in her reasoning."

that on this picture credence 1 is merely a kind of practical certainty that can be present even when something like Cartesian certainty—certainty beyond any kind of doubt—is absent.<sup>18</sup>

The two principles **BP** and **PP** do not express any kind of norm. Rather, they are metaphysical hypotheses about full and partial belief. For our present project of finding a norm for partial reliance, they can nonetheless provide some helpful clues. We already touched upon the knowledge norm for full belief (Bird 2007; Engel 2004; Smithies 2012; and Williamson 2000, 2005b):

**BN:** One should: fully believe  $p$  iff one knows  $p$ .

This norm is often cast as a mere conditional (“Believe  $p$  only if you know  $p$ ”). But on the assumption that knowledge implies belief, which I shall make for the purposes of this paper (but see Radford 1966 and Myers-Schulz and Schwitzgebel 2013), there is no harm in stating the norm as a biconditional.

One can see this norm for belief as being grounded in the corresponding norm for full reliance. If it is, first, appropriate to fully rely on  $p$  in one’s practical reasoning iff one knows  $p$  and, second, belief is grounded in being willing to fully rely on  $p$ , then it seems to follow, or at least it is strongly suggested, that one should believe  $p$  iff one knows  $p$ .<sup>19</sup> In other words, **BN** holds because of **KNF**: knowledge is the norm for full belief because knowledge is the norm for full reliance and full reliance is what grounds full belief.

As a working hypothesis, one may consider the exact same relationship for partial reliance and credence. One would hypothesize that there is a certain state which governs partial reliance and partial belief alike. This would establish an equivalence between the conditions under which relying on  $p$  is normatively proper and the conditions under which having a corresponding credence in  $p$  is normatively proper:

**Eq:** It is appropriate to partially rely on  $p$  in one’s practical reasoning iff one has a corresponding partial belief that is normatively proper.

In the light of this equivalence thesis, we can make our life somewhat easier by searching for a norm on partial belief. As I think we have a better intuitive grip on the conditions under which one should form a certain credence, this promises to be progress.

Before moving on, let me comment on one general aspect of the norm for partial reliance I am searching for. The knowledge norm on full reliance associates full reliance with a purely epistemic state, knowledge (at least to the extent that knowledge itself is not influenced by pragmatic factors). In a similar vein, the hypothesized norm for partial reliance associates partial reliance with a purely epistemic state, namely with normatively proper credences (at least to the extent that credences themselves are not influenced by pragmatic factors). This general aspects runs against more pragmatic construals of reliance. For instance, Alonso (2014, 2016) takes reliance to be partly governed by pragmatic factors. Thus, to continue with an example mentioned earlier, one may rely on a rope being strong enough to hold one’s weight when deciding to climb a small rock but not when deciding to climb a big rock. As reliance can come with various risks, a less than perfectly justified proposition may be sufficiently justified to be relied upon in a low-stakes context while being insufficiently qualified to perform this role in a high-stakes context.

On the notion of reliance I am after, I would describe the situation somewhat differently. If one only has a certain credence that the rope will hold one’s weight, one can still partially rely on this proposition in one’s decision-making. As a result, one may decide to climb the small rock because

<sup>18</sup>The present picture is similar to the account of Wedgwood (2012). For Wedgwood, there are *practical credences* as well as *theoretical credences* and outright belief is defined as practical certainty in normal situations. The credences present in a principle like **PPx** would be practical credences in Wedgwood’s sense, while something like Cartesian certainty would require theoretical certainty.

<sup>19</sup>I defend this norm against the charge of belief being a much weaker attitude in Schulz (2021b).

the associated risks are low and one will enjoy doing so (climbing may, in the situation at hand, have maximal expected value). However, even though one can partially rely on this proposition, this does not mean that one should be inclined to climb the big rock, too. As climbing the big rock has an extremely negative possible outcome, the expected value of not climbing may be higher than the expected value of climbing. For this reason, partial reliance can explain the difference in the decisions one should make without being influenced by pragmatic factors.

Nevertheless, I think there is a well-entrenched notion which is pragmatically influenced in much the same way Alonso construes reliance: the notion of *acceptance*.<sup>20</sup> In a given situation, we may treat certain propositions as if we fully believed them because in this way we can simplify our reasoning and save time and energy. One may even vary the example above to yield a situation in which I do not want to think too hard and too long about the matter of climbing the small rock. For this reason, I decide to accept that the rope will hold my weight and quickly reason to the conclusion that I may use it to climb the rock. Let me say that I think the phenomenon of acceptance is a real one. But it is not the phenomenon I am trying to capture in terms of full and partial reliance.

#### 4. Toward unification

The previous section explained that the conditions for proper partial reliance can be expected to be the same as the conditions for proper partial belief. This parallels a similar relation between full reliance and full belief: the conditions for proper full reliance are the same as the conditions for proper full belief. If this proposal is correct, then we can get a grip on proper partial reliance by facilitating our understanding of proper credence.

As it stands, the picture developed so far is open to the objection that it is bifurcated in a potentially worrisome way. On the one hand, there is full reliance, which is, just like full belief, governed by knowledge. On the other hand, there is partial reliance, which is governed by whatever governs proper credence. Is there anything that ties the two norms together?

There are, it seems to me, two dimensions along which one would like to unify the picture. The first dimension concerns the *content* of the two norms. If full belief is governed by knowledge, credence should be governed by conditions which stand to credence in a way sufficiently similar to how knowledge stands to full belief. If this can be secured, then the two norms would be unified by making a claim of broadly the same kind. But there is also a second dimension of possible comparison. This dimension concerns the *structure* of how the two norms are related. Ideally, one would like to show that there is a well-balanced division of labor between the two norms without the possibility of conflict.

Let me start by admitting that giving a fully developed account of the nature of normatively proper credences is beyond the scope of the present paper. Fortunately, there is already a rich literature on *justified credence* and *probabilistic knowledge* one can draw from. For example, Moss (2018, chaps. 5.6–5.7) argues that, on her conception of probabilistic knowledge, credences can have almost any of the properties typically ascribed to knowledge, including *safety* or *sensitivity*. More recently, Beddor and Goldstein (2021) develop a more detailed conception of safety for probabilistic knowledge (in a framework of information-sensitive semantics).

Moss as well as Beddor and Goldstein are concerned with probabilistic knowledge of the form ‘*S* knows that probably *p*.’ Nevertheless, it seems possible to adapt certain aspects of their proposals. To a first approximation, one could then say that a credence is safe iff in all nearby worlds *w* in which the credence is formed in the same (or a sufficiently similar) way, it matches (or is sufficiently close to) the epistemic probability for the agent at *w*.<sup>21</sup> In this way, safe credences would be credences

<sup>20</sup>Alonso (2014, 165) explicitly places his work in close proximity to the debate about acceptance.

<sup>21</sup>The bracketed remarks indicate possible variants of the same core idea which are more relaxed. I say a little more on epistemic probability below.

which aptly track the epistemic probabilities without involving a certain element of luck.<sup>22</sup> Of course, this can only be a first approximation to a full account, but it may suffice to indicate that there does not seem to be a principled hurdle to extend the notion of safety (or sensitivity) to credences without adopting probabilistic contents.

There is also no question that credences can be *reliable* as the works by Dunn (2015), Pettigrew (2021) and Tang (2016) indicate. Konek (2016) develops an account of proper credence based on *cognitive skill*, drawing on virtue-theoretic approaches to knowledge. Although far from being conclusive, this is further evidence that for any property potentially characterizing the nature of knowledge beyond true belief, credences can have a similar property. If so, then proper credences could be described in terms of conditions similar to the conditions required by knowledge, whatever those ultimately turn out to be. As far as their content is concerned, this would unify the norms on full belief and credence.

To guard against a possible misunderstanding, let me make explicit that assuming that there are conditions which stand to credence as knowledge stands to full belief does not require credence to be apt for having just any property knowledge might have. There are differences between full belief and credence which are likely to be reflected in differences in whatever normative conditions govern these attitudes. This point is relevant for assessing the fact that knowledge implies truth. It seems very clear that the truth of  $p$  is not a standard of correctness for a credence in  $p$ . A high credence in  $p$  is not *incorrect* just because  $p$  happens to be false (Tang 2016, 65). If one possesses misleading evidence, it rather seems correct to assign  $p$  a high credence even if  $p$  is, as a matter of fact, false. The reason for this seems to be that the function of credences is to gauge indecisive evidence. Consequently, credences between zero and one do not take a stance on the truth or falsity of  $p$ : when challenged by the falsity of  $p$ , it is always appropriate to respond that one did not commit to the truth of  $p$  because one assigned the negation of  $p$  a positive credence as well (Ross and Schroeder 2014, 276).<sup>23</sup>

Let us now turn to the structural question. Can we ensure that the norms for proper full and partial reliance cohere with each other?

To address this question, recall that full reliance has been characterized as requiring practical certainty. Knowledge in turn is supposed to be the normative standard for practical certainty. Therefore, practical certainty is, on this view, proper iff one has knowledge.

The norms for full and partial reliance make contact when partial reliance converges to full reliance and (practical) credences to (practical) certainty. In order to avoid conflict, a norm for proper credence must secure that one should be certain of a proposition  $p$  iff one knows  $p$ .

I think this requirement can be satisfied if we assume that the standard for proper credences is a probability which comes from conditionalizing on what one knows. Conditional on what one knows, one's knowledge will be certain. And conversely, everything made certain by what one knows is something one will at least be in a position to know.<sup>24</sup> Thus, if one requires that the standard for proper credences is partly set by a probability conditional on what one knows, the structure of the two norms would be set up in such a way that they do not conflict where they make contact.

<sup>22</sup>In a similar way, one could envisage a notion of sensitivity for credences which requires that one's credences would have been (significantly) different if the epistemic probability had been (significantly) different.

<sup>23</sup>This does not mean that there is no interesting relation between truth and credence: credences may well try to get as close to the truth as possible within the boundaries of what the evidence permits. See Konek (2016) for relevant discussion in the context of probabilistic knowledge.

<sup>24</sup>I deliberately use 'certain' instead of 'probability 1' in this context, for some propositions with probability 1 are not certain. This is illustrated by the familiar dart board cases. It is not certain that the dart will not hit the bull's eye. Nevertheless, given that the underlying space of possible points is (uncountably) infinite, a natural uniform distribution will assign probability 1 to this proposition.

A residual issue is what kind of probability is involved in setting the standard for proper credence. This question is not yet settled by assuming that these are probabilities conditional on what one knows as virtually any kind of probability could in principle be conditionalized on what one knows. It is, I think, very plausible to assume that a certain kind of *epistemic probability* is involved in the standard for proper credence. A proper credence in  $p$ , formed by an agent  $A$  at a time  $t$ , would at least have to match the epistemic probability of  $p$  conditional on what the agent knows.<sup>25</sup>

Taking proper credences to go with epistemic probabilities conditional on what one knows moves the proposal fairly close to the one offered by Hawthorne and Stanley (2008). Recall that for Hawthorne and Stanley, one deals adequately with uncertainty in practical deliberation when one possesses knowledge about the corresponding evidential probabilities. The difference is that for Hawthorne and Stanley (2008), one has to form an epistemic attitude *about* epistemic probabilities (at least on one of the proposals they consider; see fn. 7), while on the present suggestion one merely has to form an epistemic attitude *in accordance with* the epistemic probabilities.<sup>26</sup>

On the picture developed in this paper, one can just as well partially rely on a proposition as one can fully rely on it in one's decision-making. Full reliance is governed by proper full belief, which was assumed to be knowledge. In a similar way, partial reliance is governed by proper credence. Steps toward unifying this picture can be taken in two directions. I proposed that one should think of proper credence in this context as being characterized by conditions which stand to partial belief as knowledge stands to full belief. Moreover, conflict between these norms can be avoided if one assumes that knowledge is the standard for practical certainty, while the credences relevant for practical deliberation are governed by epistemic probabilities conditional on what one knows.

## 5. Concluding remarks

Let us take stock. On the present proposal, knowledge governs full reliance, while proper credence governs partial reliance. An account along these lines enables one to take Schiffer's example at face value. Agents can and oftentimes should rely on uncertain information. When they do, they do not have to know the proposition in question. However, just as there is a normative constraint on full reliance, so there is a normative constraint on partial reliance.

This picture will not be attractive to everyone. But it should be attractive to friends of a knowledge norm for full belief who also favor a nonreductive account of credences. On the latter kind of account, credences are not to be reduced to full beliefs about probabilities and have ordinary propositions as their content. Hence, the question of when one should have a certain credence does not reduce to the question of when one should have a certain full belief. Moreover, credences would be expected to play a role in decision-making which does not reduce to the role of full beliefs. If one shares these views, a norm on partial reliance and credences which is similar to the knowledge norm on full reliance and full belief should be attractive.

The present account is tailored to cashing out knowledge norms for practical reasoning in terms of reliance. There are similar norms in which the notion of a *reason* takes center stage. For instance, Hawthorne and Stanley (2008) hold that one should treat a proposition  $p$  as a reason for action only if one knows  $p$ . Can the present account be adapted to reason-based norms as well? As far as I can see, a positive answer to this question is possible, although I must leave the details for another

<sup>25</sup>See Williamson (2000, chap.10) for a conception of evidential probability suitable to play the envisaged role of epistemic probability. Note that as the proposal is presently set up, one is not forced to identify one's evidence with one's knowledge, although there is a pull toward such a claim. All that is strictly required is that part of what makes a credence proper is that it matches one's epistemic probability conditional on what one knows.

<sup>26</sup>What if this is in a given situation not possible because the task would be too complex? Should one then form no credence at all? The answer, in my opinion, is that in such situations the demands on the subject should be relaxed, while different situations may require different reactions. Sometimes the right reaction might be to form an *imprecise* credence, while at other times it may be permissible to *ignore* certain bits of information one cannot evaluate.

occasion. Just as one can fully or partially rely on a proposition, it seems possible to adopt a qualified attitude toward potential reasons. One can fully or partially possess a reason depending on whether one possesses knowledge or merely has a (warranted) high credence. Thus, one may propose that one should treat a proposition  $p$  as a *full reason* for action only if one knows  $p$ . But one may treat a proposition  $p$  as a *partial reason* when one only has a (warranted) high credence.

Let me close with a few remarks on how the present picture sheds light on the potential tension between the knowledge norm for practical reasoning and classical decision theory in terms of expected value maximization. To better conceptualize this issue, let us distinguish a *formal* from a *material* understanding of decision theory. On a formal conception, decision theory would just make a very minimal claim: for some agent-dependent probability function  $p$  and some agent-dependent utility function  $u$ , the action one should choose is an action which has maximal expected value relative to  $p$  and  $u$ . As it stands, a formal conception merely says something about the structure of the choices one should make. A material conception, on the other hand, adds to a formal conception constraints beyond the mere formal structure, for example, constraints on the probability function and the utility function to be used in one's calculation of expected value.

As far as I can see, knowledge norms do not speak strongly on the formal aspect of decision theory. As far as the norms on full/partial reliance and full/partial belief are concerned, various decision rules could in principle be adopted (for more on this, see Weisberg 2013). The knowledge norms just tell us to what degree one should rely on a certain proposition in one's practical reasoning. They do not tell us how to decide in the light of certain full or partial information. Therefore, these norms are fully compatible with expected value maximization (but they are also compatible with alternative decision rules). The original knowledge norm for practical reasoning is only in tension with expected value maximization if the norm is supposed to govern not only practical reasoning with full beliefs but also practical reasoning with uncertain premises (this seems to be Schiffer's [2007] point). Once it is granted that practical reasoning can involve one's credences, it is perfectly coherent to assume that the right conclusion to draw from uncertain premises is to choose an action with maximal expected utility.

So, let us suppose for the moment that expected value maximization provides the right answer to the form a decision theory should take. Let us further suppose that the norms on full belief and partial belief proposed so far are on the right track. What follows for the material aspect of decision theory?

It is at this point, I think, where the present picture no longer stays neutral. In my view, the norms on full and partial belief are best seen as providing an answer to material questions of good decision-making. The norm on partial belief suggests that a certain credence should enter the calculation of expected value only if it is "knowledge-like." Thus, this norm results in a substantial normative constraint on the probabilities entering the calculation of expected value. The knowledge norm on full belief can be seen as a limiting case of this construal. As it describes what one can fully rely upon in one's practical reasoning, it allows to rule out those possibilities which are incompatible with one's knowledge. That is, compatibility with one's knowledge would be the entry condition for a possibility to be considered in one's decision-making. A view like this has been advocated by Weatherson (2012). According to Weatherson, knowledge answers the question of which possibilities should enter one's decision matrix in the first place. The probabilities in one's calculation of expected value would then range over those possibilities compatible with one's knowledge.

The picture just sketched is in deep disagreement with radically subjective Bayesianism. According to this view, there are hardly any material constraints going beyond a merely formal constraint on one's decision-making. As long as one's probabilities and utilities are coherent, one would not violate any norm of proper decision-making when choosing an action which maximizes expected value. The picture developed in this paper opposes the radically subjective Bayesian by demanding substantial normative constraints on which probabilities to use when deciding what to do.

**Acknowledgments.** I would like to thank two anonymous referees for their very helpful comments. This work was supported by the Deutsche Forschungsgemeinschaft under grant SCHU 3080/3-1/2/3 and was conducted within the Emmy-Noether research group “Knowledge and Decision.” Many thanks to all its participants (actual or associated) for their valuable feedback!

**Moritz Schulz** is professor of theoretical philosophy at the Technische Universität Dresden, Germany. His research includes the philosophy of language and logic, epistemology as well as decision theory. He currently conducts a research group entitled “Knowledge and Decision,” which studies the role of knowledge in (rational) decision-making.

## References

- Alonso, Facundo. 2014. “What Is Reliance?” *Canadian Journal of Philosophy* 44: 163–83.
- Alonso, Facundo. 2016. “Reasons for Reliance.” *Ethics* 126: 311–38.
- Beddor, Bob. 2020. “Certainty in Action.” *The Philosophical Quarterly* 70: 711–37.
- Beddor, Bob, and Simon Goldstein. 2021. “Mighty Knowledge.” *The Journal of Philosophy* 118 (5): 229–69.
- Bird, Alexander. 2007. “Justified Judging.” *Philosophy and Phenomenological Research* 74 (1): 81–110.
- Brown, Jessica. 2008. “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.” *Noûs* 42: 167–89.
- Carnap, Rudolph. (1950) 1962. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Dunn, Jeff. 2015. “Reliability for Degrees of Belief.” *Philosophical Studies* 172: 1929–52.
- Easwaran, Kenny. 2018. “Sarah Moss: Probabilistic Knowledge.” *Notre Dame Philosophical Reviews*. <https://ndpr.nd.edu/reviews/probabilistic-knowledge/>.
- Engel, Pascal. 2004. “Truth and the Aim of Belief.” In *Laws and Models in Science*, edited by Donald Gillies, 77–97. London: King’s College Publications.
- Fantl, Jeremy, and Matthew McGrath. 2002. “Evidence, Pragmatics, and Justification.” *The Philosophical Review* 111 (1): 67–94.
- Gao, Jie. 2017. “Rational Action without Knowledge (And Vice Versa).” *Synthese* 194: 1901–17.
- Gao, Jie. 2019. “Against the Iterated Knowledge Account of High-Stakes Cases.” *Episteme* 16: 92–107.
- Gerken, Mikkel. 2011. “Warrant and Action.” *Synthese* 178: 539–47.
- Greco, Daniel. 2013. “Probability and Prodigality.” In *Oxford Studies in Epistemology* vol. 4, edited by Tamar Szabó Gendler, 82–107. Oxford: Oxford University Press.
- Hawthorne, John, and Jason Stanley. 2008. “Knowledge and Action.” *The Journal of Philosophy* 105: 571–90.
- Holton, Richard. 1994. “Deciding to Trust, Coming to Believe.” *Australasian Journal of Philosophy* 72: 63–76.
- Hyman, John. 1999. “How Knowledge Works.” *Philosophical Quarterly* 49 (197): 433–51.
- Jackson, Alexander. 2012. “Two Ways of Putting Knowledge First.” *Australasian Journal of Philosophy* 90: 353–69.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- Konek, Jason. 2016. “Probabilistic Knowledge and Cognitive Ability.” *Philosophical Review* 125: 509–87.
- Locke, Dustin. 2015. “Practical Certainty.” *Philosophy and Phenomenological Research* 90 (1): 72–95.
- Lord, Errol. 2018. *The Importance of Being Rational*. Oxford: Oxford University Press.
- Moss, Sarah. 2018. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Müller, Andy, and Jakob Ross. 2017. “Knowledge Dethroned.” *Analytic Philosophy* 58: 283–96.
- Myers-Schulz, Blake, and Eric Schwitzgebel. 2013. “Knowing That P without Believing That P.” *Noûs* 47: 371–84.
- Pettigrew, Richard. 2021. “What Is Justified Credence?” *Episteme* 18: 16–30.
- Radford, Colin. 1966. “Knowledge—By Examples.” *Analysis* 27: 1–11.
- Rich, Patricia. 2020. “The Logic of Probabilistic Knowledge.” *Philosophical Studies* 177: 1703–25.
- Ross, Jacob, and Mark Schroeder. 2014. “Belief, Credence, and Pragmatic Encroachment.” *Philosophy and Phenomenological Research* 88 (2): 259–88.
- Schiffer, Stephen. 2007. “Interest-Relative Invariantism.” *Philosophy and Phenomenological Research* 75 (1): 188–95.
- Schulz, Moritz. 2017. “Decisions and Higher-Order Knowledge.” *Noûs* 51: 463–83.
- Schulz, Moritz. 2021a. “Practical Reasoning and Degrees of Outright Belief.” *Synthese* 199 (3–4): 8069–90.
- Schulz, Moritz. 2021b. “Strong Knowledge, Weak Belief.” *Synthese* 199 (3–4): 8741–53.
- Smithies, Declan. 2012. “The Normative Role of Knowledge.” *Noûs* 46: 265–88.
- Staffel, Julia. 2013. “Can There Be Reasoning with Degrees of Belief?” *Synthese* 190: 3535–51.
- Tang, Weng Hong. 2016. “Reliability Theories of Justified Credences.” *Mind* 125: 63–94.
- Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. New York: Oxford University Press.
- Weatherston, Brian. 2005. “Can We Do without Pragmatic Encroachment?” *Philosophical Perspectives* 19: 417–43.
- Weatherston, Brian. 2012. “Knowledge, Bets, and Interests.” In *Knowledge Ascriptions*, edited by Jessica Brown and Mikkel Gerken, 75–103. Oxford: Oxford University Press.
- Wedgwood, Ralph. 2012. “Outright Belief.” *Dialectica* 66: 309–29.
- Weisberg, Jonathan. 2013. “Knowledge in Action.” *The Philosopher’s Imprint* 13: 1–23.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

- Williamson, Timothy. 2005a. "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." *The Philosophical Quarterly* 55: 213–35.
- Williamson, Timothy. 2005b. "Knowledge, Context and the Agent's Point of View." In *Contextualism in Philosophy*, edited by G. Preyer and G. Peter, 91–114. Oxford: Oxford University Press.
- Williamson, Timothy. 2017. "Acting on Knowledge." In *Knowledge First*, edited by A. J. Carter, E. Gordon, and B. W. Jarvis, 163–81. Oxford: Oxford University Press.