

KNOWLEDGE IN ACTION

Jonathan Weisberg
University of Toronto

ABSTRACT. Recent proposals that frame norms of action in terms of knowledge have been challenged by Bayesian decision theorists. Bayesians object that these knowledge-based norms conflict with the highly successful and established view that rational action is rooted in degrees of belief. I argue that the knowledge-based and Bayesian pictures are not as incompatible as these objectors have made out. Attending to the mechanisms of practical reasoning exposes space for both knowledge and degrees of belief to play their respective roles.

ACCORDING to a spate of recent proposals, one's actions should be guided by what one knows. But according to Bayesian decision theory, rational decision-making is about maximizing expected utility given one's credences. Bayesians worry that the knowledge-based picture conflicts with their credence-based view because it deals in full beliefs, as opposed to the credences, or "partial beliefs", central to Bayesian decision theory. For example, Schiffer (2007) worries that the knowledge-based picture precludes middling credences from governing action, since they fall short of full belief, and hence fail to constitute knowledge. And Douven (2008) objects that the knowledge-based picture presumes a step-by-step method of practical reasoning at odds with the Bayesian method of calculating and comparing expected utilities.

But the knowledge- and credence-based pictures are not as incompatible as it seems. They are not even theories of the same thing. The Bayesian view says which of an agent's choices are rational, while the knowledge-based view says what assumptions she may rely on in making her choice. There is a real danger of collision, but whether that danger is realized depends on what auxiliary hypotheses we accept. It depends on how one's reasons for action interact in practical deliberation, and what role credences play in that process. The knowledge-based view will collide with the Bayesian view if reasoning only on the basis of what one knows prevents one from arriving at an expected utility maximizing choice, for example. There is also

Thanks to Jane Friedman, Alex Jackson, Sarah Moss, Jennifer Nagel, and Adam Sennet for their helpful discussion and feedback. I am especially grateful to Jennifer Nagel, whose generous guidance made this paper possible.

the danger that reasoning only on the basis of what one knows will leave no role for credence to play in practical deliberation. But whether these threats become realities depends on details neither party has said much about, details to do with the rules and mechanics of good practical reasoning.

To show how these dangers can be avoided, I will develop three irenic proposals that fill in some of these details, thereby bridging the gap between the two theories. My first proposal will answer the first Bayesian worry articulated by Douven: that the knowledge-based view presupposes a step-by-step method of practical reasoning at odds with expected utility maximization. In response, I outline knowledge-based methods of practical reasoning capable of making expected utility maximizing choices. My second and third proposals will answer the second Bayesian worry articulated by Schiffer: that the knowledge-based view precludes middling credences from governing action, since they fall short of full belief, and thus of knowledge. My second proposal answers that middling credences can constitute knowledge by constituting dispositional beliefs about epistemic probabilities, in which case acting on one's knowledge just amounts to acting on one's credence. And my third proposal argues that, even when credences don't constitute knowledge of epistemic probabilities, they can still influence action by serving as weights for the reasons one's knowledge does provide.

The picture to emerge will be one of rational agents striving to maximize expected utility by engaging in a variety of reasoning processes, processes that draw on a mixture of credal and knowledge states as appropriate. Which reasoning process we should use in a given situation, and thus what mixture of credence and knowledge we draw upon, will depend on a number of factors: how important it is that an optimal decision be made, what knowledge is available, how much time the subject has, and more. Detailing the mechanics of these processes and the norms governing them is a major project. The work done in this paper makes progress towards that end, but the central aim and contribution of the paper is not to provide a complete such theory. It is, rather, to bring the knowledge- and credence-based approaches closer together by drawing out the extent to which their differences are only apparent.

In the end, I will argue, both camps have left an important cluster of questions about practical reasoning largely unanswered. The three proposals developed in this paper make progress towards answering these questions, in a way that warrants optimism about the possibility of a friendly resolution between the two camps. Whether the two approaches are ultimately compatible, however, depends in the end on how those questions are answered in full.

I will begin by outlining the knowledge- and credence-based views in §1. I'll then use Douven's and Schiffer's Bayesian objections to the knowledge-based view as foils for my three proposals. Douven's worry will serve as a foil for the first proposal in §2, and Schiffer's as a foil for the second and third proposals in §3. Finally, in §4, I'll tie things together by explaining how the tensions between the credence- and knowledge-based views derive from misconceptions about what each view provides a theory of.

1. BACKGROUND

A number of authors now endorse knowledge-based norms governing practical reasoning.¹ We will use Hawthorne and Stanley's (2008) proposal as our representative:

The Reason-Knowledge Principle (RKP): Where one's choice is P -dependent,² it is appropriate to treat the proposition that P as a reason for acting iff you know that P .

RKP embodies a sufficiency claim and a necessity claim. The sufficiency claim is that it is always appropriate to use relevant knowledge as a reason in practical reasoning. The necessity claim is that it is only appropriate to use what one knows as a reason for acting. Our main focus will be the necessity claim, as it's been the primary target of Bayesian critiques.

Why endorse RKP? Hawthorne (2004) and Hawthorne & Stanley (2008) argue that the necessity half of RKP accounts for the way we assess practical reasoning. If you decline very cheap insurance on the grounds that misfortune will not strike, you are subject to criticism, since you don't know misfortune won't strike. Stanley (2005) also defends an RKP-like claim on the grounds that it accounts for intuitive shifts in what one knows as stakes change. Under ordinary circumstances, Hannah knows that the bank is open on Saturday since she was there just a few Saturdays ago. But if a lot hangs on whether the bank will be open on Saturday—Hannah must deposit her paycheck to avoid dire financial consequences—she does not know without further evidence. These intuitions line up nicely with the thought that what she knows varies with what it would be appropriate for her to assume in practical reasoning.

In Bayesian decision theory, choices are evaluated in terms of credences rather than knowledge:

¹See (Fantl & McGrath, 2002), (Hawthorne, 2004), (Stanley, 2005), (Fantl & McGrath, 2007), (Hawthorne & Stanley, 2008), and (Weatherson, 2012).

²A choice is " P -dependent" iff which option is preferable is different given P than given $\neg P$. The restriction to P -dependent choices avoids the absurd consequence that it is appropriate to use what one knows even when it is not relevant.

The Expected Utility Principle (EUP): It is rational to choose an act only if³ it maximizes expected utility⁴ with respect to one's credences and utilities.

Two motivations for EUP have been especially influential. The first is its agreement with common-sense: EUP gives intuitively plausible verdicts in a wide variety of cases.⁵ The second argument for EUP is more technical: representation theorems show that violating EUP means having irrational preferences. If one does not maximize expected utility, one will have intransitive preferences, or violate the sure-thing principle, or something similar. (von Neumann & Morgenstern, 1944; Savage, 1954; Jeffrey, 1965; Joyce, 1999)

Much more could be said for and against each of RKP and EUP, as well as their respective motivations. But our question here is whether they are compatible. Why think they are not?

There are two broad kinds of concerns. The first are what we might call "J vs. K" issues. These arise from the fact that RKP appeals to knowledge while EUP appeals to beliefs (albeit partial ones). Many Bayesians would agree that a belief must be justified or rational in order to guide action, but must it be more than that? Must it have the extra, external features that separate knowledge from justified belief? Brown (2008a; 2008b) puts a point on this concern by comparing two subjects, both of whom justifiably believe their train comes at 12:20 P.M. Only one of them actually knows this fact however, the other being the victim of a Gettier (1963) scenario. Intuitively, these two subjects are equally entitled to assume the train comes at 12:20 P.M. in deciding when to leave the office, yet RKP says only the knower is entitled.

The J vs. K difference is an important one, and points like Brown's pose a serious challenge for RKP. But these J vs. K issues will not be our focus here. For present purposes, I will assume that RKP's proponents have adequate answers to worries like Brown's. Maybe, as Hawthorne and Stanley (2008) suggest, Brown's Gettiered subject is blameless because she makes an excusable mistake, though the need for an excuse betrays the fact that she has violated a norm, namely RKP. Or maybe some other response is correct. I won't pursue the matter here.

³EUP is a necessary but not sufficient condition because it is possible to choose an expected utility maximizing option for the wrong reasons. And if one does, the action is not rational in an important sense. Rational action, like rational belief, must be well-founded (but cf. (Maher, 1993, 148–9)). This point will turn out to be important, especially in §2.

⁴The expected utility of act A is defined $\sum_i p(O_i|A)u(O_i \wedge A)$, where $\{O_i\}$ is a partition of possible outcomes of the act, p is the subject's credence function, and u her utility function.

⁵Alleged counterexamples have been discussed (Allais, 1953; Ellsberg, 1961; Nozick, 1969). Proponents of EUP respond that these are cases where common sense goes astray, or else the principle has been misinterpreted or just needs to be reformulated (Savage, 1954; Jeffrey, 1965; Maher, 1993; Joyce, 1999).

The second class of concerns are what we might call “partial vs. full” issues, arising from the fact that EUP appeals to partial beliefs while RKP appeals to knowledge and hence, presumably, full beliefs.⁶ For example, one worry is that RKP and EUP presuppose incompatible views about how practical reasoning works. A theory based on full beliefs invites a natural deduction sort of picture, where premises are used to infer lemmas which are in turn used to infer conclusions about how to act. Whereas a theory based on partial beliefs suggests a very different sort of procedure, calculating values of actions by applying the expected utility formula.

This paper is concerned with these partial vs. full issues. My aim is to develop a picture of practical reasoning that integrates the partial and full belief perspectives, one where credences and knowledge (hence full beliefs) both have significant roles to play. The picture I will develop sees practical reasoning as drawing on a variety of cognitive processes, some more classically Bayesian in their operation than others. These processes often draw on a mixture of cognitive states, combining knowledge with credences to settle on an appropriate action. In fact, in some cases I will argue, knowledge and credal states are one and the same (or at least very intimately metaphysically related), so that acting on a piece of knowledge amounts to acting on a credence (Moss, 2013). The resulting picture will not answer every concern about RKP and EUP’s compatibility, but it will bring the two closer together, and suggest ways for further work to close the gap.

2. METHODS OF PRACTICAL REASONING

RKP is motivated in large part by its ability to explain the impropriety of certain episodes of practical reasoning, like declining very cheap insurance on the grounds that misfortune won’t strike when you don’t know it won’t. One Bayesian worry is that such examples presuppose a step-by-step, premise→lemma→conclusion form of practical reasoning at odds with the Bayesian method of expected utility calculation. To clarify the exact nature of the challenge, it will help to focus on a concrete formulation from the literature.

2.1. **Clarifying the Challenge.** Douven (2008) gives voice to this worry, using an example from (Hawthorne, 2004) to illustrate. Suppose you have the opportunity to sell your lottery ticket for a penny and you reason as follows:

- (i) My lottery ticket is a loser.
- So, if I keep the ticket, I will get nothing.

⁶This presumption is challenged by Moss (2013). We will examine it more carefully in §3.

But if I sell it, I will get a cent.

So, I should sell the ticket.

RKP garners support from such examples by agreeing with common sense that the reasoning is faulty. The RKP diagnosis locates the fault in the fact that your reasoning is based on something you don't know, namely that your ticket is a loser. But, Douven counters, Bayesianism tells us the reasoning in (i) is bad for a different reason: "it deploys the wrong kind of reasoning for the purpose at hand: a decision about whether to sell a lottery ticket (or any other decision, for that matter) is not to be taken on the basis of a deductive argument like (i), but on the basis of expected utility calculations." (2008, 107) If the highly successful and widely endorsed Bayesian theory of decision is right, then isn't RKP on the wrong diagnostic track altogether?

There is an obvious reply to this initial Bayesian challenge, but a more serious challenge is waiting in the wings. The obvious reply is that Bayesianism does not require us to calculate expected utilities, neither consciously nor even unconsciously. Bayesian decision theory says that one should choose an option that maximizes expected utility, but it says nothing about how one must arrive at that choice. EUP, for example, does not say anything about what steps one must follow in coming to a decision, so long as one ends up at a decision that maximizes expected utility. Indeed, many Bayesians explicitly acknowledge that one needn't make one's choice on the basis of expected utility calculations, since they allow that dominance reasoning is rational when applicable. I might buy fire insurance because the resulting peace of mind will far outweigh the minor financial cost, leading to a better outcome whether there is a fire or not. In this case I do not calculate expected utilities, but instead use a form of reasoning that has the same effect.⁷

Reasoning that doesn't involve actually calculating expected utilities isn't just compatible with the letter of Bayesianism, it is also compatible with its spirit. The main arguments supporting EUP are silent on the question how one should arrive at an expected utility maximizing choice. For example, the argument that EUP agrees with common sense frequently relies on dominance reasoning and other heuristics to elicit intuitions about cases. And the representation theorem argument for EUP doesn't speak to one's method of reasoning either, but instead to the coherence between one's beliefs and preferences. According to that argument, failing to choose

⁷In fact, (i) sets out precisely this form of argument, trying to establish that selling your ticket dominates. The natural Bayesian diagnosis of the flaw in this dominance reasoning is that it uses the wrong space of possibilities, or the wrong distribution of probabilities over that space. But this Bayesian diagnosis looks to be compatible with RKP. You can't exclude the possibility that your ticket will win from your decision table because you don't know it won't obtain. See Weatherston (2012) for more on the proposal that RKP complements Bayesian reasoning by shaping decision tables.

an expected utility maximizing option means having intransitive preferences, or preferences that violate the sure-thing principle, or something similarly unpalatable. But any method of forming one's preferences that maximizes expected utility will avoid these unpalatable consequences equally well.

The broad point here is that we must distinguish between what psychologists call theories of *substantive* rationality and theories of *procedural* rationality (Simon, 1982). EUP is not a theory of *procedural* rationality, a theory that says which ways of reasoning towards optimal actions are good or bad. It is a theory of *substantive* rationality, a theory that says what choice is rational given one's cognitive state. Given one's state of mind—beliefs, knowledge, credences, desires, utilities, and so on—which action is optimal? Bayesian decision theory says it is any action that maximizes expected utility. But it does not say anything about what reasoning or argumentation one may use to arrive at that choice (except that it must be reasoning that takes one to a substantively rational choice).

So why worry that EUP might be incompatible with examples like Hawthorne's (i)? Here we come to the more serious Bayesian challenge. What procedural theory could RKP's proponents have in mind such that:

- (a) practical reasoning is representable in the step-by-step, premise→ lemma→ conclusion format of (i), and
- (b) it results in expected utility maximizing choices?

One might well be skeptical that any method of practical reasoning could satisfy both (a) and (b).⁸

Here again Douven gives voice to the worry, motivating skepticism about the joint satisfiability of (a) and (b) by focusing on another of Hawthorne's examples, this time an example of apparently good reasoning. Suppose you are a person of modest means shopping in a bookstore, and you are contemplating whether to buy the local destination guide or the more expensive worldwide guide. You reason as follows:

- (ii) I won't be able to afford a trip to an exotic destination.
Thus, I won't have any use for the worldwide guide.
Thus, I should buy the local destination guide.

⁸I take this to be the challenge Douven ultimately means to raise for RKP, since he acknowledges the rationality of dominance reasoning (2008, fn. 21), and goes on to consider the possibility that RKP's proponents might prefer a procedural theory that does not involve calculating expected utilities.

RKP allows that (ii) is good reasoning, since you presumably know what you can afford, and this premise supports the subsequent lemma and conclusion. But, Douven contends, (ii) is only good if it is supplemented with a number of additional assumptions. For example, we must take you to know, or at least presuppose, that the worldwide guide is not of outstanding quality and the last copy to be found anywhere. We must also assume that you do not expect to be in a much better financial position in a couple years. Further still, we must assume that your niece is not planning a trip to an exotic destination, so that she would have use for the guide. And so on. Douven's point seems to be that for (ii) to represent a good episode of reasoning, it must be an enthymematic representation of a much more involved argument:

(ii') I won't be able to afford a trip to an exotic destination, *ever*.
 Thus, I won't have any use for the worldwide guide, *ever*.
 No one else I know will have any use for the worldwide guide, *ever*.
 I couldn't resell the worldwide guide for a profit.
 ⋮
 Thus, buying the local guide is the best option.
 Thus, I should buy the local destination guide.

But if (ii') correctly represents the reasoning in question, then RKP should deem it bad on the grounds that some of the premises are not known. What seems true is not that the myriad possibilities highlighted by (ii') are ruled out by one's knowledge, but that they will carry little weight in an expected utility calculation, since you give them little credence. In short, we are being pressed to grant that (ii) is really just a crude summary of a much more involved reasoning process, one that is more accurately represented by a thorough Bayesian decision table. The challenge is to explain how one could reason one's way to buying the local guide from the scant knowledge given in the premise of (ii), when the myriad possibilities acknowledged by (ii') cannot be ruled out by one's knowledge.

My central proposal in this section is an attempt to answer that challenge. But first let me resist the pressure to replace (ii) with (ii'). Douven presses for (ii') as the more faithful representation of your reasoning by pointing out possibilities that would make the worldwide guide the more optimal choice, insisting that you must in some sense rule them out. But consider the analogous attack on a commonplace piece of non-practical reasoning:

I remember locking my door when I left the house this morning.
 So my home and belongings are as I left them.

Maybe there is some sense in which I need to rule out possibilities that would undermine the support the premise lends to the conclusion here. For my reasoning to be good, I may need to know that my memory is working well, that forceful break-ins are rare in my neighbourhood, that my landlord is unlikely to use his key to steal my things, and so on. But these factors needn't enter my reasoning, consciously *or unconsciously*, for my reasoning to be good. My reasoning can be as simple as noting the premise and combining it with my (perhaps tacit) knowledge that the premise makes the conclusion probable. It may be that I need to know many other things in order to know (perhaps tacitly) that this premise renders this conclusion probable. But possessing that knowledge needn't complicate my actual reasoning process.

The pressure to acknowledge (ii') as a more faithful representation seems to rest on a mistaken presupposition that (ii) is intended to be deductive.⁹ The suggestion seems to be that all possibilities of error must be ruled out before the conclusion is warranted. But this supposition is one that Douven imposes on proponents of RKP, not one they have any reason to accept. It is uncontroversial that knowledge-based reasoning in theoretical domains is frequently non-deductive, and I see no reason the practical domain should be different.

Returning to our central question then: how might one's reasoning proceed in (ii)? What process could take you from such sparse information to the conclusion that the local guide is the better option, if not by taking account of the myriad possibilities acknowledged by (ii') and then calculating expected utilities?

2.2. First Proposal: Knowledge-Based Reasoning. Research in the last 40 years strongly supports the view that we do not make decisions (only) by calculating expected utilities. We also use a variety of more economical methods, methods specifically designed to make effective use of sparse information. There is a massive research program in psychology dedicated to determining what methods we use, when we use them, and how effective they are at generating expected utility maximizing choices. This program is far from complete, but it has come far enough to offer a plausible answer to Douven's challenge. It has uncovered fairly effective ways of determining which option maximizes expected utility, not by actually calculating expected utilities, but instead using reasoning that looks much like that represented in (ii).¹⁰

⁹Indeed, Douven frequently accuses RKP's proponents of presupposing a "deductive" conception of practical reasoning. But he does not say why it must be deductive as opposed to merely being representable in premise→lemma→ conclusion format.

¹⁰For some early work in this program, see (Simon, 1956; Tversky, 1972; Kahneman & Tversky, 1973). For more contemporary surveys, see (Payne et al., 1993; Gigerenzer et al., 1999; Baron, 2007). For an

One extensively studied such method is the lexicographic heuristic, LEX.¹¹ When deciding between two options, LEX considers various desirable attributes each option might have or lack, searching through them in descending order of importance until a “tie-breaker” is found. If we are deciding between restaurants *A* and *B*, and we rate in descending order of importance: price, taste, service, proximity, and ambience, LEX will run through these factors in that order until it hits on an attribute with respect to which one option does noticeably better than the other. If *A* and *B* are comparable in price, it will go on to consider taste; if they are comparable in taste, it will go on to consider service; if *B* has noticeably better service, LEX will stop and settle on option *B*. (Payne et al., 1993, ch. 2)

LEX is crude, going with the first tie-breaker it finds. Searching for a tie-breaker by order of importance is better than searching randomly, but there is still plenty of room for error. The option not chosen could actually be the one that maximizes expected utility, possessing many desirable attributes that weren’t considered because they were too far down in the search queue. One way of reducing such errors is to keep the search going until it becomes unlikely that further searching will make a difference. For example, we might keep track of how many desirable features each option has, and to what degree it has them, searching until one option passes some pre-determined threshold. When greater assurance of arriving at the expected utility maximizing option is needed, we can set the threshold high; when optimality is not so essential, it can be set lower. These enhancements of LEX yield another prominent proposal, the evidence accumulation model, or EAM. (Lee & Cummins, 2004)

Many methods besides LEX and EAM have been proposed and continue to be studied,¹² but these two will serve as examples. They show that there are simple methods for determining which option maximizes expected utility without actually calculating expected utilities. Moreover, they look to be the sorts of methods we plausibly use when reasoning as represented in (ii). Consider two competing stories about how a rational subject actually reasons when we describe her as using the reasoning in (ii). The first story, Douven’s, says that (ii) is a badly enthymematic representation of an expected utility calculation that draws on the expanded space of possibilities acknowledged in (ii’). The second story, mine, says that the subject compares her two options on a couple of the most important attributes, favouring

opposing paradigm, see (Oaksford & Chater, 1998, 2007, 2009). For a critical review of recent work in this tradition, see (Hilbig, 2010).

¹¹LEX is closely related to, but distinct from, the “take the best” heuristic popularized by Gigerenzer and his colleagues (Gigerenzer & Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer & Selten, 2001).

¹²Payne, Bettman, and Johnson (1993, ch. 2) survey other prominent proposals. For recent work on LEX, see (Gigerenzer et al., 1999; Bröder, 2000; Newell et al., 2003; Newell, 2005; Bröder & Newell, 2008), and for recent work on EAM see (Lee & Cummins, 2004; Newell, 2005; Newell & Lee, 2011).

the one that does better on both. She considers of each guide how much it costs and whether she is likely to have use of it in the near future, opting for the local guide since it does better in both respects. If more were at stake or if she were pressed to justify her choice, she might consider further attributes or even calculate expected utilities. But in the circumstances, a restricted set of considerations suffices to settle the matter. I submit that the second story is more plausible. At the very least, it would be hasty to dismiss it out of hand, and the support claimed for RKP along with it.

What if the agent had different background credences about the considerations Douven raises in (ii')? If her credences about her niece's travel plans or about the potential for profitable resale were different such that EUP recommended the world-wide guide instead, would LEX and EAM change their recommendations too? If not, Douven's worry would seem to remain unanswered: actually calculating expected utilities would remain the only known method for arriving at the decisions recommended by EUP.

In fact, methods like LEX and EAM are sensitive to such differences, in two ways.

First, such differences will affect the order of the search queue, the order in which desirable attributes are considered. LEX orders the attributes in its search queue by their importance, and having different credences about relevant matters will affect your assessment of what's important. If you believe your niece's birthday is coming up, and you suspect she might appreciate a travel guide, then which guide she would prefer becomes important. Similarly, the more strongly you suspect you could resell your purchase for a nice profit, the more importance you'll give to resale value. In fact, one way of formalizing the idea that attributes are ordered by their importance is to rank them by *validity*, where an attribute's validity is the probability that option *A* is all-things-considered preferable to *B*, given that *A* does better than *B* with respect to that attribute. Lee and Cummins adopt this formalization in their (2004) presentation of EAM. And on this approach it is particularly clear that credal differences affect the order of the search queue.

Second, background credences will affect which method an agent uses to make her decision. There is a large literature on *strategy selection*, i.e. on what factors determine what decision method we use, and how they determine that method. Many theoretical frameworks are on offer, with no universal agreement as to which is correct (see (Payne et al., 1993, 99–114) and (Bröder & Newell, 2008) for helpful surveys). But a common theme across these frameworks is that subjects select strategies roughly in accordance with their respective costs and benefits. Different frameworks emphasize different costs, like time, cognitive effort, error-proneness, and others. Benefits are understood in terms of anticipated accuracy, which could

mean anticipated conformity to a norm like EUP, or anticipated achievement of an actually good outcome.

But regardless of these details, the cost-benefit view predicts that changes in background credences of the sort we are considering will incline the agent towards reasoning that is more likely to agree with EUP. The more potential uses you foresee for your purchase at the bookstore, the more cause you have to anticipate that a more thorough examination of each option's pros and cons will lead to an accurate decision. The more strongly you suspect that your niece might enjoy a travel guide as a gift, or that you might be able to resell your purchase for a nice profit, the more probable it is that a more thorough assessment of your options and their potential uses will lead to the more optimal choice. This might mean choosing a more comprehensive method over a cruder one like LEX, or it might mean setting a higher threshold in your use of EAM. Whatever the particulars, the general effect is that increasing your credences in the kinds of possibilities raised by Douven's (ii') increases the chance that these factors will figure into your choice, making it more likely that you will choose the worldwide guide in a case where doing so maximizes expected utility.

2.3. Imperfection & Rationality. Nevertheless, methods like LEX and EAM will not respect EUP perfectly: it is possible to apply LEX or EAM correctly and still settle on an option that does not maximize expected utility. As noted earlier, LEX can select the sub-optimal option because the alternative's numerous desirable attributes were too far down in the search queue to be considered. And the same is true of EAM, especially when the pre-determined threshold is set low (because time is short, for example). The current proposal thus implies that an episode of practical reasoning need not follow a perfect method for maximizing expected utility in order to be good reasoning.

Some Bayesians will feel uncomfortable allowing the use of such imperfect methods. But many Bayesians have long held that the Bayesian canons of rationality are only idealizations, useful because they illuminate philosophical problems and serve as instructive idealizations (Horwich, 1982, 1993; Garber, 1983; Christensen, 1992, 2004, 2007; Weirich, 2004; Hawthorne, 2005; Hájek & Hartmann, 2010). And a number of authors already advocate using heuristics to approximate Bayesian ideals in the domain of theoretical reasoning. Okasha (2000), McGrew (2003), and Lipton (2004) suggest that Inference to the Best Explanation is a heuristic for approximating Bayesian updating. There is also evidence that human decision makers actually perform better when they use heuristics than when they try to calculate expected utilities (Hogarth & Karelaia, 2007; Gigerenzer & Brighton, 2009).

Even so, allowing that imperfect methods of reasoning can be rational does raise hard questions. What should we say about choices arrived at by rational means when they violate EUP: are these choices irrational yet still ‘reasonable’ in some important sense? Or should we say that they would be irrational for ideal agents, but they are rational for us, limited agents?¹³ Call these cases—where good reasoning departs from EUP—*tricky*.

Luckily, we needn’t take a stand on these thorny questions about tricky cases in order to answer the present Bayesian challenge to RKP.

First because RKP’s proponents can garner support from non-tricky cases, those where imperfect methods do respect EUP. Hawthorne’s (ii), for example, is presumably intended as such a case. And on this understanding of the case, RKP agrees with intuition; RKP finds no flaw in your reasoning, as seems right. But it would find a flaw if, say, you knew that a wealthy relative whose will features you prominently were about to undergo high-risk surgery. For then you would be reasoning based on something you do not know, namely that you won’t be able to afford a trip to an exotic destination. And that seems right too: your reasoning would then be subject to criticism for assuming that you won’t be able to afford an exotic trip.

Second, RKP garners support from tricky cases too, whether we say the choices in those cases are imperfectly rational, or that they are rational for us despite being irrational for ideally situated agents. To illustrate, consider a tricky variant of (ii). Suppose that, pressed for time, you reasonably engage EAM with a low threshold in order to decide between the two travel guides. And suppose you choose the local guide as a result, though circumstances are such that the worldwide guide would have emerged as preferable had you considered every last one of the myriad potential uses each guide might be put to. There is a clear sense in which, even if the local guide is not the rational choice, your reasoning would suffer an additional defect if you weren’t entitled to one of your assumptions. If, again, a wealthy relative whose will features you prominently were about to undergo high-risk surgery, there would be something defective in your reasoning, a defect that would not be present if you did know this premise. And RKP explains this difference.

In general, whether we deem the choices in tricky cases rationally perfect or imperfect, RKP still sorts better tricky cases from worse ones. Even if your choice in a tricky case suffers some rational defect just in virtue of contravening EUP, your reasoning can still suffer further defects. And RKP gains support insofar as it explains these further defects. If, on the other hand, choices in tricky cases suffer no rational

¹³I am grateful to an anonymous referee for pressing me to address this issue.

defect just in virtue of contravening EUP, RKP gains support insofar as it explains other imperfections when they arise.

To sum up then, Douven may be right that there is no way to always and perfectly meet the demands of EUP, except to calculate expected utilities. But most psychologists, and many philosophers, have come to suspect that doing so is far beyond our cognitive abilities. They interpret EUP as an idealization rather than as a necessary condition on good practical reasoning. And on this view, Douven's challenge can be answered. There are rational alternatives to the Bayesian, expected utility calculating procedure, and these alternatives can underwrite the conception of practical reasoning presupposed in examples like (i) and (ii).

3. REASONS & CREDENCES

A different Bayesian objection to RKP is that a knowledge-based picture of practical reasoning leaves middling credences no role to play in the governance of action. The worry begins to emerge when RKP's proponents are confronted with cases where it seems rational to act on such a credence. For example, Schiffer (2007) worries about the case of Jane, who has 0.4 credence that it will rain and consequently carries an umbrella. Her reason for carrying an umbrella is not that it will rain, for she does not believe it will. Nor does she believe it won't. She acts instead on her middling credence, and yet she seems rational.

Hawthorne & Stanley (2008) reply that Jane's reason for carrying an umbrella is that there's a decent chance of rain, which is something she does know. Here they understand 'chance' to mean *epistemic probability*, where one's epistemic probability for P is the extent to which one's total body of knowledge provides evidence for P . Epistemic probability should not be confused with credence, which is a psychological matter of one's confidence in P . Epistemic probabilities are also distinct from what philosophers of science frequently call 'objective' or 'physical' chances, which are instead the probabilities that appear in physical theories like quantum mechanics.

Schiffer worries that Jane may not possess the sophistication to distinguish epistemic probability from these other kinds of probability, and hence may not be able to know that there's a decent chance of rain in the relevant sense of 'chance'. But Hawthorne & Stanley caution against over-intellectualizing beliefs about epistemic probabilities. Folk discourse makes frequent use of expressions for epistemic probability, as in "it's likely to rain" and "there's a good chance it'll rain." So it seems that having beliefs about epistemic probabilities does not require philosophical sophistication. Thus Jane can believe, and know, that there is a decent chance of rain. And it is this knowledge that is her reason for carrying an umbrella, consistent with RKP.

This is where the real tension between the credence- and knowledge-based pictures emerges. To accommodate cases like Jane's, Hawthorne & Stanley re-describe them as cases of acting on a belief about epistemic probability, rather than acting on the basis of a credence. Beliefs about epistemic probabilities thus threaten to supplant credences in decision-making. What then becomes of credences? Must we abandon them in favour of beliefs about epistemic probabilities, jettisoning the psychological picture traditionally at the core of Bayesianism? Or can traditional Bayesian psychology be integrated with a picture where action is based on known reasons?¹⁴

I opt for integration, and will offer two proposals to that end. The first is that credences typically *constitute* beliefs about epistemic probability, and hence sometimes constitute knowledge about epistemic probability. Thus acting on one's knowledge of epistemic probabilities sometimes just *is* acting on one's credences. The second proposal is that, in many cases where one's reason is a known proposition, one still needs to determine how much weight to give that reason, and credences serve to determine that weight.

These two proposals are complementary, as we'll see. The first proposal illuminates the role that credence plays when we reason on the basis of probabilistic knowledge, knowledge of how epistemically probable a certain outcome is. The second proposal illuminates the role credence plays when we act instead on the basis of non-probabilistic knowledge, knowledge of how things are, or might be. We saw in §2 that there are multiple methods of practical reasoning, with different methods requiring different cognitive resources. This includes requiring knowledge of different kinds and contents. The two proposals together show that credence plays a role whether the knowledge deployed is probabilistic or non-probabilistic (though it plays different roles in each case). After presenting each proposal separately, in §3.1 and §3.2 respectively, I'll return to this complementary relationship in §3.3 for a more unified perspective.

3.1. Second Proposal: Credences & Beliefs About Probabilities. According to RKP, one's belief that a storm is coming can warrant staying home if that belief constitutes knowledge. Similarly, Jane's belief that there is a 0.4 chance of rain warrants carrying an umbrella if that belief constitutes knowledge. Now suppose we say Jane's 0.4 credence that it will rain constitutes her belief that there is 0.4 chance of

¹⁴There is also the worry that Bayesians of the subjectivist school will reject the very notion of epistemic probability. They will say that there are subjective probabilities in the form of degrees of belief, and maybe objective chances as discovered by physical theories like quantum mechanics, but nothing in between; there is no such thing as the extent to which one's knowledge provides evidence for *P*, there is only one's credence in *P* given one's total knowledge.

rain. Then her 0.4 credence warrants carrying an umbrella, consistent with RKP. In general, if credences can constitute beliefs about probabilities,¹⁵ and those beliefs can constitute knowledge, then RKP permits acting on some credences.¹⁶ For those credences constitute knowledge.

The view that credences can constitute knowledge comes from Moss (2013). I am defending the same claim on different grounds. The argument here turns on a metaphysical thesis about the relationship between credences and beliefs about probabilities. The claim is that, while a 0.4 credence in P may not constitute a belief that P , it can constitute a belief that there's a decent chance that P , or even a belief that P is 0.4 probable. My argument for this claim rests on two complementary supports.

First, we are generally prepared to attribute beliefs about what may well happen, what is likely to happen, or how probable a certain happening is, to subjects who have the corresponding credences. When someone gives something low credence we say that they believe it is improbable; when they give it middling credence we say that they believe it is somewhat probable; and when they give it high credence we say they think it highly probable. Similarly, when someone gives higher credence to P than to Q , we say they believe P is more probable than Q . (Yalcin, 2007, 2012) One might object that this tendency just reflects a strong psychological correlation between having a certain credence in P and also having formed an outright belief about P 's probability. Maybe, when we have a credence in P , we tend to also form an outright belief about P 's probability. This would explain the pattern just described without going so far as to say that the credences in question *constitute* the beliefs reported. But the best candidates for cases where a subject has a credence without a corresponding belief about probability do not support this conjecture. Take someone who has never considered whether P is probable, but still has some level of confidence about P , and acts on that confidence-level. It will typically¹⁷ be acceptable to say that they thought or believed P was so-and-so probable.

Second, there is a very plausible explanation for the above connection, namely that credences often constitute dispositions to form occurrent beliefs about probabilities. While one may have a credence in P without having formed any overt belief about how probable P is, one's credence can still constitute a *non-occurrent*, or *dispositional*, belief about probabilities. After all, when someone has credence x in P , they are very

¹⁵What kind of probabilities? Like Hawthorne and Stanley, I take these to be epistemic probabilities. But as we'll see, one can take an expressivist view on which these belief attributions, and the corresponding knowledge attributions, are cashed out in terms of credences.

¹⁶Viz. *kredences*.

¹⁷Exceptions may arise when the agent is seriously incoherent. If she is disposed to judge that P is highly probable even though she gives it low credence, it may not be appropriate to say that she believes P is improbable.

close, cognitively speaking, to forming an outright judgment that P is x probable. Suppose, for example, you are fairly confident your horse will lose the race. If asked whether your horse will probably lose, you can formulate your answer by consulting your credence that your horse will lose, and then translating it into an overt judgment that he will probably lose. In general, when one has credence x in P , one is disposed to judge that P is x probable.¹⁸ This close cognitive connection between credence and occurrent probabilistic belief makes it plausible that credences constitute dispositional beliefs about probabilities.

One might nevertheless be skeptical that these dispositional beliefs about probabilities can constitute knowledge. Knowledge appears to be a relation to a true proposition, but what proposition could a subject know/believe when we ascribe these sorts of beliefs about probabilities? Like Hawthorne & Stanley, I think the relevant propositions are about epistemic probabilities. Someone cannot know that P is probable unless their epistemic probability for P is high, so it is natural to say that their (dispositional) belief is a belief about their epistemic probability.

What about Bayesians of a severely subjectivist bent, for whom probability is strictly personal? Even subjectivists who reject the notion of epistemic probability can partake in the goods of the present proposal. Building on work by Yalcin (2007; 2012), Moss (2013, §4) shows how the factivity of knowledge, and other seemingly truth-directed requirements like safety and sensitivity, can be understood in an expressivist vein. This expressivist approach does not require that knowledge be a relation to truth-apt propositions, such as propositions about epistemic probability. On this view, factivity, safety, and sensitivity are instead captured by understanding knowledge as a relation to a constraint on credences.

For our purposes, either approach will serve. Friends of epistemic probability can understand the dispositional beliefs I'm defending as beliefs about epistemic probability; foes can take the expressivist path forged by Yalcin and Moss. Either choice allows a credal state to constitute a doxastic state with all the properties one might demand of knowledge: factivity, safety, sensitivity, etc. As a friend of epistemic probabilities though, I will continue to talk in those terms.

¹⁸This disposition is likely imperfect. Indeed, empirical work suggests that people's dispositions to convert credences into overt probability judgments are subject to well-known biases like overestimating low values and underestimating high ones (Hurley & Shogren, 2005). Fortunately, the current proposal can accommodate imperfection. People who are very imperfect, e.g. those disposed to judge that P is highly probable when they give it low credence, may be too cognitively incoherent to count as knowing that P is improbable, perhaps because they can't even be said to believe that P is improbable (fn. 17). But people disposed to judge that P is $x + \epsilon$ probable when they give it credence x may still believe/know that P is *roughly* x probable. And such knowledge will often be sufficient to guide action.

The present proposal has the nice feature of integrating Hawthorne & Stanley's view with the traditional Bayesian view. By saying that credences constitute beliefs about epistemic probabilities, we can allow credences to guide action while requiring that only knowledge guide action. Hawthorne & Stanley's view—that knowledge of epistemic probabilities guides action in cases like Jane's—was compatible with this possibility all along. But by spelling out the metaphysics of the connection between credences and beliefs about epistemic probabilities, we see that the irenic possibility is actually quite plausible. We thereby allay the worry that acting on knowledge about the probability of P precludes acting on one's credence in P . Because one's belief about P 's epistemic probability can be constituted by one's credence in P , one can act on one's credence in virtue of acting on one's knowledge.

But what if your credence constitutes a probabilistic belief that fails to qualify as knowledge? Suppose your credences about a coin flip are 50/50, but they are unsafe—maybe you are too easily inclined to use the principle of indifference, so that you would have the same credences even if the true epistemic probabilities were different. In such a case, EUP will give your credences the usual weight in determining what choices are rational, but RKP will forbid you from reasoning on the basis of your beliefs that heads and tails are each 0.5 probable. If you are choosing between staking \$1 on heads and \$1 on tails, EUP will say these options are equipreferable. But RKP would seem to block you from reaching this conclusion, since it forbids you from reasoning on the basis of your beliefs about the probabilities of heads and tails.

To some extent this is the kind of “J vs. K” issue we bracketed back in §1. The problem arises because some of your doxastic states, though justified, fail an external requirement on knowledge, specifically safety. But we can't just bracket the problem on the assumption that it will be fully resolved by whatever answer RKPers give to J vs. K challenges like Brown's. Recall, Brown's challenge was that the subject of a Gettier scenario seems warranted in relying on her justified true belief in her practical reasoning. If RKPers answer that Brown's subject makes an excusable mistake, but a mistake nonetheless, this may not completely resolve our problem here. Our subject, with her unsafe credences in heads and tails, might be excused for relying on her unsafe belief that heads and tails are equiprobable. But what could she do to avoid needing any excuse? Must RKPers say that she has no alternative but to make an excusable mistake?

The lessons of §2 provide the answer here. RKP blocks you from reasoning in a particular way in the present case: it forbids you from reasoning on the basis of the assumption that heads is 0.5 probable, since that's not something you know. But as we saw in §2, there are multiple methods of practical reasoning. In the next

section we'll see that you needn't reason on the basis of the assumption that heads is 0.5 probable for your 0.5 credence to influence your decision; your credence can influence your decision in another way, by acting as a weight for other reasons, things you do know. But even before we come to that proposal, the resources of §2 point to an answer to the present concern. For even if your credence in heads is forbidden from influencing your decision altogether, other methods of practical reasoning can still provide a way forward. For example, reasoning in the style of LEX or EAM can still issue in a decision. Your knowledge about the coin is symmetric, so searching for a tie-breaker between these two options will be a bust. Each option has the same desirable features as the other, at least as far as you know, so the search will reach the end of the queue without finding a tie-breaker. Neither option will thus emerge as preferable, consonant with EUP's verdict.

In this example, reasoning with LEX or EAM yields the same recommendation as EUP. There will be cases where working with the limited pool of information allowed by RKP will result in a decision that disagrees with EUP. But this is a familiar situation from §2. EUP represents the ideal case where cognitive resources are unlimited, whereas we are considering a case where the agent is obliged to work with limited resources because she has limited knowledge. Were all her credences knowledge-constituting, she would not be obliged to fall back on imperfect methods. Moreover, the proposal of the next subsection shows that, even when your knowledge is limited in this way, you needn't necessarily fall back on imperfect methods; there is other knowledge you can deploy which, when weighed according to your credences, will agree with EUP.

3.2. Third Proposal: Credences As Reason-Weights. A natural thing for Jane to say is that she brings an umbrella because it *may* rain. But the fact that it may rain does not always carry enough weight to warrant carrying an umbrella. Two subjects can each reasonably treat the fact that it may rain as a reason for carrying an umbrella, while only one of them should, on a balance of reasons, opt to carry one. 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. 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. Thus proponents of RKP and EUP can both

be right: one's reasons must be known, but one's decisions should still weigh one's credences in various possibilities.

This way of reconciling our two parties in Jane's case threatens to raise alarm bells in both camps. The proposal relies on specific claims about Jane's reasons—about what her reasons are and how they balance out. Without a general account of practical reasons and how they interact to back up these particular claims, the proposal is likely to meet with resistance.

Bayesians, being accustomed to thinking of decision-making as strictly a matter of calculating expected utilities, may be leery of any talk of “weighing reasons” for and against a course of action.¹⁹ There is only one real reason to choose an act, they may say, namely that it maximizes expected utility. And that reason is decisive. Other Bayesians may take the opposite tack: they will say this proposal gives the game away to the Bayesian, exposing the fact that Bayesianism is the one true account of reasons and their weighing. The proposal works because good practical reasoning is just a matter of weighing possibilities according to their probabilities and utilities, just as EUP says. There is, they will say, nought to reasons and their weighing but the calculation of expected utility. The only sense in which knowledge governs practical reasoning is a trivial one, namely that one should weigh those possibilities one knows about.

RKPers, on the other hand, may worry that the proposed reconciliation respects the letter of their view but not the spirit. On the present proposal, Jane's 0.4 credence in rain does not supply a reason, but instead a reason-weight. So the present proposal still allows doxastic states besides knowledge to influence decision-making. Isn't that against the spirit of the idea behind RKP, that only knowledge should be relied on in decision-making?

I'll address each of these concerns in turn. I'll first argue that Bayesians should be open to the idea that we use and weigh reasons. I'll then show that Bayesian decision theory suggests a natural account of how *some* reasons should be weighed, an account that vindicates precisely the claims about Jane the present proposal needs. This Bayesian account will prove merely partial, however, contra those Bayesians who might see nothing to the weighing of reasons but what Bayesianism tells us. Indeed, considerations introduced in §2 will demonstrate how remaining parts of the account can be filled in. Finally, as to the worry that only the letter of RKP is being respected, an examination of the motivations for RKP will show that the spirit

¹⁹As one prominent Bayesian put it to me in personal communication, “Reasons? We don't need no stinking reasons.”

of RKP is well respected by the present proposal.

3.2.1. *Reasons & Bayesians.* Some Bayesians will be wary of the present proposal, viewing any theorizing about reasons and how they balance as misguided. We already have a successful decision theory that circumvents any need for talk of reasons and their interactions, these Bayesians will say. So we should view talk of reasons and their weighing askance—as a mere relic of a defunct folk discourse, or at best as a loose way of talking about what really rationalizes decisions, namely expected utilities.²⁰

This view of things rejects the very starting point of our discussion, since it eschews any attempt to theorize about precisely that which RKP is a theory of: what an agent may use as a reason. As such, it veers close to Douven's objection to RKP from §2, and merits a similar response.

Bayesianism might tell us which decisions are rational (perhaps in some idealized sense of 'rational'), but it does not address procedural questions of practical rationality, especially for real, limited agents. Consequently, it does not tell us whether we real agents ought to decide in a way that relies on reasons and their weighing. Moreover, this dismissive attitude goes against both common sense and empirical psychology. Examples commonly cited in the RKP literature, like Hawthorne's (ii) from §2, make it plain that we commonly conceive of decisions as based on reasons, which often need to be weighed together to reach a final verdict. Moreover, the kinds of decision methods postulated by psychologists to explain experimental data, like LEX, EAM, and others, are naturally described as making use of reasons. For example, when an experimental subject uses LEX to decide between two options, choosing *A* over *B* when she hits upon a desirable feature that *A* has but *B* lacks, it's natural to say her reason for choosing *A* is that it has this desirable feature (while *B* lacks it). Indeed, psychologists frequently describe LEX as "one-reason decision making" (Newell & Shanks, 2003; Gigerenzer et al., 2008), in contrast to the "many-reason decision making" of EAM (Lee & Cummins, 2004, 350).

In fact, the decision method most commonly associated with Bayesianism suggests an account of reasons and their weighing that vindicates precisely the present proposal's treatment of Jane's case. While in §2 we rejected the assumption that calculating expected utilities is the only rational way to make a decision, we allowed that it can be *a* rational way. And if one decides in this way, it's natural to say a possible state of affairs is a reason to *A* to the extent *A*'s outcome given that state of

²⁰Compare Richard Jeffrey's (1970) dismissal of talk about full belief as opposed to degrees of belief, and David Christensen's (2004) diminution of full belief-based reasoning methods like *reductio*.

affairs has higher utility, and to the extent that state of affairs is probable.²¹ If you are contemplating buying fire insurance, that there might be a fire is a reason to buy insurance since this eventuality has higher utility if one buys insurance; and this reason is weightier the more likely a fire is to occur. Similarly, that there might not be a fire is a reason against buying insurance, one that is weightier the more credence one gives to there being no fire. These reasons can then be reconciled by the familiar additive formula of expected utility. Applied to Jane's case, this account delivers precisely the claims the present proposal needs: that it might rain is a reason for her to carry an umbrella, one that carries sufficient weight given her 0.4 credence in rain.

Emboldened by the ease with which Bayesianism yields the above account of reasons and their weighing, Bayesians might switch tacks. Rather than dismiss the discourse of reasons, they may instead say that Bayesianism is the be-all and end-all in the theory of practical reasons.

But the easy victories end with modal reasons, reasons to do with what might happen. Often our reasons are instead vanilla, non-modal facts, as when I choose the no-frills toothpaste because it's cheaper. There is actually a long Bayesian tradition of theorizing about vanilla, non-modal reasons in the epistemic domain, under the rubric of 'confirmation theory'. Indeed, the fact that Bayesians have worked hard to analyze epistemic reasons, generating a vast literature on confirmation theory,²² puts them in an even more uncomfortable position if they try to dismiss theorizing about practical reasons. This makes going back to dismissal in the face of vanilla reasons awkward. But it also hints at a way forward.

Translating from confirmation talk into the lingo of epistemic reasons, the classic Bayesian proposal is that E is a reason to believe H to the extent that $p(H|E) > p(H)$. So a natural Bayesian proposal is that R is a reason to choose A to the extent that $EU(A|R) > EU(A)$. Unfortunately, hard lessons from confirmation theory block this way forward.

One obstacle is that the classic problem of old evidence (Glymour, 1980) becomes the problem of known reasons: $EU(A|R) = EU(A)$ when $p(R) = 1$, so R fails to be a reason for A when R is already a given. This problem is especially poignant for

²¹We might also want to distinguish reasons for from reasons against, which we could do by distinguishing outcomes with positive utility from those with negative utility. Bayesian tradition has it that there is no real distinction between positive and negative utility, since preferences are unaffected by positive linear transformations of utility. But empirical research shows that actual people do distinguish between gains and losses (Kahneman & Tversky, 1979), and there is no harm in introducing a natural zero-point or status quo for the purposes of building a theory of reasons on top of our existing decision theory.

²²See (Fitelson, 2001) and (Weisberg, 2011, §6) for overviews.

us, since known reasons are what we're after.²³ Bayesian responses to the problem of old evidence can be repurposed here, but the dialectic bottoms out in a similar manner. For example, Bayesians might reply that we shouldn't assign $p(R) = 1$ since nothing is certain for fallible agents like us. But whether R is a reason to A shouldn't hang on the fact that we are fallible about R 's truth; a more ideal agent who is entitled to be certain of R would rightly regard R as a reason to A . Moreover, matters of degree still come out wrong; the extent to which R is a reason to A should not decrease as R 's probability increases, yet $EU(A|R) - EU(A)$ will decrease. One might try to refine the account, correcting for the increased probability of R by using $p(-R)$ as a normalizing factor. We would then measure R 's strength as a reason to A by the formula $[EU(A|R) - EU(A)]/p(-R)$. But this proposal runs afoul of counterexamples like those that plague the analogous move in confirmation theory. (Christensen, 1999)

A related problem is that this account can only capture all-things-considered reasons, missing out on prima facie reasons. Thus it cannot regulate the deliberative process of balancing complementary and countervailing prima facie reasons to arrive at an all-things-considered assessment. The problem arises in both the epistemic and practical domains. Sometimes E is a reason to believe H that is defeated by some further consideration, D . Likewise, R can be a reason to choose A that is defeated by some D . In both cases, when the defeater D is given the Bayesian analysis says that there simply is no reason present. In the epistemic case, $p(H|E) = p(H)$ when $p(D) = 1$ for any complete defeater D . And in the practical case, $EU(A|R) = EU(A)$ when $p(D) = 1$. So even if the problem of known reasons were resolved, the account would not deliver what we want. It does not identify a deliberating agent's prima facie reasons and counter-reasons, and then say how these reasons should be balanced to generate an all-things considered judgment or decision. It assumes instead that she has already arrived at the correct all-things-considered assessment in the course of updating on her earlier evidence, including the defeater D .

The source of these problems is similar to the source of the problem of old evidence. The proposed analysis uses the agent's present credences as the probabilities by which reason-support is measured. Intuitively though, we are trying to capture more a priori

²³Notice also, even Bayesians will agree that conditionalizing on the fact that it *might* rain will not affect Jane's expected utility for carrying an umbrella. So Bayesians would at least need a disjunctive account to handle both modal and vanilla reasons.

connections, connections that don't already take account of defeating information, which the the agent's present credences already have accounted for.²⁴

Bayesians of an objectivist bent may have recourse here, since they can appeal to logical or evidential probabilities, conditioned on some subset of the subject's knowledge that does not already include the problematic defeater *D* or known reason *R*. Similarly, RKP's proponents can appeal to epistemic probabilities. As the literature on the problem of old evidence has exposed though, it's hardly clear what subset of the agent's knowledge we should take as given when measuring intuitive confirmation relations (Maher, 1996; Christensen, 1999). And the same problem will arise for attempts to measure practical reasons using objective/epistemic probabilities. Objective Bayesians and RKPers find common ground here, since solving this problem is a project they can pursue together. For the purposes of the present discussion though, the point is that a complete theory of how practical reasons should be weighed in deliberation cannot just be read off of Bayesian decision theory.

Other decision methods we canvassed in §2 are more forthcoming with guidance on vanilla, non-modal reasons though. When deciding by a method like LEX, the natural thing to say (and what psychologists do say) is that the reasons the agent considers are facts about which of the desirable features in the search queue each option has. The agent "weighs" these factors according to their order in the search queue, albeit by a very crude mechanism, viz. giving decisive "weight" to the first tie-breaker she finds. When deciding by EAM, the reasons are the same but they are weighed differently, making the weighing metaphor more apt: the reasons for each option are added up until one option passes the pre-determined threshold.

Importantly, both decision methods search through attributes in order of their importance. A standard interpretation identifies importance with the probability that the option with that attribute will be the better option all-things-considered (Lee & Cummins, 2004, 344). Thus importance can be measured by credence—one's credence that having the attribute predicts all-things-considered preferability. Thus, when ordering her search queue, an agent is guided by her credences that various attributes are indicative of bestness. The decision that results will then be knowledge-based, yet weighted by her credences: her knowledge that one restaurant has cheaper prices, for example, is weighted according to her credence that price predicts all-things-considered superiority.

²⁴There is an old Bayesian tradition of trying to work around this problem by appealing to the agent's historical or counterfactual credences, but these attempts fail for well-known reasons (Maher, 1996).

Reasons are assessed one way by LEX, somewhat differently by EAM, and completely differently by the method of expected utility calculation. This might start to look less like an account of reasons and their balancing than a hodgepodge of psychological habits. But this appearance just reflects a central moral of §2. There are different ways of trying to maximize expected utility, and different methods will be appropriate to different circumstances, depending what resources the agent has available and how important the decision is. Because different methods will be reasonable, and because these methods will assess reasons differently, a univocal theory of how reasons ought to be assessed shouldn't be expected. Moreover, we can't expect to have a complete theory until we know more about what methods we do and can use, and when we can reasonably be expected to employ them.

But even if we cannot hope to produce a complete such theory now, we can make progress, as we have here. And importantly, this progress supports the conciliatory approach to RKP taken in this paper, and especially the present proposal's claims about Jane and her reasons.

Where does all this leave us? There's been a lot of back and forth, so let me summarize the main points of this subsection:

- Bayesians cannot dismiss theorizing about reasons and their balancing. Doing so goes against both common sense and empirical psychology. It is also in tension with the extensive body of Bayesian work on confirmation theory.
- Bayesianism actually lends itself quite readily to an account of certain, modal reasons. Moreover, that account vindicates precisely those claims about Jane the present proposal needs.
- Bayesianism does not readily provide a complete theory of reasons, however; vanilla, non-modal reasons are not easily accounted for.
- Other decision methods introduced in §2 lend themselves more readily to modelling these vanilla reasons and how they are balanced.
- No single, uniform account of reasons and their balancing should be expected; different methods of practical reasoning will (rightly) balance reasons differently.

Our discussion thus vindicates the present proposal's assumptions about Jane, while simultaneously allaying Bayesian skeptical concerns, and also fleshing out the ways that knowledge and credence interact in the deliberative process of weighing reasons.

3.2.2. *The Spirit of RKP.* Let's turn now to a worry from the RKP camp, that the spirit of RKP forbids being swayed by any doxastic state that fails to constitute knowledge. Even if Jane does not treat the proposition that it will rain as a reason, treating her credence that it will rain as a reason-weight does employ a doxastic attitude other than knowledge to guide her reasoning. And isn't that against the spirit of RKP?

Not if one endorses RKP for the reasons prominent in the literature. Stanley's (2005) reason for embracing RKP is that it matches intuitions about how people's knowledge varies with stakes. On this view, there is no reason to think the above proposal violates the spirit of RKP: increased stakes will still block knowledge in the high-stakes cases that motivated RKP. We can continue to say that, when much hangs on whether the bank is open tomorrow, Hannah does not know that it will be open tomorrow, so it is not appropriate for her to treat that proposition as a reason for waiting until tomorrow to visit. It is appropriate for her to treat the fact that it *might* be open as a reason for delaying her visit, but if she is reasonable, her credence will not be high enough to give that reason enough weight to warrant postponing her visit until tomorrow.

Hawthorne (2004) and Hawthorne & Stanley's (2008) offer as a reason to endorse RKP that it explains folk appraisals of practical reasoning. A doctor who uses a needle without knowing it safe is blameworthy. Here again, there is no reason to shy away from the above proposal: the doctor still does not know the needle is safe. She does know that it *might* be safe, but she does not know it is safe, so she probably doesn't have the very high credence it would take to outweigh the costs of error in an expected utility calculation.

A third motivation for RKP comes from Williamson's (2000) arguments that knowledge is central to our normative theorizing, and especially his arguments that evidence and knowledge are coextensive. If knowledge is the fundamental ground of theoretical reasoning, it is natural to suppose that it plays the same role in practical reasoning (Hawthorne & Stanley, 2008, 577). Yet here again, there is no tension. Knowledge may be the fundamental ground of theoretical reasoning, with credence guiding its use and impact in ways parallel to the above suggestions for practical reasoning.²⁵

One might still worry that Jane's knowledge does too little work on the present proposal. After all, only a very weak sort of knowledge appears to be operative in her decision-making, namely her knowledge of what may be. This appearance is deceptive though. Jane's knowledge about what may be is both complemented, and shaped, by a good deal of knowledge of what is. For example, Jane's decision is also

²⁵See (Jackson, 2012), especially §5, for another perspective.

influenced by her knowledge that her umbrella will not break and will thus keep her dry if it rains, as well as by her knowledge that it won't blizzard or monsoon, and that rain is not uncommon in her present location and season.

Some of this non-modal knowledge will, if she relies on it as a reason in her deliberations, also be weighted by her credences. Her knowledge that her umbrella will not break, for example, may carry more or less weight depending how confident she is in that knowledge. Thus Jane's reasons for carrying her umbrella can include both the modal reason that it might rain, and the non-modal reason that her umbrella will keep her dry, with each reason weighted according to the credence she gives it. The more credence she gives to the possibility of rain, the more confidence she may need in her knowledge that her umbrella will keep her dry. Or, another example, suppose Jane also knows that she has an important meeting to get to. In view of her knowledge that it might rain, this knowledge supplies an additional reason to carry an umbrella, a reason whose strength increases with her credence that she will make a bad impression if she shows up wet and disheveled.²⁶

Other non-modal knowledge, e.g. that it won't blizzard or monsoon, may have a more distal connection to her decision, such that it would be inappropriate to count these facts amongst her reasons. In Bayesian terms, such knowledge influences her decision by shaping her decision table, excluding possibilities like snowstorms and monsoons (Weatherson, 2012). But it remains true that the knowledge that does supply her reasons would not constitute sufficient reason to carry an umbrella were it not for this further wealth of background knowledge she possesses.

The present proposal thus meshes well with the idea that knowledge plays a central, foundational role in both theoretical and practical reasoning. To identify Jane's relatively weak, modal knowledge as a reason is not to render her other, stronger knowledge irrelevant or inessential. Rather, it is merely to identify one aspect of her knowledge that has a certain, proximal relationship to her decision such that it counts as one of her reasons.

3.3. Two Proposals, One Solution. We've now seen two proposals for reconciling the knowledge- and credence-based views in Jane's case. According to the first, Jane's credence supplies her reason for carrying an umbrella, but this is okay by RKP because her credence constitutes knowledge. According to the second, her credence

²⁶These examples expose another way in which the Bayesian account of "might" reasons sketched earlier (§3.2.1) is only partial. Modal reasons can interact with non-modal, vanilla reasons: Jane's modal reason that it may rain can interact with non-modal reasons, such as that carrying an umbrella is a nuisance, and that she has a meeting later.

does not supply her reason for acting, but rather the weight she gives to her reason, which is supplied instead by her knowledge that it may rain. Which account is right?

Both accounts are correct, they merely describe different cases. In one case Jane reasons on the basis of the probability of rain, in the other she reasons on the basis of the possibility of rain. Which description of Jane's reasoning is correct depends on the details of her psychology, i.e. on what method of practical reasoning she uses. As we've noted a few times now, different methods of practical reasoning are appropriate to different circumstances. Because different methods make use of different resources, Jane will rely on different reasons in different circumstances.

In either case, both Jane's credences and her knowledge play central roles, though they operate differently depending on the case. Her credence in rain is operative in both cases, though it plays different roles; in the first case her credence in rain supplies a reason, while in the second it supplies a reason weight. Different pieces of knowledge are operative in each case though; Jane relies on her knowledge about the probability of rain in the first case, and on her knowledge of the possibility of rain in the second.

Which way Jane ought to reason in deciding whether to carry an umbrella will depend on how important the decision is, and on what resources are available to her, including what she knows. If Jane knows the probability of rain with sufficient precision, and she has the time and other resources necessary to implement an overt Bayesian calculation, she might carry an umbrella on the grounds that there's a decent chance of rain. But if she doesn't have the time for such reasoning, or if she doesn't know enough about the relevant probabilities, she might need to reason differently; she might simply entertain the possibility of rain and respond by choosing to carry an umbrella (though she would respond differently if her credence in rain were lower).

When does Jane count as relying on one reason as opposed to the other? Some cases are easy to call. If Jane is a trained decision theorist and she calculates expected utilities on the back of an envelope, she clearly relies on the probability of rain as a reason. In this case RKP requires that her credence in rain constitute knowledge. If, on the other hand, she merely imagines the possibility of rain with a vividness proportional to her credence, and that vividness is sufficient to overcome her distaste for toting the extra baggage, it's plausible to say the possibility of rain is her reason. And in that case, RKP requires only that she know it may rain.

Other cases will be harder to call. Ultimately, the question comes down to what mental representations Jane uses in her decision method, and how she uses them. Mapping the discourse of reasons onto Jane's psychology will likely sometimes be

a fraught and vague matter. But provided we are not skeptical of reasons talk altogether, this is a practical problem of application, not a cause for skepticism about the proposals offered here. Our two proposals cover Jane's case whichever description is ultimately correct once the details of her psychology are fleshed out, whether her reason turns out to be that it may rain or that there is a decent chance of rain.

4. STEPPING BACK

Much of Bayesians' dissatisfaction with RKP ultimately stems, I suspect, from the fact that Bayesians see themselves as having a complete and time-tested decision theory embodied in EUP, while RKP barely offers a fragment of such a theory. Given an agent's beliefs and desires, Bayesianism tells her which choices are rational. But RKP only tells her what assumptions she may rely on in making her decision; it does not tell her which of these assumptions count as reasons for a given option and which count against, nor does it tell her how strongly they count, nor how these reasons ultimately combine to yield a final recommendation.²⁷ Bayesians may even feel that the fragment of theory provided by RKP is actually untestable; without auxiliary hypotheses about how reasons combine, a proposal like RKP makes no definite predictions about what a subject ought to choose.

From that vantage point, the central contribution of our discussion may appear to be the offer of such auxiliary hypotheses, thereby narrowing the gap between the achievements of the credence- and knowledge-based approaches. We developed some theory, especially in §2 and §3.2, of when a subject should treat *R* as a reason to *A*, of when *R* is a stronger or weaker reason, and of how she ought ultimately to reconcile these reasons in reaching her decision. What theory we developed was by no means complete. We examined only a few candidates for rational methods of practical reasoning, and we said very little about when it is rational to engage each method. But we made a beginning, and the beginning we did make went a long way towards answering prominent Bayesian criticisms.

But I think this is the wrong view of things, and it lets Bayesians off the hook much too easily. For another thing our discussion exposed is that Bayesians do not actually have a theory of their own on these matters. Bayesians have a theory of which choice is rational, but not a theory of the reasons for action that might lead one to that choice. Bayesianism is not a theory of what one should count as a reason for/against a given option, not to mention a theory about how those reasons interact while working towards a final decision (except that the final outcome should agree with EUP). As we saw in §3.2, part of such a theory can be read off the Bayesian formalism,

²⁷Cf. Ichikawa (2012) on the "rationalizing" relation.

but the result is hardly a complete theory. EUP suggests a theory for certain kinds of modal reasons, but even that bit of theory only applies to some agents, those who reason by expected utility calculation. And as we noted in §2, we often are not such agents. In fact, thinking about the sorts of examples of practical reasoning discussed in §2—the sort Bayesians criticize RKPers for appealing to—suggested ways of going beyond the bit of theory that we can read off the Bayesian formalism.

So I advocate the following view of things instead. Our discussion has exposed that RKP and EUP are theories of different things, albeit related ones. EUP is a constraint on what choices one may make, while RKP is a constraint on what assumptions one may rely on in arriving at those choices. Acknowledging this difference helped us unearth questionable presuppositions underlying Bayesian critiques of RKP, thereby defusing those critiques. Making a decision by appealing to reasons does not preclude one from engaging in a method of practical reasoning that maximizes expected utility (§2), nor does it preclude one's middling credences from entering into one's reasoning (§3).

More broadly though, seeing how RKP and EUP are theories of different things helps us to see how Bayesian impatience with knowledge-based proposals can be misplaced. Advocates of knowledge-based proposals are not playing catch-up with Bayesians. They are, instead, offering a piece of theory in a domain where neither party has satisfactorily established itself: a theory of what one's reasons are, what they are reasons for, and how those reasons interact. By prompting us to think about decision-making in terms of reasons, proposals like RKP push us to explore this uncharted (or under-charted: see (Horty, 2012)) theoretical terrain.

Still, it would be unfair to say that Bayesian critiques of proposals like RKP are of no service. They draw our attention to the work proposals like RKP leave undone, the work of saying when *R* is a reason to *A* and how such reasons interact. They also provide a criterion of correctness for a completed such theory, namely that it agree with EUP's recommendations (at least in ideal circumstances). Finally, Bayesian critiques also point to resources that contribute to the completion of such a theory. The psychological theory of credence and utility, as well as the method of expected utility calculation, suggest how some agents can assess and weigh some kinds of reasons to arrive at good choices, at least in some circumstances. The task of augmenting this partial theory to handle other kinds of reasons, and other circumstances, is a task both camps face. They might even face it together.

REFERENCES

- Allais, Maurice. 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica* 21(4): 503–46.
- Baron, Jonathan. 2007. *Thinking and deciding*. Cambridge University Press.
- Bröder, Arndt. 2000. Assessing the empirical validity of the “take-the-best” heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26(5): 1332–46.
- Bröder, Arndt & Ben R. Newell. 2008. Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making* 3(3): 205–14.
- Brown, Jessica. 2008a. Knowledge and practical reason. *Philosophy Compass* 3(6): 1135–52.
- Brown, Jessica. 2008b. Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs* 42(2): 167–89.
- Christensen, David. 1992. Confirmation holism and bayesian epistemology. *Philosophy of Science* 59(4): 540–57.
- Christensen, David. 1999. Measuring confirmation. *Journal of Philosophy* 96(9): 437–61.
- Christensen, David. 2004. *Putting logic in its place: Formal constraints on rational belief*. Oxford University Press.
- Christensen, David. 2007. Does murphy's law apply in epistemology? self-doubt and rational ideals. *Oxford Studies in Epistemology* 2: 3–31.
- Douven, Igor. 2008. Knowledge and practical reasoning. *Dialectica* 62(1): 101–18.
- Ellsberg, Daniel. 1961. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics* 75(4): 643–69.
- Fantl, Jeremy & Matthew McGrath. 2002. Evidence, pragmatics, and justification. *The Philosophical Review* 111(1): 67–94.
- Fantl, Jeremy & Matthew McGrath. 2007. On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research* 75(3): 558–89.
- Fitelson, Branden. 2001. *Studies in bayesian confirmation theory*. University of Wisconsin—Madison dissertation. <http://fitelson.org/thesis.pdf>.
- Garber, Daniel. 1983. Old evidence and logical omniscience in bayesian confirmation theory. *Testing Scientific Theories, Midwest Studies in Philosophy* X: 99–131.
- Gettier, Edmund L. 1963. Is justified true belief knowledge? *Analysis* 23: 121–3.
- Gigerenzer, Gerd & Henry Brighton. 2009. Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science* 1(1): 107–43.

- Gigerenzer, Gerd & Daniel G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103(4): 650–69.
- Gigerenzer, Gerd, Laura Martignon, Ulrich Hoffrage, Jörg Rieskamp, Jean Czerlinski & Daniel G. Goldstein. 2008. One-reason decision making. In Charles R. Plott & Vernon L. Smith (eds.), *Handbook of experimental economics results*, vol. 1, 1004–17. Elsevier.
- Gigerenzer, Gerd & Reinhard Selten (eds.). 2001. *Bounded rationality: The adaptive toolbox*. The MIT Press.
- Gigerenzer, Gerd, Peter M. Todd & The ABC Research Group. 1999. *Simple heuristics that make us smart*. Oxford University Press.
- Glymour, Clark. 1980. *Theory and evidence*. Princeton University Press.
- Hájek, Alan & Stephan Hartmann. 2010. Bayesian epistemology. In Jonathan Dancy, Ernest Sosa & Matthias Steup (eds.), *A companion to epistemology*, vol. 4, 93–104. Wiley-Blackwell.
- Hawthorne, James A. 2005. Degree-of-belief and degree-of-support: Why bayesians need both notions. *Mind* 114(454): 277–320.
- Hawthorne, John. 2004. *Knowledge and lotteries*. Oxford University Press.
- Hawthorne, John & Jason Stanley. 2008. Knowledge and action. *The Journal of Philosophy* 105(10): 571–90.
- Hilbig, Benjamin E. 2010. Reconsidering “evidence” for fast-and-frugal heuristics. *Psychonomic Bulletin & Review* 17(6): 923–30.
- Hogarth, Robin M. & Natalia Karelaia. 2007. Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review* 114(3): 733–58.
- Horty, John F. 2012. *Reasons as defaults*. Oxford University Press.
- Horwich, Paul. 1982. *Probability and evidence*. Cambridge University Press.
- Horwich, Paul. 1993. Wittgensteinian bayesianism. *Midwest Studies in Philosophy* 18(1): 62–75.
- Hurley, Terrance M. & Jason Shogren. 2005. An experimental comparison of induced and elicited beliefs. *The Journal of Risk and Uncertainty* 30(2): 169–88.
- Ichikawa, Jonathan Jenkins. 2012. Knowledge norms and acting well. *Thought: A Journal of Philosophy* 1(1): 49–55.
- Jackson, Alexander. 2012. Two ways to put knowledge first. *Australasian Journal of Philosophy* 90(2): 353–69.
- Jeffrey, Richard C. 1965. *The logic of decision*. University of Chicago Press.
- Jeffrey, Richard C. 1970. Dracula meets wolfman: Acceptance vs. partial belief. In Marshall Swain (ed.), *Induction, acceptance, and rational belief*, 157–85. D. Reidel.

- Joyce, James M. 1999. *The foundations of causal decision theory*. Cambridge University Press.
- Kahneman, Daniel & Amos Tversky. 1973. On the psychology of prediction. *Psychological Review* 80(4): 237–51.
- Kahneman, Daniel & Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2): 263–91.
- Lee, Michael D. & Tarrant D. R. Cummins. 2004. Evidence accumulation in decision making: Unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review* 11(2): 343–52.
- Lipton, Peter. 2004. *Inference to the best explanation (second edition)*. Routledge.
- Maher, Patrick. 1993. *Betting on theories*. Cambridge University Press.
- Maher, Patrick. 1996. Subjective and objective confirmation. *Philosophy of Science* 63(2): 149–74.
- McGrew, Timothy. 2003. Confirmation, heuristics, and explanatory reasoning. *British Journal for the Philosophy of Science* 54(4): 533–67.
- Moss, Sarah. 2013. Epistemology formalized. *The Philosophical Review* 122(1): 1–43.
- Newell, Ben R. 2005. Re-visions of rationality. *Trends in Cognitive Science* 9(1): 11–5.
- Newell, Ben R. & Michael D. Lee. 2011. The right tool for the job? Comparing an evidence accumulation and a naïve strategy selection model of decision making. *Journal of Behavioral Decision Making* 24(5): 456–81.
- Newell, Ben R. & David R. Shanks. 2003. Take the best or look at the rest? Factors influencing “one-reason” decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29(1): 53–65.
- Newell, Ben R., Nicola J. Weston & David R Shanks. 2003. Empirical tests of a fast-and-frugal heuristic: Not everyone “takes-the-best”. *Organizational Behavior and Human Decision Processes* 91(1): 82–96.
- Nozick, Robert. 1969. Newcomb’s problem and two principles of rational choice. In Nicholas Rescher (ed.), *Essays in honor of Carl G. Hempel: A tribute on the occasion of his sixty-fifth birthday* Synthese Library, 114–5. D. Reidel.
- Oaksford, Mike & Nick Chater. 1998. *Rationality in an uncertain world*. Psychology Press.
- Oaksford, Mike & Nick Chater. 2007. *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.
- Oaksford, Mike & Nick Chater. 2009. Précis of *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. *Behavioral and Brain Sciences* 32(1): 69–120.
- Okasha, Samir. 2000. Van Fraassen’s critique of inference to the best explanation. *Studies in the History and Philosophy of Science* 31(4): 691–710.

- Payne, John W., James R. Bettman & Eric J. Johnson. 1993. *The adaptive decision maker*. Cambridge University Press.
- Savage, Leonard J. 1954. *The foundations of statistics*. Wiley Publications in Statistics.
- Schiffer, Stephen. 2007. Interest-relative invariantism. *Philosophy and Phenomenological Research* 75(1): 188–95.
- Simon, Herbert. 1956. Rational choice and the structure of the environment. *Psychological Review* 63(2): 129–38.
- Simon, Herbert. 1982. *Models of bounded rationality, volume 1: Economic analysis and public policy*. The MIT Press.
- Stanley, Jason. 2005. *Knowledge and practical interests*. Oxford University Press.
- Tversky, Amos. 1972. Elimination by aspects: A theory of choice. *Psychological Review* 79(4): 281–99.
- von Neumann, John & Oskar Morgenstern. 1944. *Theory of games and economic behavior*. Princeton University Press.
- Weatherson, Brian. 2012. Knowledge, bets, and interests. In Jessica Brown & Mikkel Gerken (eds.), *Knowledge ascriptions*, 75–103. Oxford University Press.
- Weirich, Paul. 2004. *Realistic decision theory: Rules for nonideal agents in nonideal circumstances*. Oxford University Press.
- Weisberg, Jonathan. 2011. Varieties of bayesianism. In Dov M. Gabbay, Stephan Hartmann & John Woods (eds.), *Handbook of the history of logic*, vol. 10, 477–552. Elsevier.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford University Press.
- Yalcin, Seth. 2007. Epistemic modals. *Mind* 116(464): 983–1026.
- Yalcin, Seth. 2012. Bayesian expressivism. In *Proceedings of the aristotelian society*, vol. 112 2, 123–60.