



The key to the knowledge norm of action is ambiguity

Patricia Rich¹ 

Received: 1 June 2020 / Accepted: 17 May 2021

© The Author(s) 2021

Abstract

Knowledge-first epistemology includes a knowledge norm of action: roughly, act only on what you know. This norm has been criticized, especially from the perspective of so-called standard decision theory. Mueller and Ross provide example decision problems which seem to show that acting properly cannot require knowledge. I argue that this conclusion depends on applying a particular decision theory (namely, Savage-style Expected Utility Theory) which is ill-motivated in this context. Agents' knowledge is often most plausibly formalized as an ambiguous epistemic state, and the theory of decision under ambiguity is then the appropriate modeling tool. I show how to model agents as acting rationally on the basis of their knowledge according to such a theory. I conclude that the tension between the knowledge norm of action and formal decision theory is illusory; the knowledge-first paradigm should be used to actively select the decision-theoretical tools that can best capture the knowledge-based decisions in any given situation.

Keywords Knowledge · Decision-making · Ambiguity · Knowledge-first epistemology · Imprecise probability · Expected utility · Uncertainty

1 Introduction

Many epistemologists have endorsed a *knowledge norm of action*, according to which agents should base their actions only on what they know. This action norm is one part of a broader knowledge-first paradigm, which takes the concept of knowledge as basic and additionally includes knowledge norms of belief and assertion and the thesis that one's evidence coincides with one's knowledge (Williamson 2000); see e.g. (Smithies 2012; Greenough and Pritchard 2009; McGlynn 2014) for discussion. According to

This article belongs to the topical collection "Knowledge and Decision", edited by Moritz Schulz, Roman Heil, Jakob Koscholke, and Patricia Rich.

✉ Patricia Rich
patricia.rich@uni-bayreuth.de

¹ Department of Philosophy, University of Bayreuth, Bayreuth, Germany

proponents, this paradigm achieves great explanatory power and avoids the notorious Gettier problem just by putting knowledge at the center of everything.

A *prima facie* concern about the knowledge norm of action is immediate: what about the fact that humans are perpetually uncertain, but must decide nonetheless? Schiffer (2007) framed this problem as one of competition with the widely-deployed Bayesian account of rational choice under uncertainty, Expected Utility Theory (Savage 1954). Proponents of the knowledge norm clearly need a story about this well-developed existing account, either rejecting it or explaining when and how their norm can also play a role. The significance of the knowledge norm would be greatly diminished if it were to be relevant only in situations of certainty. Ideally, there would be a well-worked-out knowledge-based decision theory with the kind of general applicability that Bayesian Expected Utility Theory is taken to have (this special issue reflects this desideratum).

The key first step in defending the knowledge norm is to recognize that knowledge is not (necessarily) just the traditional outright attitude, but is in fact compatible with uncertainty. That is, epistemic states reflecting uncertainty, such as partial beliefs or credences, can constitute knowledge. Weisberg connects credences to beliefs about chances, which could constitute knowledge about chances (Weisberg 2013). Moss argues in great detail that agents can have not only knowledge of chances, but knowledge which is itself, directly, probabilistic (Moss 2013, 2017). For example—according to this view—an agent can know that it will probably rain, know that rain is between .7 and .8 likely, or have a .65 credence in rain which constitutes knowledge.

Despite this important insight, uncertainty still presents a puzzle for the knowledge norm because of the existing alternative mathematical decision theory. Expected Utility Theory appears very different from the knowledge norm, and the two normative hypotheses would plausibly have incompatible consequences. In practice, Expected Utility Theory has been taken as a source of objections to the knowledge norm of action. The standard strategy has been to defend the knowledge norm by arguing that its normative implications are in fact compatible with, or even coincide with, those of Expected Utility Theory (see e.g. Hawthorne and Stanley 2008; Weisberg 2013; Moss 2013).¹

Mueller and Ross consider this tradition of defenses of the knowledge norm of action and claim to show that the norm cannot be correct (Mueller and Ross 2017). Their thorough analysis is based on an example in which the inputs to rational choice according to Expected Utility Theory seem clearly not to count as knowledge. This is because the belief (here, a credence) on which the agent bases their (expected utility maximizing) decision is right on the borderline of supporting one choice over the other, but is intuitively unsafe. Therefore, the argument goes, decisions need not be based on knowledge.

This demonstration seems forceful because it is extremely plausible in this case that the decision-supporting belief Mueller and Ross identify would not count as knowledge. It is critical to note, however, that the argument essentially depends on the presupposition that Expected Utility Theory is the full mathematical story of rational

¹ Similarly, Jackson (2012) defends the knowledge norm as governing only full belief, without rejecting Expected Utility Theory as capturing rationality.

choice. This is a real problem: Expected Utility Theory may be the standard way to represent rational choices in epistemology, but it is certainly not the only—and often not the best—way. Indeed, from the perspective of the knowledge norm of action, a different decision theory may be required.

Importantly, while proponents of the knowledge norm of action have not moved to reject Expected Utility Theory, the assumption that rationality means maximizing expected utility with respect to precise credences is not considered critically or provided with motivation relative to the knowledge-first paradigm, either. Although Expected Utility Theory is commonly employed within epistemology with little or no argument, it was designed to apply only in particular settings [famously, in Savage's "small worlds" (Savage 1954, Chap. 2.5); see e.g. Chew and Sagi 2008; Binmore 2008; Mousavi and Gigerenzer 2017] and is generally controversial (the critiques and paradoxes are endless, but see e.g. Ellsberg 1961; Tversky 1975; Sen 1977; Buchak 2013). The problem is particularly relevant in the context of the debate over the knowledge norm of action. This is because Expected Utility Theory is especially poorly motivated from the perspective of the knowledge norm. In fact, considerations in support of the knowledge norm can be seen as directly speaking against Savage-style Expected Utility Theory for many practical purposes, and for cases like that in Mueller and Ross' key example.

Given Mueller and Ross' description and discussion of this key example, I argue that the best candidate for the agent's knowledge will be an epistemic state reflecting a situation of ambiguity. Hence, a theory of decision making under ambiguity would best characterize the knowledge-based decisions from a mathematical perspective, and arguably characterize *rational* choices in this setting. I illustrate the possibilities here using one such theory which is especially simple and amenable to use in epistemology.

The existence of this more natural alternative picture means that the knowledge norm can be defended against Mueller and Ross' challenge. Their argument in fact leaves room for two possible conclusions: either knowledge is not the norm of action, or Expected Utility Theory is often the wrong way to identify the right actions. The latter conclusion is clearly preferable from the knowledge-first perspective, and is in general appealing once the broader range of decision-theoretic possibilities is considered. The knowledge norm of action is a much stronger hypothesis once it is paired with its own mathematical characterization of rational decisions.

To be clear, the argument in this paper is meant to supplement, and not to refute or replace, defenses of the knowledge norm articulating a particular relationship, such as a division of labor, between knowledge-based reasoning and formal decision theory (e.g. Weatherson 2012; Weisberg 2013; Feldbacher-Escamilla 2018). The knowledge norm of action is often expressed in terms of what agents may take as reasons for actions (Hawthorne and Stanley 2008), and given this formulation, some story about the relationship between acting on reasons and conforming to a formal decision theory is needed. The purpose of this paper is to show that prospects for telling such a story are much better if the formal theory is chosen in accordance with the agent's level of knowledge in any given case. The focus is on explaining why and how the theory of decision under ambiguity should often be used instead of Expected Utility Theory, no matter exactly what role the formal theory is taken to play in the knowledge-first

paradigm. Similarly, I will assume for present purposes that separate debates regarding the viability of the knowledge norm of action can be resolved.

This paper is organized as follows: Section 2 presents Mueller and Ross' main argument that the rational and knowledge-based decisions need not coincide (2017). Section 3 considers in turn the key premises in their argument, and delineates the cases in which each should be accepted; the two premises are argued to be compelling in distinct and non-overlapping types of cases. The crucial observation regarding the second premise is that Savage's Expected Utility Theory is treated as the sole available formal decision theory, but this ignores the further development of numerous decision theories for situations in which there is greater uncertainty. Section 4 provides an overview of an alternative general theory which includes many familiar proposals as special cases. It shows how the theory can be applied to Mueller and Ross' example to derive unique rational choices in that example. Importantly, it would be natural to characterize agents' knowledge as the epistemic inputs to such a theory. Section 5 steps back to reflect on the morals for the knowledge norm of action and for decision theory. An important takeaway is that the knowledge norm of action can be used to adjudicate among the countless available decision theories. The proponent of knowledge norms should use the strength of their paradigm to play an active role in this old, but still tough, debate with extensive impact on epistemology.

2 Elvis' umbrella dilemma

Mueller and Ross (2017) take themselves to have demonstrated that the knowledge norm cannot be correct. The key example for us is the following choice problem:

Elvis' Shoes Elvis is concerned about his famous blue suede shoes. He considers whether it will rain today or not; rain would damage the shoes. Elvis can however take a preventative measure: he can buy an umbrella, which would cost money but protect his shoes in the event of rain. Elvis must decide whether to buy the umbrella.

Table 1 shows Elvis' utility losses for each of the two actions he considers (buying the umbrella or not) and for each of the two relevant states of the world (raining and not raining). The expected utility of Elvis' two actions would be equal if the probability of rain ($P(\text{Rain})$) were .39992.² This probability would make the two options equally good. A higher probability of rain would make the expected utility of buying the umbrella greater, while a lower probability of rain would give not buying it greater expectation.

For the example to do its job, there needs to be a conflict between the implications of decision theory and of the knowledge norm. Mueller and Ross identify "borderline probability cases" as the important test cases; in such cases, the belief that the agent acts on is just at the borderline of supporting a particular action, and so it will often be questionable whether the agent has knowledge supporting their action. Therefore,

² The precise indifference probability is .399920016, but I follow Mueller and Ross in rounding to 5 decimal places.

Table 1 Elvis fears for his shoes

	Rain	No rain
Buy umbrella	-20	-20
Do nothing	-50.01	0

the authors ask us to imagine that Elvis has credence of .4 in rain, where his credence matches the epistemic probability of rain, that is, the probability given all of Elvis' knowledge (or equivalently, his evidence). We can then reconstruct Mueller and Ross' argument against the knowledge norm of action as follows:

- P1*: Elvis does not have (probabilistic) knowledge that $P(\text{Rain}) \geq .39992$ (so his knowledge does not (unequivocally) support buying the umbrella).
P2: To determine what is appropriate for Elvis to do, we should identify the action that maximizes subjective expected utility relative to a .4 credence in rain (namely, buying the umbrella).
C: In buying the umbrella, Elvis acts appropriately without acting on his knowledge.

The argument for *P1* is that knowledge—even of the probabilistic variety—must be safe, but Elvis' credence is not. Safety is often taken to be a necessary condition for knowledge; its purpose is to rule out luckily-true beliefs by requiring that the agent also believes truly in nearby possible worlds (see Sosa 1999 for discussion of how to formulate the condition). Here, since humans cannot perfectly track epistemic probabilities, there are intuitively close possible worlds in which the evidential probability is, say, .39, but Elvis' credence is still .4. In these worlds, Elvis would falsely believe that he should buy the umbrella. But then he cannot know in this world that he should.

Regarding *P2*, we should notice that it is taken for granted that the requirements of rationality are here captured by Subjective Expected Utility Theory (as developed by Savage (1954)). Briefly, this theory says that we can represent the agent via a precise, probabilistic credence function over events and a utility function over outcomes such that the agent's choices have maximal mathematical expectation given the credence and utility functions. As Mueller and Ross explain, Schiffer (2007) points to this theory as a plentiful source of prima facie counter-examples to the knowledge norm of action, and much subsequent literature can be understood as attempting to address this objection. For the most part, these attempts do not question the force of Expected Utility Theory, but rather try to show how the knowledge norm could be compatible with it (see e.g. Hawthorne and Stanley 2008; Weisberg 2013; Moss 2013; Jackson 2012). Mueller and Ross argue in detail, however, that none of these attempts succeeds in showing that a knowledge-based construal of Elvis' decision can recapture the expected utility maximizing choice behavior. The biggest departure from the expected utility assumption is a knowledge-based variant of this theory that Mueller and Ross themselves discuss and reject, and which I will say more about in Sect. 3.3.

If it is right for Elvis to maximize expected utility relative to his .4 credence in rain, however, and if Elvis' knowledge cannot support this choice, then it can be right to act on the basis of epistemic states which do not constitute knowledge. The knowledge-based and credence-based actions would fail to coincide, and *P2* asserts that particular credence-based actions are the appropriate ones. In what follows, I will argue that this

argument fails to establish C . This is because it is doubtful that we should accept both of the premises at the same time. Once we fill in the details of the example so that $P1$ holds, $P2$ is especially poorly motivated. Cases in which $P2$ is easy to accept, in contrast, are those in which $P1$ is less plausible.

3 Alternatives

3.1 Preliminaries

For the sake of clarity in my argument, it is useful to introduce some basic notation to distinguish between the various types of probability under consideration. We will need to distinguish between credences and “true” probabilities, probabilities which count as knowledge and those that don’t, and so forth. As noted in the introduction, there are different ways of combining knowledge and uncertainty. Here, I adopt Moss’ (2017) approach under which credences and other expressions of partial belief can directly constitute knowledge. Hence, rather than always saying e.g. that Elvis knows *that* the probability is like so, a probabilistic expression will directly represent his knowledge in the same way that probabilistic credences directly represent beliefs. This helps to streamline the exposition.

Given this, the first type of probability that we need is probabilistic knowledge, i.e. probability expressions which we take to constitute knowledge. Let P_K denote probabilistic knowledge. Here, P_K is always relative to Elvis, our subject. The distinction between knowledge and belief or credence is critical to Mueller and Ross’ argument, and so we need to separately refer to probabilistic expressions which represent Elvis’ beliefs; these are denoted by P_B . Both probabilistic knowledge and belief can be precise, but need not be. So, for example, we might say things like $P_B(\text{Rain}) = .4$, $P_K(\text{Rain}) \in (.2, .8)$; $\neg P_K(\text{Rain}) \geq .39992$, and so forth. The use of \neg here and throughout is convenient shorthand; $\neg P_K(\text{Rain}) \geq x$ indicates that it is not the case that the agent knows rain to be at least x likely.

Especially for the sake of discussing whether some P_B is also a P_K , we will also need to refer to both epistemic probabilities and to the probabilistic expressions that we are treating as true or correct. Continuing the pattern, denote these by P_E and P_T , respectively. These latter two permit of more possible meanings than do P_K and P_B , and so I will say a bit more up front about what they can be taken to mean here.

There are various possible relationships between P_T , P_B , and P_E . Here, I do not assume that $P_B(\text{Rain}) = P_E(\text{Rain})$; Elvis’ subjective belief may not coincide with or track the epistemic probability. This is intuitive, and in line with Mueller and Ross’ argument. I take $P_E(\text{Rain})$ to correspond to the P_B that an ideal agent would have, if they had exactly the evidence that Elvis has. This ideal agent may be better at probability calculations or logic than Elvis is, but they are not working with more information. Multiple possibilities regarding what this P_E looks like or how it might be arrived at will be considered here. Similarly, I leave open whether $P_E(\text{Rain})$ should be treated as the $P_T(\text{Rain})$, or whether we might take something else to serve as $P_T(\text{Rain})$. Intuitively, though, P_T must reflect at least as much (and as good) information as P_E , and possibly more.

3.2 Examining P1: Does Elvis know?

3.2.1 What knowledge requires

Most of the defenses of the knowledge norm canvassed by Mueller and Ross (2017) target $P1$, arguing that Elvis does somehow have specific knowledge unequivocally supporting his purchase of the umbrella. Intuitively, in true borderline cases in which tiny differences in probabilities make a difference, it is actually unlikely that Elvis would have such knowledge. His epistemic status depends on the details of the case, however, and even if we accept $P1$, it matters exactly how the premise is true. I delineate the possibilities in this section.

The knowledge-first approach takes the distinction between knowledge and ignorance as basic, thus rejecting analyses of knowledge into a set of necessary and jointly sufficient requirements (Williamson 2017). Most nonetheless agree on three general necessary properties of knowledge, and therefore that the absence of knowledge can occur in at least three general ways. It is useful to imagine Elvis lacking knowledge in each of these ways.

Specifically, I consider how it could fail to be *true* that $P(Rain) \geq .39992$, when we would not attribute to Elvis the belief that this were the case, and the possibility that Elvis' belief is somehow faulty such that it cannot constitute knowledge. Each of these scenarios would make $P1$ true, but they support different lines of reasoning regarding $P2$ (discussed in Sect. 3.3). That is, it matters whether we accept that $\neg P_K(Rain) \geq .39992$ because $\neg P_T(Rain) \geq .39992$, because $\neg P_B(Rain) \geq .39992$, or *despite* $P_B(Rain) \geq .39992$ because of a fault with this belief. In contrast, by imagining the situation in which Elvis has the relevant true and non-faulty belief, we can see how Elvis might most plausibly have the knowledge that Mueller and Ross deny to him. This is useful because it shows when $P2$ would be most compelling.

Of course, Mueller and Ross stipulate that Elvis has a true belief, and lacks knowledge because this belief is faulty in a particular way. This is one way of filling in the details of the example, but it is instructive to step back and reflect on a broader set of possibilities. For one, this helps us to appreciate the interpretational and representational choices required to move from a real (hypothetical) scenario to a formal or abstract representation. Secondly, stepping back helps us to see the crucial point that the use of Expected Utility Theory is just another modeling choice, and one whose propriety depends very much on its fit with the rest of our interpretation of the case (*and* on whether we accept the knowledge norm of action or some other norm).

3.2.2 Truth

Knowledge is widely held to be factive, and so what Elvis can know depends on what the relevant truth or fact is. For probabilistic knowledge, the relevant background facts or truths will be probabilistic as well (Moss 2017). Mueller and Ross stipulate that in Elvis' case, $P_E(Rain) = .4$, where the epistemic probability is the probability given all of Elvis' knowledge. They implicitly take this epistemic probability to play the role of truth, i.e. they treat P_E as equivalent to P_T . I step back a bit and consider what

might be appropriate to take as the truth according to which Elvis' correctness should be judged.

There are two important possibilities for our purposes. The first (**a**) is that the appropriate truth is precise; $P_T(\text{Rain}) = .4$. The second (**b**) is that the appropriate truth is less specific than this; for example, let us imagine that $P_T(\text{Rain}) \in (.35, .45)$. Which of these cases one takes to apply would depend both on the details of Elvis' situation and on one's theoretical commitments or inclinations as an epistemologist.

Case **a** would arguably occur if Elvis has checked a reliable weather website that incorporates a lot of data and provides precise estimates of $P_T(\text{Rain})$; such a website could reveal that $P_T(\text{Rain}) = .4$. Perhaps Elvis could have softer evidence, such as experience with the weather and observations of the clouds, but still have a $P_E(\text{Rain}) = .4$ that would then serve as P_T . We can imagine different ways in which this $P_E(\text{Rain})$ might be identified.

First, suppose that $P_{B_0}(\text{Rain})$ (or, arguably better yet, the corresponding P_{K_0}) is Elvis' prior (perhaps knowledge-constituting) credence in Rain, before he looks up at the sky or thinks much about the weather for the day; let this be precise for whatever reason (e.g. it could reflect a forecast from some time ago). Suppose then that Elvis looks up at the sky and acquires knowledge via perception regarding what he sees, which we can somehow express as a proposition *Sky*. Then perhaps $P_{E_1}(\text{Rain}) = P_{B/K_0}(\text{Rain}|\text{Sky}) = .4$; i.e. $.4$ is the probability of rain calculated by Bayesian-conditioning Elvis' prior precise belief or probabilistic knowledge on his new perceptual knowledge.

An alternative way of identifying a $P_E(\text{Rain})$ which allows Elvis to update using probabilistic knowledge would be to employ Jeffrey Conditionalization³ (Jeffrey 1965, Chap. 11). This lowers the bar to incorporating knowledge by letting Elvis' knowledge have probability less than 1, but is simultaneously more demanding in requiring a much richer set of conditional probabilities. Nevertheless, in a case like the following it might be applied: Suppose again that $P_{B/K_0}(\text{Rain})$ is Elvis' prior belief or probabilistic knowledge regarding rain in Memphis today. Suppose then that Elvis gains relevant probabilistic knowledge in the form of hearing the forecast for the next town over; let's call this $P_{K_1}(\text{AdjacentRain}) = x$. We might then identify $P_{E_1}(\text{Rain}) = P_{B/K_0}(\text{Rain}|P_{K_1}(\text{AdjacentRain}) = x)$, calculated as Jeffrey proposes.⁴

In either of these scenarios, we might say that this perfectly-calculated, precise probability incorporating all of Elvis' relevant knowledge deserves to be treated as the truth relative to which the factivity of Elvis' belief is assessed, i.e. that $P_{E_1}(\text{Rain}) = P_T(\text{Rain})$. Again, the case seems stronger if P_{K_0} rather than P_{B_0} is updated—and hence if the conditional probabilities also count as knowledge—since then P_{E_1} is more thoroughly based on evidence.

A rather different line of argument for an **a**-case interpretation would appeal to extra principles to essentially force a precise P_T . Some argue that although our evidence itself typically fails to pin down a precise probability, we should apply further principles

³ Thanks to an anonymous referee for suggesting this.

⁴ Specifically, for a hypothesis H and evidence E , and given initial and subsequent times 0 and 1, Jeffrey's rule is $P_1(H) = P_0(H|E)P_1(E) + P_0(H|\sim E)P_1(\sim E)$.

such as maximizing entropy (see e.g. Jaynes 1968; Williamson 2010) in order to derive precise probabilities. Someone taking such a view might argue that this precise probability should be treated as the background fact as well. Yet another perspective would be that the true probability of rain would be the evidential probability given all potentially-available evidence, e.g. all the evidence a professional forecaster and experienced sky observer might have. Then we could have $P_T = .4$ without $P_E = .4$; the evidential probability could be less specific (for reasons considered presently), or precise but simply different.

The situation could also well be a **b** case. The available evidence may be sparse and simply underdetermine the probability, and one might not see sufficient justification that any precise probability—even the most seemingly-neutral of the options—could be appropriately treated as *true*.⁵ If Elvis saw a weather report several days ago, or just looked up at the sky, it is very plausible that the probability of rain given his evidence is not precisely determined. Realistically, even if Elvis is an expert and attentive observer of wind and sky, it may be most appropriate to treat his evidence as at best supporting a near-term probability of rain somewhere between .35 and .45.

Consider again the proposals for identifying $P_E(Rain)$, and the observation that it is more reasonable to treat P_E as P_T the more thoroughly P_E is grounded in knowledge. That is, if the evidential probability is the product of conditioning prior probabilistic knowledge on further (perhaps also probabilistic) knowledge—appealing to conditional probabilities which are also known—then we have the strongest grounds for maintaining that this P_E could itself count as knowledge as far as factivity is concerned. If P_B matched such a P_E , then that P_B should count as correct. The difficulty is that the extensive and precise probabilistic knowledge required for this would be extremely hard to come by. Realistically, Elvis may know that if it's very likely to rain in the adjacent town, then rain is more likely in Memphis, but not know exactly what that conditional probability is. His prior probabilistic knowledge may be imprecise, for example if his friend simply warned him that it would probably rain before he left for the trip, and so he has a prior $P_K(Rain) > .5$. Similarly, his new evidence regarding rain may be imprecise, for example if the umbrella salesman tells him that he expects rain. The standard for knowledge is high, and we may well think that in non-artificial cases, it will typically not be reasonable to attribute precise probabilistic knowledge to Elvis. But if Elvis' knowledge is ambiguous, then whatever epistemic probabilities we use it to generate will typically be ambiguous as well. This does not mean that no sensible quantification is possible, but it does mean that we may be able to defend something like $P_E(Rain) \in (.35, .45)$, but nothing narrower than this.

In short, if we follow Mueller and Ross in taking epistemic probability as the standard for knowledge (playing the role of truth), require this epistemic probability to reflect only knowledge (in contrast to, say, subjective beliefs), and then take a realistic view of what Elvis can plausibly know, then both the epistemic and “true” probabilities are unlikely to be precise. Furthermore, Mueller and Ross explain epistemic probabilities as those the agent ought to have (Mueller and Ross 2017, ft. 4). They later avail themselves of an ought-can principle, and throughout the paper make use

⁵ Bertram's Paradox (Bertrand 1888; see Shackel 2007 for discussion) and the well-known problems with logical positivism (Friedman 1991; see Feldbacher-Escamilla 2018 for discussion relative to the knowledge-first program) provide further, more concrete reasons for skepticism.

of the fact that Elvis (as a human) cannot perfectly track the probabilities supported by his evidence. Together, this makes **b** a very reasonable interpretation of the case as they describe it.

Nonetheless, whether skepticism about precise probabilities is productive may depend on whether it prevents us from moving forward, or instead suggests new avenues for progress. In this case, as the paper aims to show, embracing concerns about the standard, precise picture shows us a potentially very productive way forward, and even suggests a way to expand and develop the knowledge-first paradigm.

At any rate, in case **b**, $P1$ is true not (or at least, not only) for the reason identified by Mueller and Ross. Rather, if $\neg P_T(Rain) = .4$, then $\neg P_K(Rain) = .4$, via factivity.⁶ What we identify as the background truth makes a difference: we will see in Sect. 3.3 that the reason behind $P1$ makes a difference to how we assess $P2$, too.

3.2.3 Belief

Let us next consider what Elvis believes in our scenario. We need to imagine what subjective epistemic state he might be in, and how that $P_B(Rain)$ is best captured and represented for purposes of formal decision-theoretic analysis. Since the point here is precisely *not* to take Savage-style Expected Utility Theory for granted, we must consider a slightly enlarged set of possibilities from the epistemology and decision theory literatures. The relevant broad possibilities here are (**c**) that the best representation of Elvis is with the precise $P_B(Rain) = .4$, and (**d**) that he is best represented as believing something less specific, for example by $P_B(Rain) = .4 \pm \epsilon$, $P_B(Rain) \geq .3$ or $P_B(Rain) \in (.3, .5)$.

There are some scenarios which would uncontroversially be counted as **c** cases. Elvis may have recently seen a weather report which stated that $P_T(Rain) = .4$, and he may have taken this report on board. Whether for this or some other reason, if Elvis himself would assert that $P_B(Rain) = .4$ if asked—and especially if this assertion would hold up under prodding—then Elvis is best represented as having this precise belief or credence. Elvis takes himself to be in a classical situation of risk.

There is also a well-known and important literature arguing that if Elvis' comparative beliefs or preferences conform to certain axioms of rationality, then it is possible for us to represent him as having a precise credence, even if Elvis himself has no introspective access to such numbers (see Fishburn 1986 for an overview).⁷ Especially in cases in which there is a single probability distribution representing Elvis' comparative likelihood judgments, it is natural to think that we should make use of this ability and treat Elvis as having, for example, the posited .4 credence in rain. Similarly, if one construes Elvis' situation as an **a** case with a precise $P_E(Rain) = .4$, then one

⁶ That factivity would not be satisfied is shown by Moss' illustrative examples (2017) and by the formal account of factivity for probabilistic knowledge worked out in Rich (2019). The intuition is that if it is only *true* e.g. that rain *may* fall [i.e. $P_T(Rain) > 0$], but $P_B(Rain) = .4$, then Elvis is wrong: the probability is not .4. This is akin to believing A when only $A \vee B$ holds.

⁷ See Joyce (1998) and subsequent literature regarding an alternative approach to probabilistic representability based on credal accuracy. Considering this approach would affect the details, but not the substance, of the arguments below.

could also imagine that Elvis' belief state perfectly matches this ideal, so that also $P_B(\text{Rain}) = .4$. Then the scenario would be construed as a **c** case.

There are three main reasons for treating the scenario as a **d** case instead. The first is skepticism regarding the normative arguments in favor of **c** cases, the second is skepticism that those arguments reflect psychological reality, and the third is that **d**-case representations may better serve our present purposes.

First, the normative arguments that Elvis should be representable by a probability distribution—and especially their interpretation—are contested. The set of conditions which are necessary for Elvis' probabilistic representability are much more straightforward and less controversial than the enlarged set of conditions (Kraft et al. 1959) which would also guarantee that Elvis' beliefs can be represented by a probability function (see Fishburn 1986 for an overview and Konek 2019 for further discussion). Even when there is a probabilistic representation of Elvis, there will be other, non-probabilistic representations of his beliefs as well; this makes it difficult to explain why we should take the probabilistic representation especially seriously, favoring it over alternatives (Zynda 2000). Finally, and perhaps most importantly, there will often be a set of probability functions which could be used to represent Elvis' beliefs. Yet it is precisely the uniqueness of the proposed credence function which would provide the strongest prima facie reason for a **c**-case over a **d**-case interpretation (see Fishburn 1986; Zynda 2000 for discussion).

Whether the balance of normative considerations requires **c**-case interpretations or instead supports **d**-case agents is a matter of long-running and ongoing debate; see for example Suppes (1974), Elga (2010), Joyce (2010), Bradley and Steele (2014), Schoenfield (2017) and Konek (forthcoming). A review of the literature or a rehearsal of the various arguments and critiques thereof is beyond the scope of this paper, and I engage with this debate only indirectly. What matters for our purposes is just that it is extremely plausible that Elvis' belief is best construed as involving a degree of uncertainty, and many people have defended this position from a variety of perspectives and using diverse formal tools. The counter-proposals include various axiomatic systems supporting the representability of agents via ranges of probabilities. For example, Elvis may conform to rationality postulates which allow him to be represented as having $P_B(\text{Rain}) \in (.3, .5)$. A key claim of this paper is that the knowledge norm of action provides a new consideration in favor of such a representation.

Second, we can separate the questions of what ideal agents would or could believe and what real agents do believe. In asking what Elvis believes, we ask a descriptive question. The empirical literature leaves little doubt that real agents' epistemic states sometimes reflect ambiguity. Best known are Ellsberg's classic experiments (1961), which showed that people do not always attribute precise probabilities to events beyond what their information or evidence warrants, and spurred significant research into ambiguity. The enduring influence of these experiments reflects the assessment that in many situations the theory of ambiguity better reflects agent's actual epistemic states than does earlier theory based on precise probabilities.

Third, it is important to bear in mind that which modeling tools we use is a choice. We may be able to represent Elvis as in either a **c** or a **d** case—especially given such a sparse description of his situation and no way of testing him—and need to decide on the basis of our theoretical purposes. From the perspective of knowledge-

first epistemology, there is reason to choose a formal representation of Elvis' beliefs such that those beliefs will be good candidates for knowledge. A **d**-case representation may allow us to identify relevant items of knowledge that would not show up with a **c**-case modeling approach. For example, we may be able to formally represent Elvis via either a precise $P_B(\text{Rain}) = .4$ or a range $P_B(\text{Rain}) \in (.3, .5)$, but if only the latter would count as knowledge for Elvis, then the latter representation would be the one of interest for us here.

Mueller and Ross treat Elvis as being in a **c** case since they have Savage's Expected Utility Theory in mind, which uses precise credences. Their discussion, however, strongly encourages a **d**-case interpretation, in particular because they discuss the intuitive psychology behind Elvis' assessment of the evidence. They explain that humans like Elvis are not perfectly accurate in this capacity, and that they therefore cannot track the epistemic probabilities perfectly. Elvis, being aware of his status as a human, would probably not be perfectly confident of whatever probability of rain he would try to estimate. From their description of Elvis, it's clear from the start that Mueller and Ross do not take any precise credence in rain to constitute knowledge for him here, and for that matter, the precise credence does not seem like Elvis' belief, either; one is tempted to call it so only from the Expected Utility perspective. In light of the foregoing, this observation suggests interpreting Elvis as considering a range of probabilities as possible, reflecting the margin of error of his calculations. Put differently, Elvis can measure the probability, but not with arbitrary precision.

Under our alternative modeling possibility—in **d** cases— P_1 would be true for a second reason: $\neg P_K(\text{Rain}) \geq .39992$ because $\neg P_B(\text{Rain}) \geq .39992$. Elvis' belief would not unequivocally support buying the umbrella. Mueller and Ross cannot easily dismiss this as a live possibility because they themselves motivate it (albeit using different language) and consider what would follow (2017, Sect. 4). They argue that the proponent of the knowledge norm cannot provide a valid decision principle once we admit that Elvis is uncertain about whether it is likely enough to rain, but nowhere argue that Elvis must have a precise belief. Since the present focus is on characterizing Elvis' mental state, let us set aside their argument about the decision principle until Sect. 3.3.

3.2.4 Epistemic virtue

Beliefs, whether of the **c** or **d** variety, cannot constitute knowledge if they are faulty in some pertinent way; in the traditional analysis, one would label the relevant property "justification." While there is wide variety in the proposed concepts of epistemic goodness, we can again consider two rough possibilities: (**e**) Elvis' belief exhibits no relevant fault, or (**f**) the belief is instead importantly faulty.

According to various accounts, the difference between **e** and **f** cases will have to do with how appropriately Elvis responds to his evidence, how well his beliefs track the truth (Roush 2005), or how reliable his belief-forming process is (Goldman 1979). Generally speaking, only non-faulty beliefs are candidates for knowledge. Importantly, Moss (2017, Sect. 5.7) demonstrates that this basic protocol applies to probabilistic beliefs just as to outright ones.

Mueller and Ross seem to suggest that in borderline probability cases such as this, Elvis could do everything right but still lack knowledge because his belief that the probability warrants buying the umbrella cannot be safe. This concern seems real if we have taken Elvis to be in a **c** case, as Mueller and Ross do. That is, they take Elvis to rationally have a precise belief but not to have the precise knowledge that would support his action. Let's note that if Elvis is additionally in a **b** case, such that his precise belief is false, then he is already disqualified from knowing. Given the discussion in Sect. 3.2.2, he has arguably done something wrong as well; if the epistemic probability plays the role of truth, for example, then Elvis' belief would be more precise than warranted by the evidence.⁸ Then we need to consider whether Elvis may find himself in an **a–c–e** case but still lack knowledge.

There are two possible responses to Mueller and Ross' claim about Elvis doing everything right but failing to have knowledge. The first response would be to deny their premise: if Elvis' belief is not safe, and if unsafe beliefs cannot be knowledge, then Elvis has not done everything right. This is a natural response for the knowledge-firster, given that they endorse a knowledge norm of belief as well as action: in Smithies' terms, for example, "knowledge is the *aim* of belief, which sets a normative standard of correctness and a corresponding normative standard of justification" (Smithies 2012, p. 266). Perhaps safety is not actually necessary (see Bogardus 2014; Broncano-Berrocal 2014 for discussion) or does not apply in probabilistic cases (it's not clear how it should, if probabilistic knowledge is not to be too hard to come by). Then Elvis may know after all, and be in an **e** case. Perhaps instead Elvis really is in an **f** case, and therefore would act inappropriately in acting on his .4 belief. Either way, the claimed dissociation between knowledge and the correct basis for action would not arise.

The second response, really a variant of the first, would be to say that Elvis is in the same kind of situation as a Gettiered agent, in that he intuitively has done as well as we could ask of him, but nonetheless lacks knowledge. In this case, the analysis put forward by Mueller and Ross essentially pinpoints the same issue as Brown's objection to the knowledge norm (Brown 2008a, b). Briefly, Brown points out that a Gettiered agent does not qualify as knowing the proposition in question, but nonetheless appears to reason and act appropriately on its basis. The knowledge-firster typically replies that the agent does not in fact reason and act appropriately, but they have a good excuse; they did as well as one could expect given the circumstances (Williamson 2011; Hawthorne and Stanley 2008). The status of this kind of response is debated (see e.g. the discussion in Gerken 2011; Comesaña 2018; Heil 2019b), but I assume for the sake of argument that the knowledge-firster can adequately address the type of situation raised by Brown. The focus here is on the benefit of bringing the theory of ambiguity into the knowledge-first toolkit.

The same reply would then in principle be available regarding Elvis, if we see no way for him to have done better. Elvis would have an excuse for not believing appropriately. In this case, though, the proponent of the ambiguity analysis will of course point out that Elvis could have acted on a less precise, hence less problematic belief; he would arguably be blameworthy. Either way, we would classify Elvis as being in an **f** case.

⁸ See Rich (2019) for a formal elaboration of the intuition that Elvis believes falsely in this case.

We can then summarize the ways in which $P1$ can be true or false using the traditional analysis of knowledge as true belief which is virtuous in some appropriate way. Again, this traditional analysis serves only a heuristic function, helping us to systematically elaborate the cases in which ' $P_K(Rain) = .4$ ' is true or false. First, $P1$ could be true because it could be false that $P_T(Rain) = .4$, so the factivity condition for knowledge would not be satisfied. The background fact against which we judge Elvis' belief might more reasonably be expressed as ambiguous (a **b** case). Second, $P1$ could be true because it could be false that $P_B(Rain) = .4$; Elvis' belief might also be better represented as ambiguous (a **d** case). Third, $P1$ could be true because Elvis' credence exhibits some defect, e.g. the credence might not track the truth, or it might have been formed by an unreliable process (an **f** case). Finally, if we decide that Elvis indeed has a "true" and virtuous $.4$ credence in rain—including the absence of further problems with his belief requiring that we somehow excuse him—it's less clear how we could still deny him knowledge. In this last (**a–c–e**) case, we would reject $P1$.

3.3 Examining P2: The status of subjective expected utility maximization

The previous section delineated many possibilities regarding what Elvis might believe or know, and why. These different possible cases will have different implications for $P2$, which asserts that we should determine Elvis' appropriate action by modeling him as maximizing expected utility relative to a $.4$ credence in rain. Although the normative force of Expected Utility Theory is commonly assumed for applications in epistemology, this is also a strong assumption and one which is contested for many reasons in many contexts. It is important to bear in mind that formal decision theories are modeling tools rather than absolute truths, and we must always check whether we are using the appropriate tool for the job at hand. Regarding Elvis, I will argue that $P2$ is only well-motivated (from a knowledge-first perspective) when $P1$ is false; if Elvis lacks knowledge of the precise probability of rain, then there are better decision theories to apply to his choice. This means that we have not been given reasons to accept both premises of the argument simultaneously. The conflict between the knowledge norm of action and correct behavior, expressed by the conclusion, disappears accordingly.

The case in which $P1$ is false—the **a–c–e** case—is relatively straightforward. I argued that the knowledge-firster should resist the dissociation between unproblematic true belief and knowledge, one way or the other, and so we should take cases in which Elvis really believes unproblematically to be knowledge cases. If $P_K(Rain) = .4$, then there is good reason to take Elvis' correct action to be the expected-utility maximizing action relative to this knowledge (i.e. buying the umbrella).⁹ The case in which Elvis acts on precise probabilistic knowledge is the kind of case in which Expected Utility Theory functions the best; we could then even apply the more straightforward and less controversial von Neumann-Morgenstern axioms for given probabilities (Von

⁹ Arguments for non-expected utility risk preferences, such as in Buchak (2013), are orthogonal to our interests here. Such arguments have to do with what kinds of preferences are admissible given that the agent perceives the situation as risky. The knowledge norm of action does not bring new reasons to question the form of Elvis' preferences in this case.

Neumann and Morgenstern 1947) rather than the more demanding and contentious Savage axioms (or similar) (Savage 1954).¹⁰ Indeed, the arguably most influential challenge to the theory—namely (Ellsberg 1961)—puts pressure precisely on what happens when such knowledge is absent and uncertainty is greater.

When P_1 is true, it matters why it is true. We have agreed that $\neg P_K(\text{Rain}) = .4$, but the reason for this influences our reasoning about how best to model Elvis' decision problem so as to identify the action he should take. Let's first consider **f** cases. If $P_B(\text{Rain}) = .4$ but Elvis lacks knowledge because this belief is defective, then—given Mueller and Ross' arguments that the expected utility maximizing choice cannot be recaptured via a knowledge-based decision procedure—it won't be appropriate for Elvis to maximize expected utility relative to this belief. Then we should reject P_2 , since applying Expected Utility Theory will not identify Elvis' appropriate action. At least, the proponent of the knowledge norm must say so. Elvis will not act appropriately no matter what decision principle he follows, if we take him to act on the basis of a faulty belief and not on the basis of knowledge. The literature reflects the conviction that fully appropriate action does depend on the quality of the underlying epistemic state.¹¹

If Elvis maximizes expected utility on the basis of an epistemically flawed $P_B(\text{Rain}) = .4$, then, his action is also defective. Appealing to the line of reasoning discussed above, this is true even if Elvis has an excuse for the belief in question. If we follow Mueller and Ross' description closely, however, then it would seem to be clear to Elvis himself that he lacks the corresponding knowledge $P_K(\text{Rain}) = .4$. He would therefore avoidably violate the knowledge norm of action if he simply maximized expected utility relative to this probability. Importantly, Elvis is not stuck acting on a non-knowledge-constituting belief for which we would need to excuse him. According to Mueller and Ross' description, for example, Elvis does seem to know something like $P_K(\text{Rain}) = .4 \pm \epsilon$ for some ϵ . In general, if we weaken Elvis' belief enough, we will find something that does count as knowledge.¹² So, Elvis can rescue himself from inappropriate action simply by acting on the basis of what he

¹⁰ Since the von Neumann-Morgenstern axioms take probabilities to be given, they pertain only to preferences and they only need to deliver a representation of the agent's utility function. The Savage axioms, in contrast, need to deliver a simultaneous representation of *both* the agent's beliefs (via a probability function) and preferences (via a utility function). The Savage axioms, furthermore, need to be strong enough to ensure that the agent acts as if they assigned precise probabilities to every event, i.e. they need to rule out behavior indicative of perceived ambiguous probabilities. Savage therefore requires a number of so-called "structural axioms" imposing particular kinds of complexity on the world in addition to the less demanding "rationality axioms" (Suppes 1974). Many have argued that it is more defensible to take rationality as requiring less of agents, given how the world actually is. The upshot is that von Neumann and Morgenstern only need four simple, intuitive axioms which are easy to endorse on normative grounds, whereas Savage needs seven axioms which more readily invite objections. If we can treat the agent as having precise probabilistic knowledge, then we can use the normatively more forceful set of axioms to infer that they should be (representable as) maximizing expected utility.

¹¹ For example, Mueller (2017) has recently delved into this fairly standard presumption, arguing that decisions based on irrational epistemic states are not only epistemically faulty, but practically irrationally as well. Similarly, Weisberg (2013) defends expected utility maximization as only a necessary condition for rational action.

¹² The weaker the (probabilistic) proposition, the easier it is to know all else being equal; at minimum, Elvis should have $P_K(\text{Rain}) \in [0, 1]$ (see, Rich 2019).

does know—and there are plenty of formal theories available to capture such action. Once we recognize that Expected Utility Theory is not mandatory, $P2$ is no longer very compelling—certainly not from the knowledge-first perspective—in this case. If we're seeking to identify Elvis' knowledge-based decision, and what Elvis knows is not captured by a specific, unique probability function, then we should not take a modeling approach that takes such an epistemic state as an input. But if we reject $P2$, how should we model Elvis' decision problem? This gets us to the **d–e** cases in which Elvis acts on the basis of non-faulty, but non-precise, beliefs.

These final cases are also the most realistic and most important ones: Elvis does not know that the probability of rain warrants buying the umbrella because he is in fact uncertain about this (as most of us would be). We can formalize Elvis' epistemic state while capturing his uncertainty in various ways, for example as Mossean probabilistic knowledge (Moss 2017), as imprecise probabilities (Walley 1994), or using belief (lower probability) and plausibility (upper probability) functions in the Dempster-Shafer tradition (Dempster 1967; Shafer 1976).¹³ Any of these would allow us to formally represent Elvis as taking the probability of rain to be between .3 and .5, between .39 and .41, etc. as the case warrants. In this case, given that one seeks to determine Elvis' rational choice via a decision theory, standard Expected Utility Theory is simply not sufficient. Elvis' epistemic state (whether we consider P_B or P_K) is not among the possible inputs to an expected utility calculation, and so $P2$ should be rejected. This is much easier to see and do with an alternative representation of his epistemic state (and an alternative decision theory) on the table. $P2$ is not false because mathematical decision theory has no role to play here, but because the knowledge-focused articulation of the case makes calculating the expected utility given a .4 credence the wrong option.

An alternative is illustrated in Sect. 4, but in further support of this line of argumentation, it is helpful to consider Mueller and Ross' own discussion of Elvis' uncertainty, and their response to it. Critically, Mueller and Ross acknowledge that Expected Utility Theory may be too demanding in this circumstance and discuss the possibility of determining Elvis' rational choice on the basis of Elvis' actual knowledge. They consider a particular alternative decision principle that could be employed when Elvis is uncertain as to whether $P(\text{Rain})$ is above or below the threshold, in their framing (Mueller and Ross 2017, Sect. 4). According to this principle, Elvis would be permitted to either buy the umbrella or not, since either would be compatible with some $P_T(\text{Rain})$ that he does not rule out. Mueller and Ross reject this decision rule. It does not yield a unique verdict as to what Elvis should do, and worse, they show that Elvis could be turned into a money pump because he would be permitted to continually buy the umbrella at one price and then sell it back at a small loss.

This is a familiar argument pattern, and much could be said about it. Even if we agree to reject the permission-based decision principle, however, it would not do much to support $P2$. Mueller and Ross refer to the choice between expected utility maximization and their permissibility rule as a dilemma, but there are certainly more

¹³ See Wakker (2008), especially the subsection on *multiple priors*, for a useful overview of approaches, oriented towards the more developed economics literature.

than these two options here.¹⁴ The option space is limited in this way only because the debate over the knowledge norm has tacitly neglected most of decision theory. While permissibility rules may have their place, the most attractive alternative to P_2 is not to retreat from formal decision theory and to give up on determinate judgments about what Elvis should do. Expected Utility Theory was a brilliant starting point rather than an end-point; in the intervening decades, a plethora of formal decision theories and game-theoretic principles have been developed for exactly the kind of uncertainty-reflecting epistemic states P_K that we would expect Elvis to have. These alternative theories *will* deliver precise verdicts about what Elvis should do, with no need to consider multiple options to be permissible. In light of these offerings, P_2 is severely undermotivated from the knowledge-first perspective.

The foregoing should suffice to show that the conclusion of Mueller and Ross' argument is not established by their example. Once we consider ways of filling in the details of Elvis' situation, we see that for each characterization, either the knowledge-firster should grant that Elvis really does have knowledge $P_K(Rain) = .4$, or we need/should not model his choice as maximizing expected utility with respect to $P_B(Rain) = .4$. In short, we should take Elvis' pertinent knowledge—whatever it may look like—as the epistemic input to a decision theory for which that knowledge is a suitable input, at which point the mismatch between the knowledge norm and decision theory disappears. Hence it has not been demonstrated that Elvis' appropriate choice is not a knowledge-based choice. In the remainder of the paper, I elaborate on the salient alternative to Mueller and Ross' picture, in which Elvis is modeled as choosing rationally based on his knowledge.

4 Decision under ambiguity

4.1 Alternatives for epistemologists

In light of the role played by the presupposition of Expected Utility Theory in the dialectic, simply demonstrating an intuitive and easy-to-use alternative can go a long way towards supporting the knowledge norm of action. It will show just how low the barrier to employing a formal decision theory taking knowledge as its input actually is. In this way, we can have all the benefits of a well-worked out and understood formal theory without compromising or casting doubt on the requirement that decisions rest only on knowledge.

As already noted, there are many different ways to formally represent Elvis' uncertainty-manifesting epistemic state as we intuitively understand it. Similarly, there is a true wealth of alternative theories of decision-making given that ambiguity is present; these range in complexity from the simple permissibility rule discussed by Mueller and Ross to extremely complex and mathematically sophisticated theories developed in the economics literature. Nothing like an overview or survey will be attempted here; Wakker (2008), Etner et al. (2012), Walley (2000) and Rinard

¹⁴ Heil (2019a) makes this point as well. He suggests using imprecise credences based on a margin of error of $\pm\epsilon$ and then applying a variant of the permissibility rule considered by Mueller and Ross.

(2015) collectively provide a good theory-oriented picture, while Trautmann and van de Kuilen (2015) review and discusses the experimental literature.

This paper proceeds to demonstrate the deployment of one particular theory, the so-called “meta-utility theory”; we have developed variants in Grant et al. (2020a, 2021), but I present the version in Grant et al. (2020b), which is best suited for use in epistemology. The basic idea of meta-utilities is that the agent’s valuation of an ambiguous act (such as buying an umbrella, perhaps) is a function of the worst and best possible expected utility of the act. In contrast to theories in which the agent may also have information or beliefs about how likely particular prior probability distributions are (see e.g. Klibanoff et al. 2005), only these lower and upper bounds affect the agent’s preferences. The agent’s ambiguity preferences essentially capture how the agent values trade-offs between the worst and best possible expected utilities, which we describe and model as “goods.” This valuation must be systematic, i.e. representable via a neat mathematical function, but the theory does not assume that the function has a particular form, for example that the agent likes or dislikes ambiguity, or that the agent’s attitude towards ambiguity does or does not depend on the stakes involved. In this way, the theory encompasses a variety of previous proposals, recognizing that agents are likely heterogeneous in their responses to uncertainty. The theory will be explained in a bit more detail, and applied to Elvis, presently. First, I explain why such a theory is appealing.

The meta-utility approach has many virtues for present purposes. It provides rigorous mathematical foundations—including a representation theorem for preferences via an effectively unique utility function—in contrast to the permissibility-based approaches discussed in Mueller and Ross (2017) and Rinard (2015). In addition to the familiar benefits (highlighted for example by Weisberg (2013) and Feldbacher-Escamilla (2018) in the knowledge-first literature, albeit with reference to Expected Utility Theory), this means that money pump objections (and other problems raised for the various permissibility rules in the philosophy literature) cannot arise.

Compared to many of the theories in the economics literature, however, the meta-utility theory is both relatively simple and very transparent. This is important because complexity can present a serious practical barrier; it is critical that a theory’s assumptions are easily evaluated and defended. In presenting the meta-utility theory, we are explicit about the goal of making models of ambiguity more accessible, and emphasize that the theory is constructed from a small number of simple and familiar tools.¹⁵ Grant et al. (2020b) has only four axioms. The time may be ripe for epistemology to extend its option space and embrace the theory of ambiguity. (And indeed, if the reader is not convinced of the need to consider the effect of ambiguity on Elvis’ umbrella dilemma, then perhaps reflecting on the many decisions we must make on the basis of possible, uncertain climate change consequences will be more motivating.)

In the end, however, the most important goal of this paper is not to convince knowledge-firsters (and other epistemologists) to apply any particular theory of ambi-

¹⁵ In fact, the ambiguity literature is seen as abstruse and unwieldy even among economists. For example, in their survey, Etner et al. (2012, p. 234) write that “[t]he non specialist observer might be overwhelmed by the number of different models aimed at capturing how ambiguity can affect decisions.” They aim to provide a “(more or less) user friendly guide” (2012, p. 235) so that the practical application of models of ambiguity can increase, commensurate with the real-world relevance of ambiguity.

guity, but rather to embrace the presence of ambiguity in agents' epistemic states and to characterize their decisions accordingly. Even further pursuing coarse permissibility rules is compatible with this primary objective. For this reason, no attempt will be made to discredit alternative approaches to decision under ambiguity. It is in fact reasonable to think that different decision principles are appropriate in different contexts, for example because the amount and structure of agents' information varies from case to case. What matters is that significant progress can be made; theories of ambiguity are a viable resource.

In the remainder of this section, I briefly explain the meta-utility theory for ambiguity and show how the knowledge-firster can appeal to it to address Mueller and Ross's (2017) example. Section 4.2 provides an overview of the theory, explains the axioms, and describes the representation theorem. Section 4.3 applies the theory to Elvis, showing how our specific hypotheses about his knowledge and preferences combine to determine his rational action in his umbrella dilemma.

4.2 The meta-utility theory of ambiguous decisions

4.2.1 How it works

The meta-utility theory provides a set of axioms capturing coherent preferences between available actions, given that the agent may perceive some decision-relevant events as ambiguous (Grant et al. 2020b). This means that in contrast to the Savage axioms (Savage 1954), the axioms do not impose the requirement that the agent treats all situations as merely risky. The axioms provide necessary and sufficient conditions for the agent's choices to be representable via a meta-utility function. The agent is thereby represented as an expected utility maximizer for situations the agent perceives as risky, while when the agent perceives ambiguity, their evaluation of actions is fully determined by the worst-case and best-case expected utilities of the action, given the probability bounds they are able to place on the ambiguous events in question.

The representation of the agent then consists of three components: a prior probability function, a (Bernoulli) utility function, and a meta-utility function. I describe each in more detail before giving an overview of the axioms. A formal description of the representation and axioms would require rehearsing the full formal apparatus of that paper, which would be a distraction. I therefore keep the discussion here rather informal. The reader interested in the formal details should consult the original paper (Grant et al. 2020b), while the reader interested in the main ideas and intuitions may just note that the general framework is that of Savage,¹⁶ with expansions for ambiguity as developed by Gul and Pesendorfer (2014).

The agent's prior π , first of all, is the familiar probability distribution over events, where the events form a σ -algebra. Since the agent may perceive some events as ambiguous, however, some decision-relevant events may not be in the domain of the prior. The agent may not be able to directly assign a probability to those events. For the knowledge-firster, the key is to let the prior encode the agent's knowledge, as we will

¹⁶ The new and aptly-named paper *Savage for Dummies and Experts* provides a useful overview and discussion (Abdellaoui and Wakker 2020).

do with Elvis presently. This application fits well with the original interpretation of the prior as reflecting the agent's (lack of) information (Grant et al. 2020b). The agent does not have sufficient information to assign a precise probability to every event, but they can use the information they do have to place lower and upper bounds on the probabilities of the events in question.¹⁷

Next, the agent's representation includes a Bernoulli utility function u assigning personal utilities to real outcomes. Together, π and u determine the agent's expected utilities for risky situations. When the agent's prior does assign a precise probability to all action-relevant events, the agent also evaluates the act according to its expected utility. It is only when the agent cannot do this that Expected Utility Theory does not determine the decision, and the agent's meta-utility function for ambiguity comes into play.¹⁸ Instead, then, the agent's prior and utility function combine to determine the *expected utility bounds* on the agent's available acts. So, while an action to which ambiguity is relevant does not have an expected utility, the agent can attach minimum and maximum expected utilities to it, based on their information.

Finally, the agent's representation includes their meta-utility function $U(\underline{u}, \bar{u})$, a utility function which does assign a utility to every act. The evaluation or ranking of an act is a function of its expected utility bounds, the worst-case and best-case expected utilities as described above. When ambiguity is not relevant, this collapses to the expected utility of the action. In this way—by stipulating that $U(x, x) = x$ —a unique meta-utility function is fixed to represent the agent. The meta-utility function captures how much the agent likes or dislikes ambiguity, including how this may depend on the range of expected utilities on the table. It captures, for example, how much best-case expected utility a cautious agent would be willing to give up in exchange for a given improvement on a given worst-case expected utility, which they would like to avoid.

4.2.2 Axioms and representation

I provide the substance of the axioms, again to avoid a detour through the many definitions needed to state the axioms formally. Since the framework is Savage-style, the axioms pertain to a weak preference relation \succsim and preferences of the form $f \succsim g$, where f and g are arbitrary acts. The acts map states to outcomes; for example, the act of not buying an umbrella yields a particular monetary outcome according to whether Elvis is in a rain or non-rain state of the world.

In Grant et al. (2020b), we present four axioms which are necessary and jointly sufficient for a meta-utility representation. Axioms 1–3 are familiar. Axiom 1 is that \succsim is complete and transitive, while Axiom 2 is that \succsim is monotonic. That is, the agent has a weak preference ordering over all acts, and if one act yields an objectively better

¹⁷ Good (1966) defends a theory along these lines. He addresses potential complaints, for example that we cannot expect the lower and upper probability bounds to be sharp.

¹⁸ Other kinds of challenges to Expected Utility Theory are not considered in this paper, since they have not been an important issue within epistemology, and in particular the parties to debate about the knowledge-first paradigm have not questioned Expected Utility Theory. It is worth noting, however, that nothing about the meta-utility approach described here depends on risky options being evaluated as in Expected Utility Theory. In fact, Grant et al. (2020a) presents a version of the meta-utility theory which allows non-expected utility risk preferences as well as loosening the requirements on the prior probability.

payoff in every state of the world, the agent strictly prefers it. Axiom 3 is a technical axiom, requiring continuity.

The new and critical axiom, Axiom 4, is the key to the result that an act's meta-utility is fully determined by its worst-case and best-case expected utilities. The axiom requires that for any two (potentially ambiguous) acts f and g , "if both the lower and upper bounds for the certainty equivalent of an act f are at least as large as those for another $[g]$, then f must be weakly preferred to $[g]$ " (Grant et al. 2020b, p. 7). In terms of preferences, an act's lower bound is determined by looking at the set of sure outcomes which are weakly preferred to all risky acts dominated by it. Similarly, the act's upper bound is determined by considering the set of sure outcomes weakly dispreferred to the risky acts dominating it. The upshot is that if f has both a better (higher) worst-case and a better best-case expected utility than g , then f has higher meta-utility (overall utility) than g .

To see the key axiom in action, suppose the worst-case for an act f is worth €5 to me, while f 's best-case is worth €100. Then my uncertainty is great, but I can still rank this option as better than an alternative g whose worst-case is worth €4 and whose best-case is worth €95. Whichever aspect of my information I would emphasize or focus on, f appears better than g . What is thereby ruled out is that the agent uses (or acts as if they had) further information about, say, the lower bound of f being more likely and the upper bound of g being more likely in this case, and so forth. As noted previously, this makes a lot of sense when we're imagining agents to lack information. In particular if we want to see the consequences of agents acting on their knowledge, it is reasonable to assume that their knowledge (and hence the determinants of their choices) usually does not go beyond the range or the margin of error within which they can make probabilistic judgments. In cases in which the agent would have such additional knowledge, they would be entitled to use it, and a theory incorporating such information would be called for.

The representation theorem in Grant et al. (2020b) then establishes that the agent's preference relation \succsim conforms to Axioms 1–4 as described above if and only if the agent is representable as a meta-utility maximizer. To represent the agent as a meta-utility maximizer is to represent them as maximizing a meta-utility function $U(\underline{u}, \bar{u})$, relative to some prior π and (Bernoulli) utility function u . As noted, the agent's meta-utility function is unique once we stipulate that $U(x, x) = x$.

4.3 Meta-utility representation of Elvis' knowledge-based decisions

With a suitable decision theory on the table, we can re-evaluate what Elvis should do in a way that takes his knowledge as an input to his choice. The objective is to make a specific judgment about what Elvis should do, given his epistemic state and preferences, just as under Expected Utility Theory. Since the right decision does depend on these facts about Elvis, we need to consider specific assumptions about Elvis' knowledge and preferences, and see what follows from them. Note that this is exactly what Mueller and Ross do in attributing hypothetical beliefs and a utility table to Elvis. The only difference is that we must now consider how Elvis responds to ambiguity, too.

Table 2 Elvis fears for his shoes, gain version

	Rain	No rain
Buy umbrella	30.01	30.01
Do nothing	0	50.01

Let's begin with Elvis' prior, π_E . By hypothesis, the event that it rains is not in the domain of the prior. Nonetheless, Elvis presumably has information allowing him to bound the probability of rain. Let's imagine that he did see a weather forecast online, but about a week ago, and he was not paying very much attention. He does, however, remember thinking to himself: "Hmm, it may rain in Memphis next weekend. I had better bring an umbrella to that show." Unfortunately, he did not remember this when packing his suitcase, but only now as he looks up at a gray sky, then across the street at a convenience store, realizing that this is his only opportunity to buy an umbrella. Elvis has some experience with Memphis weather. Let's suppose he combines his fuzzy memory of the old forecast with his current sense data and memory of similar days, and determines the probability of rain to be at least .35 and at most .45 (we can express this as $\underline{\pi}_E(Rain) = .35$ and $\bar{\pi}_E(Rain) = .45$). Let's suppose that this is the strongest thing that Elvis knows.

Next, we need to provide Elvis with a hypothetical simple utility function u_E . Mueller and Ross (2017) provide a utility table in terms of losses (Table 1 above). Since the meta-utility theory works with positive utilities, I provide an alternative utility matrix for Elvis (see Table 2). For convenience, however, the new decision table retains the property that the two actions are equally good when $P_T(Rain) = .39992$.

We can now determine the minimum and maximum expected utilities (call them \underline{u} and \bar{u}) for each of Elvis' available actions (buying and not buying the umbrella), given u_E and that $\underline{\pi}_E(Rain) = .35$, $\bar{\pi}_E(Rain) = .45$. In our example, letting *Buy* be the action of buying the umbrella and \neg *Buy* the action of doing nothing, these expectations are as follows:

$$u(\textit{Buy}) = \bar{u}(\textit{Buy}) = 30.01 \quad (1)$$

$$u(\neg\textit{Buy}) = 27.51; \quad \bar{u}(\neg\textit{Buy}) = 32.51 \quad (2)$$

These expectations reflect the fact that buying the umbrella could be better or worse than not buying it, depending on the probability of rain. Either case is compatible with Elvis' evidence, so he does not know one action to be better from the expected utility perspective. His knowledge, instead, is represented by these bounds on what is possible.

To model Elvis as a knowledge-based decision-maker, then, we model him as choosing rationally—coherently—on the basis of his preferences regarding his uncertainty as we have represented it. His preferences are encoded in his meta-utility function U_E . It tells us Elvis' relative valuations for upper and lower expected utility. Many such functions are possible for Elvis; the theory does not tell us that Elvis must be maximally cautious or even that he must be averse to ambiguity. Nor does it tell us that his relative concern about the best versus the worst expected utility cannot depend on

the level of expected utility at issue, for example that Elvis cannot be more cautious in more important situations. These are all questions of taste, not of rationality. His preferences need only be coherent, systematic, describable by a function of worst and best.

It is nevertheless useful to consider specific hypotheses about Elvis' preferences. Several such hypotheses are familiar from the literature, especially the classic proposals discussed in Luce and Raiffa's textbook on decision-making (1958). I will next show how Elvis' rational choice is fixed given each of three different meta-utility functions that might reasonably capture his preferences.

First, according to the Maxmin rule originally proposed in Wald (1950), for any act f ,

$$U(f) = U(u(f), \bar{u}(f)) = u(f) \tag{3}$$

and so

$$U_E(Buy) = 30.01; \quad U_E(\neg Buy) = 27.51 \tag{4}$$

so that Elvis prefers to play it safe and buy the umbrella.

Second, according to Hurwicz' theory (1951), Elvis is not necessarily so conservative; his preferences are defined with the help of an additional parameter α reflecting the relative (fixed) weight he places on the lower expected utility. According to this theory,

$$U(f) = U(u(f), \bar{u}(f)) = \alpha u(f) + (1 - \alpha)\bar{u}(f) \tag{5}$$

and so for example if Elvis is neutral regarding ambiguity and has $\alpha = .5$,

$$U_E(Buy) = 30.01; \quad U_E(\neg Buy) = .5(27.51) + .5(32.51) = 30.01 \tag{6}$$

meaning that the two actions are equally good and Elvis is indifferent between buying and not buying. If Elvis were characterized by a higher (respectively, lower) α then he would prefer to buy (not buy) the umbrella.

Third, according to the geometric α -maxmin theory that Binmore defends in his book (2008, Sect. 9.2), Elvis is similarly characterized by an α parameter, but his utility is defined by

$$U(f) = U(u(f), \bar{u}(f)) = u(f)^\alpha \bar{u}(f)^{1-\alpha} \tag{7}$$

and so here if Elvis has $\alpha = .5$,

$$U_E(Buy) = 30.01; \quad U_E(\neg Buy) = \sqrt{27.51} \sqrt{32.51} = 29.91 \tag{8}$$

so in this case Elvis prefers to buy the umbrella (by a small margin).¹⁹ The advantage of Binmore's proposal is that it relaxes the artificial and perhaps dubious assumption that an agent's relative concern for lower and upper expected utilities is a fixed constant. A

¹⁹ These three proposals are all special cases of the meta-utility theory, but they are not the only possibilities for systematic decision principles based on lower and upper expected utilities. We also present a more general functional form, with parameters that can be set to capture each of these three proposals as special cases (Grant et al. 2020b). Specifically, the above three proposals are members of the class of *constant elasticity of substitution* meta-utility functions, a well-behaved set of functions commonly used in

Maxmin agent would prefer an act with a sure expected utility equivalent to a million dollars to an act with unknown expected utility ranging from slightly less than a million to ten trillion dollars. A Hurwicz agent is less extreme, but may still treat the lower expected utility as three times more important than the best case when such magnitudes are in play. In contrast, the geometric α -maxmin agent displays the intuitive behavior of focusing less and less on the worst case as that worst case becomes less bad. When the lower expected utility is very low, the agent would be willing to give up a lot of upper expected utility in exchange for even a small improvement in the lower expected utility. As the lower expected utility improves, the agent gives it less and less relative weight. When the ensured minimum expected utility is high enough, the agent becomes basically ambiguity neutral. This utility function thus avoids the intuitive counterexamples thanks to which the first two, simpler proposals have been viewed with skepticism. At the same time, the agent's ambiguity preferences are systematic, represented by a straightforward and well-behaved utility function.

Absent compelling reasons to insist that rational agents must be cautious or optimistic to some particular degree or that they must react to ambiguity in the same way whether the stakes are enormous or trivial, it is best not to insist that Elvis be modeled using a particular one of these theories or with a particular α value a priori. How Elvis responds to ambiguity is a matter of his personal preferences, just as is how he responds to risky situations (where the preferences underlie the utility function we have been working with). The meta-utility theory allows us to impose significant coherence on Elvis' choices without imposing a particular functional form.

Recalling the money pump argument that Mueller and Ross (2017) use to dismiss the ambiguity-based solution to the problem they describe, we see that a full mathematical theory of ambiguity allows Elvis to be both reasonably uncertain about the events relevant to his choice and to act decisively as a result of his attitudes toward that uncertainty. The permissibility rule invites the money pump objection because for a range of prices, the rule permits but does not require Elvis to buy the umbrella. Therefore, Elvis would be permitted to buy the umbrella at one price, and then sell it back at a slightly lower price—perhaps repeatedly—resulting in a loss. If we represent Elvis as choosing on the basis of a specific meta-utility function, then there is a stable fact as to whether he should buy the umbrella at each possible price. Suppose though that Elvis is a Hurwicz- α -0.5 agent, and so he is indifferent about buying the umbrella at the price which his utility function reflects. Now if he does buy it, lowering the price will not prompt him to sell it at a loss; he may wish he had waited to buy, but at the cheaper price he strictly prefers to have it. Since Elvis has complete and transitive preferences, a clever umbrella tout cannot exploit his weather uncertainty for profit.

In summary, this section has presented an alternative way of modeling knowledge-based decisions which allows us to acknowledge agents' uncertainty about the true probabilities. The example of the meta-utility theory for ambiguity should make clear

economics. Suppressing the act f , such meta-utility functions have the general form:

$$U(u, \bar{u}) = [\alpha u^\rho + (1 - \alpha)\bar{u}^\rho]^{1/\rho} \quad (9)$$

The paper includes graphical representations of the special cases considered here, which is helpful especially in assessing Binmore's proposal.

that to reject Expected Utility Theory does not mean to give up on having a well-founded and mathematically precise theory which can yield specific verdicts regarding what agents should choose, given their epistemic state and personal preferences. While I have described one example, the alternative modeling options are numerous. When the goal is to model knowledge-based decisions, furthermore, theories which better accommodate the kind of uncertainty that agents typically face in real life are by far the more natural option.

5 Conclusion

This paper began by describing the challenge of reconciling the knowledge norm of action with the fact that agents must regularly act in the face of significant uncertainty. While many authors have recognized that knowledge is compatible with some uncertainty, a fully satisfying account of the knowledge norm must include a story about how it relates to existing accounts of rational choice under uncertainty. In particular, Savage's Subjective Expected Utility Theory has been taken to determine the correct actions for uncertain agents (Schiffer 2007), and Mueller and Ross (2017) show that (despite many attempts by defenders of the knowledge norm) agents' expected utility maximizing choices will not always be supported by their knowledge.

The main goal of this paper was to defend the knowledge norm of action from the Bayesian challenge, and from Mueller and Ross' argument in particular, by showing that the knowledge norm is in fact compatible with formal decision theory. Importantly, two possible conclusions could be drawn from Mueller and Ross' argument: either the knowledge norm is incorrect, or we should not apply Expected Utility Theory as the authors do. The paper defends the latter, arguing that the proponent of the knowledge norm can and should determine an agent's appropriate action in any given case by applying a decision theory such that the agent's correct choice according to the theory would count as a knowledge-based choice.

The paper was structured around Mueller and Ross' example of Elvis deciding whether to buy an umbrella. The details of his decision problem could be filled in in many different ways, but the paper organized these possibilities into two categories. In the first category are the cases in which Elvis has a true, unproblematic belief best represented via a precise probability of rain. In such cases, it was argued that we should count this belief as knowledge; i.e., we should accept that $P_K(\text{Rain}) \geq .39992$. Hence, Elvis' knowledge supports his expected utility maximizing choice. In the second category are the cases in which, for one reason or another, Elvis does not have this particular item of knowledge. In such cases, it was argued that modeling the problem as an expected utility maximization problem is not the best way to identify Elvis' appropriate action. Importantly, the proponent of the knowledge norm of action should first identify the most specific item of probabilistic knowledge regarding rain for Elvis; Elvis' relevant probabilistic knowledge will typically reflect ambiguity in this case, for example that $P_K(\text{Rain}) \geq .3$ and $\leq .5$. Then, a decision theory should be chosen such that Elvis' correct action according to that theory is a function of Elvis' probabilistic knowledge and his personal preferences. For both kinds of cases,

the knowledge-based choice coincides with the choice identified by a well-worked out formal decision theory.

All of this shows that the knowledge norm is not as easy to reject as it may have appeared. Expected Utility Theory is not necessarily a good basis for rejecting the knowledge norm because it cannot be used as a fixed point and applied regardless of what the agent knows and why. What we think about Elvis will ultimately depend on our positions about many other things, including our assessment of various Bayesian arguments, what theory of knowledge is correct, how that theory applies to uncertainty, and so forth. Different (complete) packages of background assumptions will support different detailed analyses of Elvis' epistemic state and his appropriate action. Here, I hope to have shown that there are compelling packages which would make an agent's knowledge a direct input to an appealing formal decision theory. Therefore, we need not accept conflict of the sort Mueller and Ross describe.

An important general take-away is that the proponent of knowledge norms can do something much stronger than to merely cast doubt on the arguments of belief-firsters by pointing to disagreement and alternative modeling possibilities such as ambiguous epistemic states. The knowledge-firster can instead argue from a philosophy of science perspective for the re-evaluation in their favor of relevant long-standing disagreements. The argument for this broader lesson draws on Williamson's picture of knowledge-first as a paradigm which stacks up well with respect to standard theoretic virtues. As noted in the introduction, taking 'knowledge' as a primitive enables a simple and highly coherent picture according to which a plausible natural kind (knowledge) serves as the common normative standard of belief, assertion, and action. The paradigm enables the explanation of a much broader range of empirical data than could other equally simple accounts, while avoiding "overfitting" or the addition of "epicycles" to accommodate contrived or unusual cases, as Williamson might put it.

In light of this, it would be reasonable and justified for the knowledge-firster to see their paradigm as being on the offensive, rather than the defensive, regarding the objection from Expected Utility Theory and others like it. This means taking the primacy of knowledge as a reason for adopting additional friendly (or at least compatible) positions. There has long been disagreement regarding the best representation of uncertainty, the rationality (or compulsion) to act as if one had precise credences regardless of one's information quality, and the status of particular proposals for decision-making given uncertainty (see especially Al-Najjar and Weinstein 2009; Joyce 2010 and the many responses to these arguments for a good overview of the ambiguity debate). Instead of merely pointing out that there are positions in these debates which fit better with the proposed knowledge norms, the knowledge-firster could use coherence with their paradigm as a reason in favor of those positions. In the case of Elvis, if the Expected Utility account of his decision is incompatible with a knowledge-based account, this provides a strong reason to represent Elvis' choice problem in a different way.

We can go further by considering an additional theoretical virtue served by embracing the theory of decision under ambiguity. Specifically, if the knowledge-first paradigm can incorporate or even necessitate a theory of decision under ambiguity, then it achieves external validity. That is, the epistemological paradigm coheres nicely with prominent decision theories and their application in the neighboring field

of economics. There, decision under ambiguity is a flourishing field. While some still question the rationality of ambiguity sensitivity due to puzzles that it poses (see again Al-Najjar and Weinstein 2009 and responses), theories of ambiguity are agreed to be necessary for the tasks of explanation and prediction, both in the laboratory and in the real world (recall again the impact of Ellsberg 1961 and see Grant et al. (2020b, Sect. 4) for discussion of an economically important example). In brief, a neighboring discipline finds theories for ambiguity to be indispensable, and it speaks in favor of the knowledge-first paradigm that it suggests the same conclusion. It could help us to rebuild the coherence between neighboring and overlapping research areas (namely, epistemology in philosophy and decision theory in economics) which have slowly drifted apart over the past half-century.

In summary, the knowledge-firster should turn the tables on Mueller and Ross' argument, arguing that if there is a conflict between Expected Utility Theory and the knowledge norm of action, then it is Expected Utility Theory which is the problem. Like all models, the expected utility model has an appropriate context of application, and when an alternative better matches the problem at hand, we should use it. It is a shame that so much further development in decision theory has been ignored in epistemology; much does not serve our purposes, but some does, as I have shown. Starting from the perspective of a well-supported epistemological commitment like the knowledge norm is an excellent way to determine which of the countless available modeling tools are appropriate for answering our questions about rational action.

I close by pointing to some further lines of inquiry which would be useful for the knowledge-firster to pursue, especially if they seek to apply particular decision theories to identify correct knowledge-based choices as I endorse. First, more will have to be said about when an agent counts as having a given item of probabilistic knowledge. According to the knowledge-first approach, knowledge is basic; we recognize it when we see it. When it comes to probabilistic knowledge, however, we do not seem to have such clear intuitions about what agents know. Moss (2013, 2017) has gotten the ball rolling by defending probabilistic knowledge and giving many examples to show how it is supposed to work, but intuitions and understanding will need to be built up. This may involve, for example, working to translate key concepts such as safety, sensitivity, and reliability into the probabilistic knowledge setting.

Secondly, the question of how knowledge is related to representation theorems for decisions and belief states may become important for determining when an agent has probabilistic knowledge. Familiar approaches to determining what agents believe include assessing the betting odds considered to be fair or attributing to an agent credences (and a utility function) such that the agent's observed choices are utility maximizing. Is the agent's knowledge some generalization or weakening of these credences? Under what conditions can the same representation theorems be interpreted as directly identifying agents' knowledge?

Finally, and relatedly, there will be a lot of work to be done in determining which of the countless decision theories are best interpreted as producing knowledge-based decisions, and under which circumstances. Section 4 of this paper described the application of one decision theory in some detail; this, however, was mainly for the sake of having a concrete example of an alternative decision theory and to show that such a theory can be both as simple to apply and as determinate in its verdicts as the more

familiar Expected Utility Theory. Even among theories incorporating ambiguity, some will be more reasonably interpreted as knowledge-based than others; this might have to do with where they lie on the spectrum from normatively- to descriptively-motivated, or the extent to which the agent is incorporating “objective” information. Furthermore, it is important to note that this paper does not argue that a general knowledge-based decision theory must be a theory for ambiguity. Instead, the key point is that whatever epistemic state the decision theory uses should count as knowledge for the agent, if indeed knowledge is the norm of action. This leaves a very wide range of options open.

In short, insofar as formal tools are going to play an important role in epistemology, their interpretation and suitability needs to be evaluated carefully. A formal apparatus cannot on its own be used to tell us whether a substantial epistemic theory is correct or not. Broader theory also needs to be used to judge the suitability of the formal tools. For knowledge-first epistemology, I have argued that this means reconceiving itself as having a more active role to play where decision theory is concerned than has so far been the case.

Acknowledgements The author gratefully acknowledges previous funding support through the German Research Foundation Emmy Noether Project “Knowledge and Decisions,” Project Number 315078566; this paper is the product of many discussions as part of that project. She especially thanks Moritz Schulz, Roman Heil, Andy Mueller, Alexander Dinges, and Jacques Vollet for helpful comments. Cooperation with Simon Grant and Jack Stecher has also obviously greatly influenced the present paper and the author’s thinking. An audience at the University of Bayreuth provided useful early feedback. Finally, two anonymous reviewers for Synthese provided valuable suggestions that improved the manuscript greatly.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The author declares that she has no conflict of interest.

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

References

- Abdellaoui, M., & Wakker, P. P. (2020). Savage for dummies and experts. *Journal of Economic Theory*, 186, 104991.
- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: A critical assessment. *Economics & Philosophy*, 25(3), 249–284.
- Bertrand, J. (1888). Calcul des probabilités: Seconde note sur la probabilité du tir à la cible. *Comptes Rendus Hebdomadaires des Séances de l’Académie des Sciences*, 106, 387–392.
- Binmore, K. (2008). *Rational decisions*. Princeton, NJ: Princeton University Press.
- Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research*, 88(2), 289–313.

- Bradley, S., & Steele, K. (2014). Should subjective probabilities be sharp? *Episteme*, 11(3), 277–289.
- Broncano-Berrocal, F. (2014). Is safety in danger? *Philosophia*, 42(1), 63–81.
- Brown, J. (2008a). Knowledge and practical reason. *Philosophy Compass*, 3(6), 1135–1152.
- Brown, J. (2008b). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Buchak, L. (2013). *Risk and rationality*. Oxford: Oxford University Press.
- Chew, S. H., & Sagi, J. S. (2008). Small worlds: Modeling attitudes toward sources of uncertainty. *Journal of Economic Theory*, 139(1), 1–24.
- Comesaña, J. (2020). A plea for falsehoods. *Philosophy and Phenomenological Research*, 100(2), 247–276.
- Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics*, 38(2), 325–339.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosopher's Imprint*, 10(5), 1–11.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4), pp 643–669.
- Etner, J., Jeleva, M., & Tallon, J. M. (2012). Decision theory under ambiguity. *Journal of Economic Surveys*, 26(2), 234–270.
- Feldbacher-Escamilla, C. J. (2018). Knowledge first and rational action. *Teorema*, 37(2), 31–54.
- Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science*, 1(3), 335–358.
- Friedman, M. (1991). The re-evaluation of logical positivism. *The Journal of Philosophy*, 88(10), 505–519.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178, 529–547.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: Springer.
- Good, I. (1966). Symposium on current views of subjective probability: Subjective probability as the measure of a non-measurable set. *Studies in Logic and the Foundations of Mathematics*, 44(C), 319–329.
- Grant, S., Rich, P., & Stecher, J. D. (2020a). Bayes and Hurwicz without Bernoulli. *Journal of Economic Theory*. <https://doi.org/10.1016/j.jet.2020.105027>.
- Grant, S., Rich, P., & Stecher, J. D. (2020b). Worst- and best-case expected utility and ordinal metautility. Available at SSRN 3369078.
- Grant, S., Rich, P., & Stecher, J. D. (2021). Objective and subjective rationality and decisions with the best and worst case in mind. *Theory and Decision*, 90, 309–320.
- Greenough, P., & Pritchard, D. (Eds.). (2009). *Williamson on knowledge*. Oxford: Oxford University Press.
- Gul, F., & Pesendorfer, W. (2014). Expected uncertainty utility. *Econometrica*, 82(1), 1–39.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571–590.
- Heil, R. (2019a). Knowledge, action and partial belief, ms.
- Heil, R. (2019b). Knowledge-based decision theory, reasonable false beliefs and guidance, ms.
- Hurwicz, L. (1951). Optimality criteria for decision making under ignorance. *Cowles Commission Discussion Paper, Statistics*, 370.
- Jackson, A. (2012). Two ways to put knowledge first. *Australasian Journal of Philosophy*, 90(2), 353–369.
- Jaynes, E. T. (1968). Prior probabilities. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(3), 227–241.
- Jeffrey, R. C. (1965). *The logic of decision*. New York, NY: McGraw-Hill.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24, 281–323.
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849–1892.
- Konek, J. (2019). Comparative probabilities. In: J. Weisberg & R. Pettigrew (Eds.), *The open handbook of formal epistemology* (pp. 267–348). Philpapers. <https://philpapers.org/rec/PETTOH-2>.
- Konek, J. (forthcoming). Epistemic conservativity and imprecise credence. *Philosophy and Phenomenological Research*.
- Kraft, C. H., Pratt, J. W., & Seidenberg, A. (1959). Intuitive probability on finite sets. *The Annals of Mathematical Statistics*, 30(2), 408–419.
- Luce, R. D., & Raiffa, H. (1958). *Games and decisions: Introduction and critical survey*. New York, NY: Wiley.
- McGlynn, A. (2014). *Knowledge first?*. Hampshire: Palgrave Macmillan.
- Moss, S. (2013). Epistemology formalized. *Philosophical Review*, 122(1), 1–43.
- Moss, S. (2017). *Probabilistic knowledge*. Oxford: Oxford University Press.

- Mousavi, S., & Gigerenzer, G. (2017). Heuristics are tools for uncertainty. *Homo Oeconomicus*, *34*(4), 361–379.
- Mueller, A. (2017). How does epistemic rationality constrain practical rationality? *Analytic Philosophy*, *58*(2), 1–17.
- Mueller, A., & Ross, J. (2017). Knowledge dethroned. *Analytic Philosophy*, *58*(4), 283–296.
- Rich, P. (2019). The logic of probabilistic knowledge. *Philosophical Studies*, *177*, 1703–1725.
- Rinard, S. (2015). A decision theory for imprecise credences. *Philosopher's Imprint*, *15*(7), 1–16.
- Roush, S. (2005). Tracking truth: Knowledge, evidence, and science. Oxford University Press on Demand.
- Savage, L. J. (1954). *The foundations of statistics*. New York, NY: Dover Publications Inc.
- Schiffer, S. (2007). Interest-relative invariantism. *Philosophy and Phenomenological Research*, *75*(1), 188–195.
- Schoenfeld, M. (2017). The accuracy and rationality of imprecise credences. *Noûs*, *51*(4), 667–685.
- Sen, A. K. (1977). Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy & Public Affairs*, pp. 317–344.
- Shackel, N. (2007). Bertrand's paradox and the principle of indifference. *Philosophy of Science*, *74*(2), 150–175.
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press.
- Smithies, D. (2012). The normative role of knowledge. *Noûs*, *46*(2), 265–288.
- Sosa, E. (1999). How must knowledge be modally related to what is known? *Philosophical Topics*, *26*(1/2), 373–384.
- Suppes, P. (1974). The measurement of belief. *The Journal of the Royal Statistical Society B*, *36*(2), 160–190.
- Trautmann, S., & van de Kuilen, G. (2015). Ambiguity attitudes. In G. Wu & G. Keren (Eds.), *The Wiley Blackwell handbook of judgment and decision making* (pp. 89–116). Wiley.
- Tversky, A. (1975). A critique of expected utility theory: Descriptive and normative considerations. *Erkenntnis* pp. 163–173.
- Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.
- Wakker, P. (2008). *New Palgrave dictionary of economics*, 2nd edn. Palgrave Macmillan, Basingstoke, chap Uncertainty, pp. 428–439.
- Wald, A. (1950). *Statistical decision functions*. New York, NY: Wiley.
- Walley, P. (1994). *Statistical reasoning with imprecise probabilities*. New York, NY: Chapman and Hall.
- Walley, P. (2000). Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, *24*(2–3), 125–148.
- Weatherson, B. (2012). Knowledge, bets and interests. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 75–103). Oxford: Oxford University Press.
- Weisberg, J. (2013). Knowledge in action. *Philosopher's Imprint*, *13*(22), 1–23.
- Williamson, J. (2010). *In defence of objective Bayesianism*. Oxford: Oxford University Press, Oxford.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2011). Knowledge first epistemology. In S. Bernecker & D. Pritchard (Eds.), *The Routledge companion to epistemology, Routledge philosophy companions* (pp. 208–219). London: Routledge.
- Williamson, T. (2017). Acting on knowledge. In J. Carter, E. Gordon, & B. Jarvis (Eds.), *Knowledge first* (pp. 163–181). Oxford: Oxford University Press.
- Zynda, L. (2000). Representation theorems and realism about degrees of belief. *Philosophy of Science*, *67*(1), 45–69.