



The knowledge norm of apt practical reasoning

Andy Mueller¹

Received: 30 March 2020 / Accepted: 12 January 2021 / Published online: 27 January 2021
© The Author(s) 2021

Abstract

I will argue for a novel variant of the knowledge norm for practical reasoning. In Sect. 2, I will look at current variations of a knowledge norm for practical reasoning and I will provide reasons to doubt these proposals. In Sects. 3 and 4, I develop my own proposal according to which knowledge is the norm of apt practical reasoning. Section 5 considers objections. Finally, Sect. 6 concerns the normativity of my proposed knowledge norm and its significance.

Keywords Knowledge norm · Epistemic norms · Practical reasoning · Reasons · Epistemic norms for practical reasoning · Virtue epistemology

1 Introduction and overview

I will argue for a novel variant of the knowledge norm for practical reasoning. The suggested knowledge norm applies to a certain kind of decision making that involves the use of reasons for action. For example, your end is to have an enjoyable Friday night out. You now have to decide to which one of two bars in town to go to pursue this end. The first bar serves your favorite beer, which is a reason to go there. The second bar also serves your favorite beer, which is a reason to go there. But there is also live music at the second bar, which you enjoy, and which is a further reason to go to the second bar. Given that the reasons to go to the second bar outweigh the reason to go to the first bar, you decide to go to the second bar.

Such ordinary decision-making based on considering reasons should be familiar to all of us and is also well entrenched in theories. In the following, the term ‘practical reasoning’ will be used to refer to the kind of decision-making based on reasons just outlined. However, the reasons-based approach is not the only available approach to decision making. According to standard rational choice theory, practical rationality is a matter of maximizing expected utility. Reasons play no role in

✉ Andy Mueller
andreas.mueller.uniffm@gmail.com

¹ Department of Philosophy, Goethe University, Frankfurt, Germany

standard rational choice theory; it relies only on credences and utilities to calculate expected utilities.

I will not discuss whether the reasons-based approach or the rational choice approach is correct. Much like Weisberg (2013), I believe that in our decision-making toolkit, there is room for both approaches and that different tools can come in handy for different decision problems. Having said that, my focus here will be on practical reasoning and which, if any, role knowledge plays for practical reasoning.

Here is how I will proceed. In Sect. 2, I will look at current variations of a knowledge norm for practical reasoning and I will provide reasons to doubt these proposals. In Sect. 3 and 4, I develop my own proposal according to which knowledge is the norm of apt practical reasoning. Section 5 considers objections. Finally, Sect. 6 concerns the normativity of my proposed knowledge norm and its significance.

2 Current variations of the knowledge norm of practical reasoning

Williamson (2005), Hawthorne and Stanley (2008), Jackson (2012) and Mehta (2016) have argued that there is an important normative connection between knowledge and practical reasoning. I will quickly introduce these different proposals for a knowledge norm of practical reasoning and give reasons to doubt them.

Hawthorne and Stanley argue for the following principle:

The Reason-Knowledge Principle (RKP): Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p.¹

By 'appropriate', Hawthorne and Stanley mean 'rationally permissible', so the RKP specifies an epistemic condition on when it is rationally permissible to treat p as a reason for action. When one satisfies the RKP, that is if one knows that p, one's reasoning cannot be criticized because of having a lacking epistemic relation to p that would make treating p as a reason impermissible. The amendment "where one's choice is p-dependent" is supposed to ensure the practical relevance of the proposition. I will omit a discussion of this amendment here, as all cases I will use will involve practically relevant propositions. Given the centrality of reasons to RKP, it is obvious that it fits the reason-based approach to decision-making.

We can familiarize ourselves with the RKP by starting with what I call the good case. Jason has a strong desire to have a glass of milk. He remembers that he bought milk last night and placed it in the fridge. Since his memory is reliable and the milk is still in the fridge, Jason comes to know that there is milk in the fridge. Treating this proposition as a reason to go to the fridge, he reasons further to the intention to go to the fridge to satisfy his desire for milk. Epistemically speaking, this piece of practical reasoning is sound because it is rationally permissible to treat the

¹ Hawthorne and Stanley (2008: p. 578).

proposition as a reason since it is known. However, there are three bad cases that suggest that knowledge is not necessary for rationally treating a proposition as a reason.

Case F: Jason remembers that he bought milk last night and placed it in the fridge. But unbeknownst to Jason, a burglar stole the milk while he slept. Nonetheless, Jason has a justified belief that there is milk in the fridge, but he lacks knowledge because there no longer is milk in the fridge. Yet, it is rationally permissible to treat this false justified belief as a reason to go to the fridge.²

Case G: A similar scenario as in Case F obtains. Unbeknownst to Jason, the milk was stolen. But meanwhile, his roommate John, who usually only buys soy milk, bought some regular milk and placed it in the fridge. When Jason remembers that he bought milk last night he ends up with a justified true belief. But since he is in a Gettier scenario, he lacks knowledge. Yet it is rationally permissible for him to treat his belief that there is milk in the fridge as a reason to go to the fridge.³

Case P: Jason is at work and plans to have pancakes after work. For this, he needs milk. He remembers buying some last night, but sometimes his other roommate, Tim, takes some without asking. He tries to call Tim, but cannot reach him. Therefore, Jason is unsure about whether there is enough milk to make pancakes in the fridge and he forms only a partial belief that there is.⁴ But nothing depends on having pancakes after work. Jason would enjoy having some, but he would not be utterly miserable if his desire for pancakes were frustrated. So Jason treats his partial belief that there is enough milk in the fridge as a reason not to go to the store. And again, despite a lack of knowledge, as knowledge requires full belief, this is rationally permissible.⁵

Hawthorne and Stanley have several strategies to deal with such counterexamples. In regard to Case F and Case G, they hold that Jason is excused for his failure to live up to the RKP, but that he nonetheless broke the norm. In regard to Case P, they hold that Jason has knowledge of probabilities in virtue of which he satisfies RKP, for example that he knows that it is likely (enough) that there is enough milk in the fridge. But these strategies to deflect the counterexamples against the necessity direction of the RKP have received criticism. The excuse-maneuver is rejected in Brown (2008: p. 173), Fantl and McGrath (2009: p. 125), Gerken (2011), and Locke (2015). The knowledge-of-probabilities maneuver is rejected in Mueller and Ross (2017). I will not engage with further epicycles in the debate about the RKP. I believe that the given cases are sufficient to cast doubt on the necessity direction of RKP.⁶ Perhaps, this version of a knowledge norm for practical reasoning can be defended. But given the state of the debate, it seems doubtful and thus it seems

² Such cases can be found in Fantl and McGrath (2009) and Gerken (2011).

³ Such cases can be found in Brown (2008), Neta (2009) and Smithies (2012).

⁴ A partial belief is to be distinguished from a full, or outright belief. Commonly, partial beliefs are also referred to as 'credences'. I use the term partial belief to contrast it clearly from full beliefs.

⁵ Such cases can be found in Jackson (2012) or Smithies (2012).

⁶ I also set aside counterexamples to the sufficiency direction of the RKP due to Brown (2008).

doubtful that there is an important normative connection between knowledge and practical reasoning.

I want to quickly look at two alternative proposals to the RKP, which are not explicitly formulated in terms of reasons, yet still can be said to uphold the ambition of outlining of an important normative connection between knowledge and practical reasoning. Williamson (2005: p. 227) argues that knowledge entitles one to rely on a proposition in practical reasoning. Suppose I have a false justified belief about the right return date for a book from the library. I might be excused for relying on this false belief in my practical reasoning about when to return the book, but I am not entitled to rely on it. Unfortunately, Williamson does not explain how the relevant entitlement differs from rational permissibility and why no lesser state than knowledge suffices for it. For example, why would a justified true belief about the return date not suffice for the entitlement? Why does the entitlement require an anti-Gettier condition? Without such clarification, Williamson's proposal runs into problems with scenarios such as Case G.

Both, Jackson (2012: p. 364) and Mehta (2016: p. 682) hold that yet another evaluative dimension, epistemic success of practical reasoning, is subject to a knowledge norm. However, neither Jackson nor Mehta explains what they mean by epistemically successful reasoning. Both give examples of successful actions, for example picking up a visitor from the airport and driving them to the city. If one knows that X is the visitor to be picked up, one's action of driving the visitor to the city will be successful. It would not be successful if one had a false justified belief that X is the visitor. But they do not discuss that knowledge might not be linked to either epistemic success or successful action. As far as successful action is concerned, one will be equally successful in driving the visitor to the city if one merely has a true belief that X is the visitor. So why not just adopt a truth norm for practical reasoning that suggests that only true propositions are to be treated as practical reasons? Why demand knowledge on top of truth? Neither Jackson nor Mehta demonstrate that there is a strong normative link between knowledge and epistemically successful reasoning, which could not also be accounted for by a truth norm.

These objections do not conclusively establish that these suggested knowledge norms are false. However, they are sufficient for raising reasonable doubt about them. My aim here is to argue that there is indeed an important normative connection between knowledge and practical reasoning, a version of the knowledge norm, that is immune to the doubts raised in this section. I think this proposal is of interest in its own right, even if some of the proposals mentioned in this section would turn out to be defensible after all.

Finally, I want to further situate the proposal I am about to develop within the current literature. My proposal draws on ideas from virtue epistemology and knowledge-first epistemology. There is currently a trend in epistemology to combine virtue epistemology and knowledge-first epistemology, see Miracchi (2015), Kelp (2016), Silva (2017). In some sense, my views are in the vein of these authors, while in others, they are not. My project differs from them insofar as it is not a piece of virtue epistemology. I merely use the performance normativity framework used by many virtue epistemologists and apply it to practical reasoning. In this respect, it is similar to Mantel (2013), though she is concerned with *acting* for a reason, not

practical reasoning. The label ‘knowledge-first’ is ambiguous. Insofar as it used for defenses of knowledge-norms, my project fits the label. But it doesn’t fit given other uses of ‘knowledge-first’, for example to express the idea that knowledge is not analyzable or the project to account for epistemic justification in terms of knowledge. Simion (2019) engages in this latter project. While she thus pursues a different knowledge-first project than I do, one that she herself contrasts with virtue-based approaches, we do share the use of the distinction between evaluative and prescriptive norms. Finally, Schellenberg (2018) distinguishes her capacitism about perception from knowledge-first and virtue epistemology. Yet, they also bear similarities and perception certainly not only informs action, but also practical reasoning. My proposal does bear at least a family resemblance to views mentioned in this paragraph. In light of these similarities, it is perhaps not surprising that I do not find anything in the mentioned body of work that undermines my proposal. However, having explained how my project differs from those mentioned, I will set aside these related views.

3 The dimensions of practical reasoning

In the debate about epistemic norms for practical reasoning, usually curiously little is said about practical reasoning. I think that this is a real oversight. Looking closely at practical reasoning reveals a new way to argue for a claim that can rightly be called the knowledge norm of practical reasoning.

I propose to think of practical reasoning as an intellectual performance. It is a performance of determining which course of action to take. This is a performance in the same sense as one’s heart pumping blood is a performance, or an archer’s shot at a target. Practical reasoning is a large performance that can be broken down into several smaller parts. Treating something as a reason is part of the large performance of practical reasoning. The whole performance of practical reasoning ends with drawing a conclusion based on those reasons, which is an intention to act.⁷ I assume that the various parts of practical reasoning affect how we evaluate the performance of practical reasoning as a whole. If a smaller part of practical reasoning falls short in some regard, so does the whole performance. My concern will be which epistemic standards one must meet when one treats a proposition as a reason so that the resulting piece of practical reasoning has no shortcomings. But first, I will spell out the idea that practical reasoning is an intellectual performance.

Sosa (2011) proposed an analysis of the performance of forming beliefs based on observations of performances generally that provides also a useful toolkit for analyzing practical reasoning. Performances can be assessed along various dimensions: accuracy, adroitness, and aptness. A performance is accurate if it is successful. A performance is adroit if it is skillful, that is ably executed. A performance is apt if it is successful *because* it was skillful, so aptness is success through ability. For

⁷ Some hold that the conclusion of practical reasoning is an action. My arguments are compatible with this view as long as the action is also accompanied by an intention to act.

example, the archer's shot is adroit if it manifests the archer's shooting competence. The shot is accurate if the archer hits the intended target. And the shot is apt if the archer hits the intended target *because* of a manifestation of her shooting competence. Practical reasoning as an intellectual performance can also be assessed along these three dimensions.

Practical reasoning is adroit only if the exercise of one's intellectual abilities determines an intention to act. For example, through an exercise of one's intellectual abilities, one discerns considerations to be treated as reasons and reasons from them to an intention. I will not offer a full account of the relevant intellectual abilities for practical reasoning. Perhaps obviously, not just any use of intellectual abilities will suffice for adroitness. Even silly reasoning, such as reasoning that is entirely oblivious of actual means-end-relations, will involve some use of intellectual abilities. But silly reasoning is not adroit. What matters for my purposes is that the intellectual abilities involved in practical reasoning and its evaluation need not be unique to practical reasoning. Consider a comparison with an archer's shot. It would be absurd to hold that an ability that is involved in the archer's shot cannot influence the evaluation of the shot unless it is unique to this activity. For example, an exercise of one's perceptual abilities is certainly part of taking a shot, but not unique to taking a shot. Likewise, we should not expect that all intellectual abilities involved in practical reasoning and influencing its evaluation must be unique to it. Since practical reasoning depends on doxastic states, the performances that brought them about affect one's practical reasoning performance, just as the exercise of perceptual abilities affects one's shooting performance. Plausibly, both the specific ability involved in formation of a doxastic state as well as its resulting epistemic status affect the status of one's practical reasoning performance and whether it counts as adroit.

What constitutes success and thus accuracy in practical reasoning? The aim of practical reasoning is to find a suitable course of action relative to the reasoner's end. This aim defines what counts as an accurate episode of practical reasoning. Practical reasoning is accurate if it leads to an intention to act in a way that can realize the reasoner's end given the reasoner's actual circumstances.⁸

A further clarification on accuracy: practical reasoning is an intellectual performance. It is accurate when it determines a course of action that can realize one's end. The subsequent action need not actually realize it, as for that an able execution of the chosen course of action is also necessary. But a failure in execution is not a failure of practical reasoning. Suppose that Jason in the good case trips on the way to the fridge, breaks a leg and then makes a trip to the ER instead of having that glass of milk. Still, his practical reasoning satisfies the above sense of accuracy. The intended action was in principle fitting to realize the end given his actual circumstances, even if its failed execution does not realize his end.

Generally, the 'can' introduced here is weak, since, given the right circumstances, an action can realize an end even if it has an extremely low probability of doing so.

⁸ This sense of accuracy does not imply that the chosen way to act will be the most effective means to achieve the end. I assume that choosing effective means is a matter of adroit practical reasoning, but will not discuss this further, as it is irrelevant to my main claim.

One's reasoning can be accurate even if the intended course of action requires luck to realize one's end. However, and I will get back to this point, this sort of luck is different from being lucky in arriving at an intention to act in a way that is suitable to the realization of one's end. For example, there is a difference between a surgeon that consciously intends to perform a complicated surgery that requires some luck to save the patient and a surgeon that luckily ends up intending to perform such a surgery.

Finally, practical reasoning is apt if it is accurate because it was adroit, that is the success is due to an exercise of the intellectual abilities of the reasoner. Aptness is success through one's own efforts.⁹ Apt practical reasoning is an intellectual achievement that is creditable to the reasoner.

Let us go back to the good case to explicate all notions just introduced. Jason has a desire to have a glass of milk and he goes to the fridge because he remembers that there is milk in it. Jason's reasoning is accurate. His practical reasoning leads to an intention to act suitable to fulfill his desire. The reasoning is adroit, because it manifests his intellectual abilities in the sense that his intention is based on reasonable beliefs relevant to his desire. And his practical reasoning is apt because Jason, through an exercise of his intellectual abilities, reasons accurately, which makes his success creditable to him.

But the three dimensions of assessment can come apart. It is possible that one's reasoning is accurate without manifesting any intellectual ability, just as an archer's shot can hit the target without an exercise of her shooting ability. On the other hand, Case F is an example of adroit but inaccurate reasoning. In this case, going to the fridge is not a fitting action to fulfill his desire for milk in Jason's actual circumstances, as a burglar stole the milk. This renders Jason's practical reasoning that concludes with the corresponding intention to act adroit but not accurate.

In Case G, Jason's reasoning is both adroit and accurate, but not apt. Jason displays intellectual ability and his intended course of action is suitable to fulfill his desire. His intellectual abilities are involved in determining his course of action. But the use of his intellectual abilities does not bring it about that he ends up choosing a course of action that can realize his end. Rather, the interplay of good and bad luck—the milk being stolen and roommate John accidentally buying some—plays a significant role in explaining why Jason's reasoning is accurate. Since Jason's reasoning is not accurate *because* it was adroit, it is not apt.

This case is analogous to Sosa's archer case in which the archer exercises her shooting competences and hits the target, but the arrow was affected by two gusts of wind that luckily cancel each other out. Such a shot is not apt as it is not accurate *because* of an exercise of the archer's shooting competences.¹⁰ In the milk case, the burglar stealing the milk and the roommate accidentally buying milk are the two gusts of wind that luckily cancel each other out. They make it that Jason's practical reasoning is accurate and adroit, while it is not true that his accuracy is achieved through the use of his intellectual abilities.

⁹ I borrow this metaphor from Wedgwood (2002: p. 283).

¹⁰ See Sosa (2011: p. 4).

In light of Case F, one may want to question my claim that Jason's reasoning is not accurate and my suggestion regarding accuracy. Here is an alternative suggestion: a successful performance of practical reasoning concludes with a rational intention, as reasoning in general is successful if it is rational. Accuracy does not require an external form of success, such as the suitability of the intention to realize the agent's end, which depends on external circumstances. Hence, one's reasoning is accurate if it is rational.

There certainly is a sense in which one's practical reasoning is successful as long as it is rational. But this alternative suggestion itself fails to provide a strong reason to abandon my suggestion about accuracy. Within the performance normativity framework, it is natural to evaluate a performance along external dimensions. The archer shoots to hit some target, hence her shot cannot only be evaluated for how skillful it was, but also as to whether it hits the target. Similarly, we engage in practical reasoning to bring about a desired end; at least most of us do not merely engage in practical reasoning to be rational. Arguably, one can engage in practical reasoning for other purposes, such as making a defensible choice. But ultimately one makes a choice to pursue an end, which one then executes. And hence, practical reasoning can be evaluated as to whether it succeeds in choosing a course of action suitable to bring about the intended end. This mirrors what one may call theoretical reasoning about what to believe. Just as we can evaluate whether such reasoning is rational, we can evaluate whether it succeeds in leading to true beliefs. Surely, this is a quite natural evaluation of theoretical reasoning, even if somewhat external. The same holds for my suggestion about accurate practical reasoning.

Furthermore, my proposal treats practical reasoning analogous to other performances, such as shooting arrows, as it suggests that the various evaluative dimensions can come apart. On the alternative suggestion, all three evaluations coincide, as adroitness would then be equal to accuracy (and therefore aptness). This makes practical reasoning as a performance peculiar from other performances.

While the alternative suggestion is coherent and natural, my account of accuracy is just as natural and coherent, and overall allows for a more comprehensive evaluation of practical reasoning. I admit that this is not a conclusive reason against the alternative suggestion and further debate is desirable, though beyond the scope of this paper. At the same time, the mere existence of the alternative suggestion is not a conclusive reason to abandon my account of accuracy and I assume that what I said is sufficient to stick with it in the meantime.

Having three dimensions of assessment on the table, adroitness, accuracy, and aptness, which one should we strive to meet? Adroit practical reasoning lives up to one standard of assessment, but it may fail to live up to another standard that is just as important: accuracy. As suggested above, it is natural to see the point of practical reasoning in securing success when it comes to acting in the world. To be indifferent about accuracy is to be blind to the practical significance of reasoning. Inaccurate but adroit practical reasoning is somewhat defective practical reasoning. Similarly, accuracy alone cannot be all that matters. After all, practical reasoning is an intellectual performance. Accuracy can be achieved in the absence of the proper use of one's intellectual abilities. But such reasoning is clearly defective.

Not even the conjunction of adroitness and accuracy is fitting as the proper standard we should strive for. Both dimensions are met in Case G. Yet, the reasoning, as an intellectual performance, is defective because the exercise of Jason's intellectual abilities is disconnected from the facts that lead to his accuracy. It is simply a matter of luck that Jason is accurate. But practical reasoning is an intellectual performance. There would be little point in exercising intellectual abilities if they would not contribute to being accurate. So one's intellectual abilities should figure in the explanation of one's accuracy. The connection between one's reasoning abilities and one's accuracy is only present in cases in which one's reasoning is apt, because aptness is accuracy through ability. Because anything less than apt practical reasoning falls short in some respect, aptness is the dimension of assessment we should strive to meet, which simply means that our practical reasoning ought to be apt. I will say more about the normativity of a norm that makes aptness central in Sect. 6, where I suggest that the sense of 'ought' is evaluative rather than prescriptive.

What is the upshot of all these considerations? It has been assumed that assessments about doxastic states bear on how we assess practical reasoning that relies on these doxastic states. For example, the doxastic states one has towards propositions treated as reasons affect how we assess the resulting practical reasoning. Adding the thought that practical reasoning ought to be apt suggests that an epistemic norm that concerns what is treated as a reason should provide an epistemic condition for apt practical reasoning. Aptness differs from the notion of rational permissibility, entitlement or epistemic success discussed in the previous section. Aptness might entail all of them, while none of them entails aptness. In this sense, I am proposing a novel variant of the knowledge norm, as I will argue that knowledge is the epistemic condition for apt practical reasoning:

3.1 Knowledge norm for apt practical reasoning (KAP)

Practical reasoning in which p is treated as a reason can be apt only if one knows that p .

What one treats as a reason affects the larger performance of practical reasoning. KAP holds that only if we treat as a reason a known proposition, will there be no negative downstream effect so that one's reasoning can be apt. This is of course compatible with the possibility that even when one knows, one's reasoning can fail to be apt (perhaps the reasoning towards the conclusion was defective). If our reasoning ought to be apt and if KAP is true, then it follows that we ought to treat a proposition as a reason only if we know it. In this sense, KAP expresses a norm.

4 A knowledge norm for apt practical reasoning

My aim in this section is to offer an abductive argument for KAP based on a number of cases. Apt practical reasoning was defined as accuracy that is achieved through the use of intellectual abilities. While we have an intuitive grasp of this notion, as the examples suggest, it is extremely difficult to provide an account of 'success

through the use of intellectual abilities' in terms of positive conditions.¹¹ But we can still approach this task indirectly. The following cases suggest that what we can broadly call anti-luck conditions on knowledge and apt practical reasoning are either identical or co-extensional. I will show that standard cases of knowledge-undermining-luck, e.g. Gettier cases, fake barn cases, and lottery cases, are all cases in which one's practical reasoning fails to be apt. The best explanation of why one's practical reasoning fails to be apt is that aptness requires knowledge. This supports KAP.

Let us consider Case G once more. While Jason has a justified true belief, the influence of luck keeps him from having knowledge. Up to this point, this is just a regular Gettier case. But once Jason treats the content of his Gettiered belief as a reason, his further reasoning determines an intention to act in a way that can realize his end, but this is significantly due to luck, not solely because he exercised intellectual abilities. Thus, as I explained in the previous section, this is not a case of apt practical reasoning.

The moral of this case generalizes. Take any Gettier case and let the Gettiered subject treat the content of her Gettiered belief as a reason. The resulting practical reasoning will not be apt because one's success will be due to luck, not due to an exercise of one's intellectual abilities. The knowledge-undermining-luck brings it about that one's reasoning fails to be apt. Therefore, beliefs may not be subject to knowledge-undermining-luck in order for them to be capable of leading to apt practical reasoning. Clearly, any piece of knowledge will meet these requirements. Although caution is warranted here, it seems that any belief that meets these requirements will simply be knowledge. This suggests that knowledge is necessary for apt practical reasoning, which is what KAP holds. This can further be bolstered if we go back to the good case in which there is no knowledge undermining luck. Jason's belief based on his memory would have been knowledge and it would no longer seem that his practical reasoning is successful because of luck.

Next, let us consider a variant of the fake barn scenario. Henry, unbeknownst to him, drives through barn facade county. Henry's visual perception works normally. When he drives by a building that looks like a barn to him, he rationally forms the belief that there is a barn in front of him. His belief happens to be true because he is looking at the only real barn in barn facade county. Yet, his belief is not knowledge because in this abnormal environment, it is too lucky to be knowledge.

Now assume that Henry treats his belief that there is a barn ahead as a reason to go and seek shelter there for the night, and that he forms the corresponding intention based on this reason. His practical reasoning is adroit and accurate, after all he is looking at a real barn. But it is not apt practical reasoning. It is due to luck that his reasoning is accurate given the environment that he is in. To achieve success through intellectual ability sometimes requires very specific abilities. Sometimes, an ability that would be sufficient in ordinary circumstances just won't do. Henry's circumstances are special in just this way. Henry's perceptual abilities do not enable him to distinguish between real and fake barns. Even though he relies on an otherwise

¹¹ See Turri (2011) and Miracchi (2015) for a detailed exposition of this problem, see Wedgwood (forthcoming) for the latest suggested solution.

reliable ability, the fact that this ability produces a true belief in this situation is accidental. Likewise, it is accidental that Henry's reasoning is accurate, as he simply does not have the ability required to make his success non-accidental and creditable to him. Thus, just like in Case G, while Henry exercises an intellectual ability, his success is not due to this intellectual ability. Consequently, Henry's practical reasoning is not apt.

Just like in the Gettier case, what stands in the way of Henry knowing is exactly what stands in the way of his reasoning being apt. That Henry is in barn facade county and that he so easily could have ended up with a false belief undermines both him having knowledge and his practical reasoning being apt. If he had not been in barn facade county, Henry's belief would have been knowledge and his practical reasoning would have been apt. This difference between these two cases can be explained by KAP.

Finally, consider reasoning based on lottery propositions. Let us assume that I have a ticket in a lottery with a million tickets, just one of which is a winner. The winning ticket gets one million dollars. The winner has been drawn, but not yet announced. Assume that I believe that my ticket is a loser based on the statistical unlikelihood that I have the winning ticket. Suppose further that I treat this consideration as a reason to sell the ticket for one penny, because I would prefer to have the safe penny over the very unlikely million dollars, and form the corresponding intention. Setting aside whether such a piece of reasoning is adroit, it is not apt.

It is relatively uncontroversial that one cannot know that one's ticket is a loser just by relying on statistical facts about the likelihood of having a winning ticket.¹² Admittedly, it sounds odd at first, but in some sense, one ending up with a true belief is lucky, although highly likely. And the same goes for relying on this belief in practical reasoning. While it is highly likely that one's reasoning will be accurate if one treats this proposition as a reason, this success is not achieved through the relevant intellectual abilities to make this a case of success through ability—just as in the fake barn case. In order for my accuracy not to be lucky and due to an exercise of intellectual ability, I would need an ability to discriminate winning from losing tickets. But I do not have an ability to do so based on merely statistical considerations. As Hawthorne (2004: p. 15) puts the point: if one divides the space of possibilities into subcases, one's epistemic standing in each subcase is not appreciably different. I simply do not have an ability to discriminate between a subcase in which I have a winning ticket and one in which I do not. And because I lack such an ability, the reasoning cannot be successful through an exercise of intellectual ability, and thus not apt. This is another case in which luck that undermines having knowledge also undermines apt practical reasoning. Following the reasoning explicated above, KAP provides an explanation for this.

To avoid misunderstanding: the notion of luck involved in my arguments is not to be equated with probabilistic unlikelihood. Some events that are objectively highly likely can in some sense still be lucky. In this sense, the lottery case is analogous to the fake barn scenario. Just as one is lucky to pick the one real barn from all the fake

¹² See Hawthorne (2004) for a book length treatment of this topic.

ones, one is lucky to have a losing ticket even though this is overwhelmingly likely. One is lucky because the intellectual abilities exercised can no more tell fake barns from real ones than they can tell winning from losing tickets.

The three cases given in this section demonstrate that where there is knowledge undermining luck, there is also apt practical reasoning undermining luck. In all cases mentioned, one's success is not due to one's intellectual abilities, but due to the same form of luck that also keeps one from knowing the proposition in question. Whatever this elusive notion of luck is that is ruled out by having knowledge, this epistemic luck at least coincides with practical luck undermining one's practical reasoning from being apt. KAP offers a simple explanation of this pattern. Because only knowledge rules out the sort of luck that is present in Gettier cases, fake barn cases, and lottery cases, we have reason to assume that knowledge of propositions one treats as reasons is necessary for apt practical reasoning. This suggests that KAP is true.

5 Four objections

In this section, I will address four objections to KAP. The first is what I call the triviality concern. One might say that if one adopts Sosa's framework in which apt belief is knowledge and applies it to practical reasoning, which depends on beliefs, then it is trivial that apt practical reasoning will require knowledge.

But this triviality concern is ill-founded. It seems to be based on the assumption that if there are two performances, in this case belief and practical reasoning, such that the latter depends on the former, then there must a match between the evaluative status that both performances attain. While this principle might be intuitively appealing, it is false, as I will demonstrate below.

More importantly though, the objection assumes that apt belief is knowledge. While this is Sosa's view, I do not endorse it. I merely argued that apt practical reasoning requires knowledge and I used the performance normativity framework to do so. I have not argued or assumed that apt belief is knowledge. In fact, the fake barn case I used for my case for KAP is classified as a case of apt belief, and thus as a case of knowledge, in Sosa (2007). If this was true, this would obviously undermine my case for KAP. Hence, I must at least reject Sosa's verdict for the fake barn case, or reject his equation of apt belief and knowledge. Given the common view that the fake barn case is not a case of knowledge, this divergence is not dramatic, I assume. But it shows that KAP is not a trivial truth, as it is not trivial how to apply Sosa's performance normativity framework to beliefs (is the fake barn case a case of apt belief or not?), nor is it trivial whether to accept Sosa's equation of apt belief and knowledge.

The second objection is based on what one might call the truth concern. Assume that I desire to have a painted wall in my apartment but I am indifferent about the specific color. I have a justified but false belief that the bucket that I bought contains blue paint, which I treat as a reason to paint my wall with the paint from this bucket. However, the bucket contains red paint. Nonetheless, given my indifference about the color, my false belief is still sufficient for choosing a course of action suitable to

realize my end of having a painted wall. Thus accuracy in practical reasoning does not always require true beliefs. Since aptness entails accuracy, we should be suspicious about whether apt practical reasoning requires true belief. Consequently, we should be suspicious about KAP, because a demand for knowledge entails a demand for truth.

This objection does suggest that accurate practical reasoning does not always require true beliefs. But I can concede this without undermining KAP, as it is not plausible that the case is an instance of apt practical reasoning. It is not because the success is not attributable to the intellectual ability actually exercised. The intellectual abilities actually involved in the reasoning make it that I treat a false proposition as a reason. But the chosen course of action is not successful because one treats this false proposition as a reason, as there are many ways of this proposition being false that would not result in having a painted wall, i.e. that the bucket contains blueberry jelly. Rather, it is because the false proposition happens to entail a true one that luckily suffices for making the chosen course of action sufficient to achieve my end. But my success cannot be attributed to the intellectual abilities actually exercised, as it is due to a lucky entailment. Since the case is not an instance of success through ability, it is not a case of apt practical reasoning and therefore does not undermine KAP. This also partly justifies the assumption that apt practical reasoning requires true beliefs—at least we have not seen any evidence to the contrary.

Here is a variant of the case. Assume that everything is as above, but that I also know that the bucket contains paint. So it could not turn out that it is false that the bucket contains blue paint, but that it does not contain paint at all. Nonetheless, this variant can be dealt with similarly. Even if I know that the bucket contains paint, my reasoning, qua stipulation, remains based on the falsehood that the bucket contains blue paint. Therefore, my success is still due to a lucky entailment. This is true even if I know said entailed proposition. Since, qua stipulation, knowledge of the entailed proposition does not figure in my practical reasoning, it is due to luck that my practical reasoning is successful, as there are many possibilities in which the relied upon falsehood would not have lead to the intended outcome. That one knew the entailed proposition does not change this, as the known proposition did not figure in my practical reasoning. If it did, then my practical reasoning may be apt, but then this variation would not contradict KAP.

The truth concern is also useful to demonstrate that the triviality concern is based on the assumption of a false principle. The paint case demonstrates that the evaluative dimensions of belief and practical reasoning do not always neatly align so that the argument for KAP would be trivial. In the paint case, my belief is not accurate because it is not true. Yet, my practical reasoning is accurate, as my intended action is suitable to realize my end. Therefore, the case demonstrates that accurate practical reasoning does not necessarily require accurate beliefs. Consequently, we should also not assume that apt practical reasoning requires apt belief, which gives us a further reason to assume that KAP is not trivially true.

The third objection concerns case pairs involving a non-traditional factor. Take the much discussed bank cases as one example. In a low stakes scenario, it is not important when Keith deposits his paycheck and he knows that the bank is open on Saturday based on his memory of a previous visit. If he treats this proposition

as a reason to postpone his visit till Saturday, then this could be an instance of apt practical reasoning. In contrast, many hold that Keith's belief that the bank is open on Saturday is not knowledge in a high stakes scenario where it is extremely important to deposit the paycheck by Saturday. According to KAP, his reasoning based on the belief would not be apt. But let us assume that Keith maintains his belief and treats it as a reason to postpone his visit till Saturday. Why, if Keith's reasoning is accurate, is his success not attributable to an exercise of his intellectual abilities, and thus apt? Why would a change in a practical factor between low and high stakes affect whether one's success is attributable to an exercise of intellectual abilities? This raises doubts about whether apt practical reasoning always requires knowledge.

There are several strategies available to deal with this objection. For all shifty epistemologists that endorse the shift in knowledge between the two scenarios,¹³ a strategy akin to the fake barn county scenario is available. For the fake barn scenario, I suggested that special scenarios call for special abilities instead of ordinary abilities. The same could be said about the high stakes scenario. Here, success from ability requires a specific ability different from the one that suffices for apt practical reasoning in the low stakes scenario. Shifty epistemologists can offer this explanation of why practical factors can also affect the aptness of practical reasoning.

Traditional views do have resources to account for the apparent shift in knowledge in the bank cases, while maintaining that Keith actually knows in both scenarios. However, they face pressure to explain why, despite his knowledge, Keith's practical reasoning in high stakes seems deficient, which puts pressure on KAP.

There are at least two strategies available to traditional views of knowledge that do not conflict with KAP. Traditionalists could maintain that Keith's reasoning is indeed apt, yet it might seem deficient because it fails to be meta-apt. Sosa (2011: p. 9) suggests that a performance is meta-apt if its aptness stems from a competence to assess risk properly. Arguably, Keith's risk assessment in high stakes seems questionable. Therefore, any seeming deficiency could be explained as a lack of meta-aptness.

However, the traditionalist could explain the seeming deficiency as a lack of aptness. This does not undermine KAP, which merely suggests that knowledge is a necessary condition for apt practical reasoning. KAP does not hold that knowledge is a sufficient condition. This reply does not undermine the previous argument for KAP, which relied on the idea that the anti-luck conditions on knowledge and apt practical reasoning are identical. The traditionalist could argue that the problem in high stakes is not that Keith does not meet the anti-luck condition, which he does, as he knows. But rather, his reasoning is not apt because it is not adroit. It is not adroit because while Keith's belief is justified, as he knows, his degree of justification is still too low to rationalize his intended course of action. Consequently, his reasoning is not adroit, and therefore also not apt.

¹³ See DeRose (1992) or Stanley (2005). I use the label 'shifty epistemology' to stay neutral on the question whether the alleged shift in knowledge-attribution is indicative of semantic context-sensitivity of 'knows', as suggested by DeRose, or whether the knowledge-relation itself is sensitive to practical factors, as suggested by Stanley.

To sum up: KAP can handle the bank cases in a number of different ways, accommodating a wide range of epistemological views, so that these case pairs do not amount to a decisive objection to KAP.

The fourth objection concerns partial belief and thus brings us back to Case P from Sect. 2. It is commonly held that knowledge requires full belief. But if KAP is true, then practical reasoning based on a proposition that is only partially believed cannot ever be apt. This cannot be right, can it?

This is right, given what it means for one's practical reasoning to be apt. Aptness is more demanding than rational practical reasoning, thus cases of partial belief do not have the same bite as used against RKP. In cases of partial belief, one's evidence is not sufficient to license a degree of confidence that leads to a belief that *p*. If one's evidence is insufficient for full belief, then treating *p* as a reason may be rational. But, to belabor the archer analogy, it is like an archer taking a lucky shot in the dark when she cannot fully see the target. Even if one ends up being successful, one's success is to a large extent due to luck and not fully creditable to one's intellectual abilities. And therefore practical reasoning in which *p* is treated as a reason, but *p* is only partially believed, cannot be apt. This is compatible with the possibility that one's reasoning was adroit and perhaps even accurate.

These considerations also explain why reasoning based on justified beliefs that fall short of the level of justification required for knowledge is not apt. Given the weakness in one's justificatory position, even if one's reasoning is accurate, it is more akin to a lucky shot in the dark than a case in which one's success is due to one's intellectual abilities.

Still, one might wonder. Suppose I have a sufficiently high enough credence for justified belief in the proposition that it will rain today, while falling short of knowledge-level justification. It seems fine to treat this as a reason to take an umbrella, while it would be rather odd to point to one's lack of knowledge as a reason to refrain from taking the umbrella. The former reasoning seems fine, while the latter does not. But why is that if the latter is concerned with meeting a knowledge norm?

To fully explain this, I will need to expand on the normativity of the suggested knowledge norm of apt practical reasoning. I will do so in the next section and return to this case. Meanwhile, note that the latter reasoning does not seem apt, as it seems not adroit. If one's credence is high enough for justified belief, then it is rather odd to weigh the fact that one does not know that it is raining against the fact that one is justified in believing that it is raining, which means that it is more likely that it will rain than not rain. This odd weighing displays a certain kind of incompetence in reasoning. Hence, the latter reasoning is not adroit, while no such shortcoming is present in the former reasoning. Still, one can wonder how a knowledge norm of apt practical reasoning fits in with these judgments. I will get to this in the next section.

Similarly to the fourth objection based on Case P, one might worry that KAP is simply too demanding. Suppose that somebody engages in a very long and complicated chain of practical reasoning, which relies on many different propositions and which requires skills of reasoning way above ordinary reasoning skills. They know all propositions except one, for which they have a Gettiered belief. If KAP was true, then the resulting practical reasoning could not be apt. Isn't this too harsh, if not false, because they seem to display outstanding abilities altogether?

I stand by the harsh verdict that this practical reasoning, even if in many respects outstanding, is not apt. We should not confuse aptness with other evaluations. The fact that the reasoner might have reasoned in an outstanding manner does not guarantee that the reasoning was apt. Given my specific understanding of aptness, it is fitting to distinguish outstanding performances from apt performances.

This harsh verdict, as well as the general demand for knowledge that KAP places on practical reasoners, may give rise to concerns about how KAP can be a genuine norm. The next section, among other things, answers such concerns.

6 The normativity of the knowledge norm and its significance

Even if we accept KAP as true, an opponent might question whether it is indicative of a knowledge norm for practical reasoning. One might question that the claim that one's reasoning ought to be apt, fulfills criteria of genuine norms. In this section, I will consider and answer arguments that raise this issue. The central idea to be developed is that the knowledge norm I suggest is to be understood as an evaluative norm.

The first argument is based on the principle 'ought implies can' (OIC). Aptness entails accuracy, which is not under the control of an agent. It will be unavoidable that at some point in time, one will rely on a false justified belief in practical reasoning and thereby engage in reasoning that will not be accurate and therefore not apt. But given OIC, and since we cannot always reason aptly, aptness cannot be normative and we cannot derive a genuine norm from KAP.

The problem with this argument is that OIC has a number of different readings because 'ought' and 'can' have several meanings, as Wedgwood (2013) and others have noted. On some readings, OIC is not a universally accepted principle.¹⁴ One needs additional arguments establishing that genuine norms express an 'ought' that implies a certain kind of 'can', and that my claims violate this principle. I will return to what sort of 'ought' my claims do express. Meanwhile, the objection, as an unspecific appeal to OIC, can be reasonably set aside.

According to the second argument, there is a strict distinction between the evaluative and the normative.¹⁵ Perhaps it is good to reason aptly, but this is merely an evaluative claim, as the talk of evaluative dimensions of practical reasoning suggests. But it is not the case that there is a genuine 'ought' associated with this evaluation that suggests that one's practical reasoning ought to be apt or that one ought to treat *p* as a reason only if one knows that *p*.

This argument poses no genuine problem because there are evaluative norms. For example, Chrisman (2009) and McHugh (2012) both defend the truth norm for belief as an evaluative norm. McHugh (2012: p. 9f) writes that evaluative norms are concerned with what is good, but that they do not presuppose that one is blameworthy for violating them. The notion of evaluative norms suggests that the distinction

¹⁴ Rejections of certain forms of OIC can be found in Sinnott-Armstrong (1984) and Graham (2011).

¹⁵ This strategy is employed against a truth norm for belief in Gibbons (2013).

between the evaluative and the normative is not as clear-cut as some assume. This also sits well with the observation that there is an evaluative sense of ‘ought’.¹⁶

I suggest that we understand the normativity of a knowledge norm derived from KAP in terms of evaluative norms, inspired by Chrisman and McHugh, but also defendants of a knowledge norm of belief (see Simion et al. (2016)). Following in their footsteps, I suggest that an evaluative norm is concerned with an ought-to-be, whereas prescriptive norms are concerned with an ought-to-do. To be concerned with an ought-to-be means to set a standard for what counts as good, where goodness, following Geach (1956), is understood in an attributive sense. An evaluative norm for practical reasoning determines what counts as good practical reasoning, just like an evaluative norm for knives determines what counts as a good knife (a sharp one). The phrase ‘good practical reasoning’ though is ambiguous. There certainly is a sense in which one’s reasoning is good if it is adroit. Yet, as I have explained before, merely adroit practical reasoning falls short in some regards. I understand good practical reasoning as practical reasoning without any shortcomings, and I suggest that the standard for good practical reasoning in this sense is aptness. KAP tells us which epistemic requirement one must meet for apt practical reasoning: one must know. Evaluative norms do not entail prescriptive norms, since from the fact that something ought to be a certain way, it does not follow that somebody ought to do something specific. This also provides a neat way to avoid any objection from OIC because it is plausible that an evaluative ‘ought’ does not always imply ‘can’. Nobody would deny that it would be good if we were to stop global warming, and thus that we ought to stop it, even if it turns out that we cannot stop it.¹⁷

Note that this move to understand the knowledge norm as an evaluative norm cannot simply be used to defend any of the criticized proposals from Sect. 2. For example, one might be tempted to say that the RKP is an evaluative norm. However, the RKP is understood to provide a condition on what is rationally permissible to be treated as a reason. Case F, G, and P suggest that rational permissibility does not require knowledge. Hence, it seems unmotivated to suggest that the RKP should be understood as an evaluative norm, as it seems that it can be equally good, given the concern with rational permissibility, to rely on beliefs that do not amount to knowledge. Things are different with KAP. According to KAP, knowledge of a proposition that is treated as a reason is a necessary condition on the resulting practical reasoning to be apt. Therefore, if aptness is the ultimate evaluative dimension of practical reasoning, as I argued in Sect. 3, then knowledge is the norm for apt practical reasoning. One ought, understood evaluatively, to treat p as a reason for action only if one knows that p .

In closing, I want to point out one upshot of the overall picture I have been arguing for. An observation used to motivate traditional accounts of a knowledge norm

¹⁶ See for example Schroeder (2011) or Wedgwood (2013) for the claim that there is an evaluative sense of ‘ought’.

¹⁷ McHugh (2012: p. 19) denies that evaluative norms imply ‘ought’, but I take him to mean that they do not imply a prescriptive ‘ought’, an ‘ought’ that implies blameworthiness for norm violations.

like the RKP is that 'knows' figures prominently in our assessment of practical reasoning. Even critics concede that this calls for an explanation if it is not true that knowledge is the norm for practical reasoning (see for example Locke (2015: p. 92). On my account, this fact needs not to be explained away.

If KAP is true, then there is a variant of the knowledge norm that is true. The fact that 'knows' figures prominently in our appraisals of could simply be due to the fact that our actual practices are sensitive to this knowledge norm. We implicitly know that practical reasoning based on knowledge is optimal practical reasoning. It guarantees that the objectives of our intended actions will not be frustrated because the world is not as we take it to be. It is thus very natural that we criticize people for not knowing something that is relevant for their choice of a course of action when such knowledge is in principle attainable in their situation. They are worthy of criticism because they needlessly settled for a suboptimal practical reasoning process. However, my suggestion also helps to explain why not just any lack of knowledge will lead to criticism, as Case F and Case G demonstrate. As McHugh writes, evaluative norms do not imply that we are always blameworthy for failures to live up to them. This explains why we do not criticize behavior in these cases—because our lack of knowledge is not something that we are blameworthy for and we only inadvertently fail to engage in optimal practical reasoning. In those cases it remains true that our practical reasoning ought to be apt, understood evaluatively, and it will also be true that our reasoning is in some sense defective because it is not apt. Yet, we are not worthy of criticism because in these cases, we do not blameworthy fail to live up to the norm of apt practical reasoning. This explains why 'knows' figures prominently in our appraisals, but also why we may refrain from using it in our appraisals in specific situations.

Case P requires a slightly different treatment, as here one's lack of knowledge is quite transparent. Yet, criticizing one for not knowing seems out of line. KAP as an evaluative norm can help explain this as well. Aptness, that is optimal practical reasoning, is a demanding standard. It is useful to have such a norm as a guide as to what to aim for. But particularly when we face inconsequential decisions, as in Case P, one can reasonably settle for less than the optimum because obtaining knowledge might be more bothersome than a failed action. Hence, criticizing a lack of knowledge seems out of line, while one can concede that the reasoning fails to live up to a norm and is in some sense defective.

Finally, we can deal with the case from the previous section that involves a high credence in a proposition that is sufficient for justified belief, while falling short of knowledge-level justification. Understanding KAP to be an evaluative norm helps to explain why one need not refrain from treating such a proposition as a reason. To hold that our practical reasoning ought to be apt does not imply that one must always avoid reasoning that fails to be apt, as evaluative norms do not entail prescriptive norms. When aptness is out of reach, it can be perfectly permissible to settle for reasoning that is not apt. It can be perfectly permissible to treat a proposition as a reason when one is justified in believing but acknowledges a lack of knowledge, for example when one is forced to make a decision immediately without the possibility of further inquiry. In such a situation, criticizing a lack of knowledge seems out of line. Still, there also remains a sense in which one's practical reasoning

would be better if it were based on knowledge. In some sense, one's reasoning here is defective, though not worthy of criticism, and the suggested evaluative reading of a knowledge norm of apt practical reasoning captures this.

7 Conclusion

I have argued for a novel variant of a knowledge norm: knowledge is the norm for apt practical reasoning, understood as practical reasoning that is successful through the exercise of intellectual abilities. I suggested that we understand this norm to be an evaluative norm and how this explains why knowledge prominently figures in ordinary appraisals of practical reasoning.

Acknowledgements Predecessors of this paper were presented at conferences in Calgary, Salzburg, and Cracow. I am thankful for helpful comments from these audiences, as well as to those from two anonymous reviewers for Synthese. Special thanks to Susanne Mantel, Jake Ross, Ralph Wedgwood, and Marcus Willaschek for commenting on earlier versions of the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL.

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

- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Nous*, 42(2), 167–189.
- Chrisman, M. (2009). Ought to believe. *Journal of Philosophy*, 105(7), 346–370.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178, 529–547.
- Gibbons, J. (2013). *The norm of belief*. Oxford: Oxford University Press.
- Graham, P. A. (2011). 'Ought' and abilities. *The Philosophical Review*, 120(3), 337–382.
- Geach, P. (1956). Good and evil. *Analysis*, 17(2), 33–42.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Jackson, A. (2012). Two ways to put knowledge first. *Australasian Journal of Philosophy*, 90(2), 353–369.
- Kelp, C. (2016). Justified belief: Knowledge first-style. *Philosophy and Phenomenological Research*, 93(1), 79–100.
- Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research*, 90(1), 72–95.
- Mantel, S. (2013). Acting for reasons, apt action, and knowledge. *Synthese*, 190(17), 3865–3888.
- McHugh, C. (2012). The truth norm of belief. *Pacific Philosophical Quarterly*, 93(1), 8–30.

- Mehta, N. (2016). Knowledge and other norms for assertion, action, and belief: A teleological account. *Philosophy and Phenomenological Research*, 93(3), 681–705.
- Miracchi, L. (2015). Competence to know. *Philosophical Studies*, 172(1), 29–56.
- Mueller, A., & Ross, J. (2017). Knowledge Dethroned. *Analytic Philosophy*, 58(4), 283–296.
- Neta, R. (2009). Treating something as a reason for action. *Nous*, 43(4), 684–699.
- Schellenberg, S. (2018). *The unity of perception*. Oxford: Oxford University Press.
- Schroeder, M. (2011). Ought, agents, and actions. *The Philosophical Review*, 120(1), 1–41.
- Sinnott-Armstrong, W. (1984). ‘Ought’ conversationally implies ‘can.’ *The Philosophical Review*, 93(2), 249–261.
- Silva, P., Jr. (2017). Knowing how to put knowledge first in the theory of justification. *Episteme*, 14(4), 393–412.
- Simion, M. (2019). Knowledge-first functionalism. *Philosophical Issues*, 29, 254–267.
- Simion, M., Kelp, C., & Ghijsen, H. (2016). Norms of belief. *Philosophical Issues*, 26, 374–392.
- Smithies, D. (2012). The normative role of knowledge. *Nous*, 46(2), 265–288.
- Sosa, E. (2007). *A virtue epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2011). *Knowing full well*. Princeton: Princeton University Press.
- Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosopher’s Imprint*, 11(8), 1–11.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16, 267–297.
- Wedgwood, R. (2013). Rational ‘ought’ implies ‘can.’ *Philosophical Issues*, 23, 70–92.
- Wedgwood, R. (2020). “The internalist virtue theory of knowledge”, *Synthese*, 197(12), 5357–5378.
- Weisberg, J. (2013). Knowledge in action. *Philosopher’s Imprint*, 13(22), 1–23.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213–235.

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