ORIGINAL ARTICLE

Insufficient reasons insufficient to rescue the knowledge norm of practical reasoning: towards a certainty norm

Jacques-Henri Vollet[1]

Abstract

A certain number of philosophers are attracted to the idea that knowledge is the epistemic norm of practical reasoning in the sense that it is epistemically appropriate to rely on p in one's practical reasoning if and only if one knows that p. A well-known objection to the sufficiency direction of that claim is that there are cases in which a subject supposedly knows that p and yet should not rely on p. In light of the distinction between sufficient and insufficient reasons, some philosophers contend that these cases are inconclusive. In this paper, I argue that this insufficient reason manoeuvre is defective because it either misconstrues the relevant cases or is at odds with strong intuitions about how we (should) reason. I then put forward further considerations relative to the instability of some pieces of reasoning and show how they can be explained by a certainty norm for practical reasoning.

Keywords Knowledge · Norm · Certainty · Action · Reason · Reasoning

1 The main objection to the sufficiency of knowledge for practical reasoning

A certain number of philosophers are attracted to the idea that there is an epistemic norm of practical reasoning in the sense that it is permissible to rely on a proposition in one's practical reasoning only if one satisfies a certain epistemic condition with respect to this proposition. It seems clear, for example, that if you have a time-limited reservation for a restaurant and go on your hunch that the restaurant is down the street, instead of asking where the restaurant is, you can be prima facie criticised (Hawthorne & Stanley, 2008: 571).

Still, what this epistemic condition amounts to has been the object of numerous debates (see Gerken & Petersen, 2020). A very influential proposal is that this

✉ Jacques-Henri Vollet
  jacquesvollet@yahoo.fr

1   Collège de France, Paris, France

condition is knowledge (Hawthorne, 2004; Hawthorne & Stanley, 2008; Stanley, 2005; Williamson, 2005):

(KN) It is epistemically permissible to rely on p in your practical reasoning if and only if you know that p.

There is, however, a strong and well-known objection to the sufficiency direction of that norm.[1] There are cases in which a subject supposedly knows that p without being thereby in a good enough epistemic position to appropriately rely on p (Brown, 2008, Reed, 2010, Gerken, 2011: 228–229, Beddor, 2020). One such case is proposed by Neta (2009: 688):

HISTORY EXAM. You are taking an oral history exam, and you come across the question in what year was Abraham Lincoln assassinated? You know that the answer is 1865. But you are momentarily struck by a neurotic diffidence about your memory of this historical fact. After you hear the question, the first thought that goes through your mind is: I believe that the answer is 1865, but of course I don't know that it is! Although the latter conjunct is false -- you do know that the answer is 1865 – it's also true that you believe that you don't know that the answer is 1865, and it's furthermore true that you don't believe that you do know that the answer is 1865. A fortiori, you do not justifiably believe that you know that the answer is 1865. In the end, you decide to answer the question by saying 1865, on the basis of the proposition that the answer is 1865 as a reason for so answering the question.

It seems that it is inappropriate for you to treat the proposition that the answer is 1865 as a reason for answering the question, even if you do know that the answer is 1865.

Neta uses HISTORY EXAM to argue in favour of a justified-belief-that-one-knows norm of practical reasoning (2009: 686):

(JBKN) Where S's choice is p-dependent, it is rationally permissible for S to treat the proposition that p as a reason for acting if and only if S justifiably believes that she knows that p.

Consider also this most discussed case, from Brown (2008, 176):

SURGEON. A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:
- Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

---

[1] There are also objections to the necessity direction. See Brown (2008), Neta (2009), and Gerken (2011). For possible replies, see for example Hawthorne and Stanley (2008) Kelp and Simion (2017), Williamson (forthcoming), and Vollet (2022).

- Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

While SURGEON challenges the sufficiency direction of KN, it also challenges the sufficiency direction of JBKN. Indeed, not only does the surgeon know that she ought to remove the left kidney, she also seems justified in believing that she knows that. Yet she should double check before operating, which strongly suggests that she does not satisfy the epistemic norm of practical reasoning with respect to this proposition.

In the remainder of this paper, I consider the reply based on the distinction between sufficient and insufficient reasons advocated by some proponents of KN or JBKN (Sect. 2). I argue that it is defective because it either misconstrues the relevant cases or is at odd with strong intuitions about how we (should) reason (Sect. 3). I then put forward further considerations relative to the instability of some pieces of reasoning and show how they can be fully explained by a certainty norm for practical reasoning (Sect. 4).

2 The insufficient reason manoeuvre

Some writers respond to the objection arising from HISTORY EXAM and SURGEON by appealing to the well-established distinction between sufficient and insufficient reason.[2] Let us say you have a good reason to go to the swimming pool: you like doing sport. Still, you may have a stronger reason to stay at home: you ought to look after your children. The first reason is an insufficient reason you have to go to the swimming pool and the second reason is a sufficient reason you have to stay at home.

With this distinction in place, these theorists maintain that the alleged counterexamples to KN (or to JBKN) can be understood as cases in which the target proposition (or fact) is not in itself a sufficient reason to act, rather than as cases in which the epistemic norm of practical reasoning is not satisfied.

To illustrate the general thought, consider the following case:

JUDGE. John is a judge assessing Bill's culpability. An evil demon will kill John if John sentences Bill by relying on non-entailing evidence. In fact, Bill is guilty, and John acquires more and more evidence that Bill is guilty.

Arguably, at some point, that Bill is guilty is sufficiently supported by John's evidence for John to be epistemically permitted to rely on this proposition in order to φ, for any φ (at least assuming that the epistemic norm of practical reasoning is not Cartesian certainty). At that point, this proposition is a reason John has to sentence Bill. Still, it would not be appropriate for John to sentence Bill by relying on the

---

[2] A different line of response invokes antiluminosity (Williamson 2005). For criticisms, see Gao (2019) and Vollet (2023).

proposition that Bill is guilty. The evidence he has is not entailing and if John sentences Bill by relying on non-entailing evidence, the evil demon will kill him.

It seems clear that John does not have the epistemic position that is, overall, required with respect to Bill's culpability to treat the proposition that Bill is guilty as a sufficient reason to sentence Bill. However, that does not mean that John does not satisfy the epistemic norm of practical reasoning with respect to that proposition. In fact, that Bill is guilty may well be a sufficient reason John has to do other things, for example, to make someone else sentence Bill. The problem is just that, given John's practical situation, this reason is not sufficient for him to sentence Bill.

This case illustrates a situation in which a subject satisfies the epistemic norm of practical reasoning with respect to p (whatever that norm is) although she does not meet the overall epistemic requirement for appropriately relying on p in a particular course of action.

These considerations strongly suggest that the following thesis is false (see also Crisp, 2005):

> **(Equivalence Thesis)** The epistemic position that is, *overall*, required with respect to p for epistemically appropriate reliance on p is equivalent to the epistemic position required by the epistemic norm of practical reasoning.

Rejecting the equivalence thesis amounts to granting that there are cases where your epistemic position with respect to p is insufficient for you to rely on p, even if p favours doing A and p is a reason you have to do A because you satisfy the epistemic norm of practical reasoning with respect to p.[3] That paves the way for an explanation of HISTORY EXAM and SURGEON friendly to KN (and JBKN). One can suggest that, in those alleged counterexamples, the subject satisfies the epistemic norm of practical reasoning with respect to the target proposition, while maintaining that her epistemic position is *overall* insufficient to turn the reason she has into a sufficient reason to act.

Thus, Neta acknowledges that the surgeon knows the target proposition and that she is justified in believing that she knows. She satisfies JBKN and it is therefore epistemically permissible for her to rely on that proposition. Yet, according to Neta, it is also permissible, and in fact, it is required, to double-check that proposition.

Here is Neta's reconstruction of the surgeon's reasoning (Neta, 2009: 697):

SURGEON'S REASONING
Premise 1: The patient's left kidney is diseased.
Premise 2: Diseased kidneys must be transplanted
Lemma: Therefore, transplant the patient's left kidney.
Premise 3: Before any kidney transplant, it is absolutely imperative to double-check which kidney needs to be transplanted.
Conclusion: Therefore, double-check the patient's records.

---

[3] On the favouring relation, see Dancy (2004, ch. 3) and Skorupski (2010, ch. 2).

As we can see, the (alleged) fact that it is epistemically appropriate for the surgeon to rely on the proposition that the patient's left kidney is diseased (in Neta's words, to treat it as a reason) does not imply that she should not double-check it (Neta, 2009: 697).

Ichikawa (2012) makes a similar claim:

> It's entirely open to the defender of the knowledge norm to argue that knowledge of p is sufficient for p to be a reason, but that in this case, p isn't a good enough reason for action (…) [P]erhaps (…) only a proposition like the surgeon knows that the disease is in the left kidney would be a good enough reason for such a decisive action as removing a kidney. (Ichikawa, 2012, 51)

The insufficient reason manoeuvre is used by proponents of KN and JBKN primarily against objections arising from cases like SURGEON. However, it is worth noting that it can easily be generalised by proponents of KN to deal with objections based on cases like HISTORY EXAM.[4] Indeed, friends of KN may suggest that, in the context in question, that the answer is 1865 is not a sufficient reason to answer "1865", and that a sufficient reason should be something like "I know that the answer is 1865". If so, since the subject does not know that she knows that the answer is 1865, there is not a sufficient reason for her to answer "1865".

3 The failure of the insufficient reason manoeuvre

According to Ichikawa, we have in SURGEON merely "a pair of intuitive verdicts":

> 1. The surgeon knows (before reading the chart) that the disease is in the left kidney.
> 2. It would be inappropriate for the surgeon to remove the left kidney without first collecting more evidence.

If these are the only verdicts we have, it seems true that these cases do not speak against KN (or JBKN). Yet, these are not the only verdicts we have. There is the further verdict that it would be inappropriate for the surgeon to remove the left kidney without first collecting more evidence *because there is a costly possibility of error*. Brown's case does not merely suppose that the surgeon should double check, or that it is appropriate to collect more evidence before removing the kidney. It supposes that the surgeon should double-check *due to the cost of a possible error*. That is clear from what the nurse says: "imagine what it would be like if she removed the wrong kidney".[5]

To faithfully reconstruct the surgeon's reasoning, while assuming KN or JBKN, the proposed manoeuvre should then replace premise 3 by premise 3*:

---

[4] This point is surprisingly overlooked by Neta 2009, who offers HISTORY EXAM against KN and in favour of JBKN, and uses the insufficient reason manœuvre to defend JBKN against SURGEON.

[5] See also Brown (2011, 267): "the thought seems to be that knowledge is insufficient because of the chance of error".

SURGEON'S REASONING*
Premise 1: The patient's left kidney is diseased.
Premise 2: Diseased kidneys must be transplanted
Lemma: Therefore, transplant the patient's left kidney.
Premise 3*: Given that an error would be costly, it is absolutely imperative to double-check which kidney needs to be transplanted.
Conclusion: Therefore, double-check the patient's records.

Premise 3 was supposed to be true in virtue of rules of surgery, for example, or in virtue of the fact that surgeons must operate on what the chart indicates. But in virtue of what is premise 3* supposed to be true? Is there a general norm saying that when the cost of error is high we should double-check? It does not seem so. When the cost of error is high, it makes sense to check only if there is a (relevant) risk or possibility of error.[6] If so, premise 3* is true (and known) only if there is a (relevant) possibility that premise 1 is false.

But this creates a trouble for the insufficient reason manoeuvre. For, what is it to rely on p or to use p as a premise in reasoning? As Locke writes:

> Roughly, to premise that p in a particular piece of deliberation is to deliberate in a way that ignores, at least for the purposes of that deliberation, the possibility that not-p (Locke, 2015, 87).

There are at least two ways of vindicating that claim. First, as Beddor observes, there is a striking difference between relying on "p" and relying on "it is very (/0.99) likely that p". This difference is hard to explain if relying on p is compatible with taking seriously the possibility that not-p.

Second, as emphasized by Fantl and McGrath, we do not in fact find ourselves weighing "p" as a reason to φ against "there is a serious risk that not-p" as a reason not to φ (Fantl & McGrath, 2009: 80). Your reasoning is somewhat instable if you rely at the same time on "p" and "it might be that not-p". For example, suppose the surgeon premises, on the one hand, that the patient's left kidney is diseased and, on the other hand, that there is a chance that it is not the patient's left kidney which is diseased. The reasoning does not seem coherent. There is a simple explanation if to rely on p (or to use p as a premise) implies ignoring the possibility that not-p, but that is difficult to explain otherwise.[7]

Now, if relying on p implies ignoring the possibility that not-p, it is appropriate to rely on p only if it is appropriate to ignore the possibility that not-p. But if there is a relevant possibility that premise 1 is false, as an advocate of the insufficient reason manoeuvre has to grant to motivate the appropriateness of premise 3*, how can it be appropriate to use premise 1 in the reasoning, as this advocate also suggests? For, if there is a relevant possibility that premise 1 is false, it is not appropriate to ignore the possibility that not-p. Hence, it is not appropriate to rely on p.

---

[6] See Vollet (2023).

[7] If ignoring the possibility that not-p amounts to having no doubt as to whether p, or being certain that p, as Beddor (2020) suggests, to rely on p implies to be certain that p. It is important to note, however, that the present argument does not need to assume that relying on p implies being certain that p.

In summary, either the insufficient reason manœuvre avoids appealing to chances of error and it misconstrues the counterexamples to KN (and JBKN), or it appeals to chances of error and it must deny that to rely on p is to ignore the possibility that not-p (in a piece of reasoning). In the first case, the manoeuvre does not provide an adequate defence of KN against the proposed counterexamples. In the second case, it faces the challenge of providing an alternative explanation both of the intuitive difference between relying on "p" and relying on "p is very likely" and of the intuitive instability of relying on "p" and "it might be that not p" in the same piece of reasoning.[8]

4 Towards a certainty norm for practical reasoning

As the counterexamples to KN (and JBKN) show, at least sometimes, an epistemic position stronger than knowledge is required before relying on a given proposition. While several candidate norms can accommodate this fact, I will argue that a closer look at the phenomenon of instability points towards a certainty norm for practical reasoning.

To begin with, let us consider the following reasoning:

INSTABLE REASONING 1
Premise 1: If it is raining, I ought to take an umbrella.
Premise 2: If it is not raining, I absolutely ought to avoid taking an umbrella
Premise 3: It is raining.
Lemma: Therefore, I ought to take an umbrella.
Premise 4: But there is a possibility that it is not raining.
Conclusion: Therefore, I ought not to take an umbrella.

This reasoning clearly looks instable. As we have observed, this instability can be explained by the assumption that to rely on p is to ignore the possibility that not p. Putting aside this instability, the reasoning does not seem that bad, insofar as the proposed premises more reasonably support the conclusion than its negation. This means that premise 4 destabilizes premise 3, rather than the other way round. Is that due to the order of the premises?

Compare with the following reasoning:

INSTABLE REASONING 2
Premise 1: If it is raining, I ought to take an umbrella.

---

[8] At this stage, friends of KN might be tempted to abandon the insufficient reason manoeuvre and to embrace a shifty view of knowledge. Pragmatic encroachers, for example, say that given the cost of error in SURGEON, higher epistemic standards for knowledge are in place. They can then say that the surgeon in fact does not know the target proposition. While this move may have some plausibility in SURGEON (perhaps in that case the use of `know' is not literal, see Fantl and McGrath, 2009: 62) it is far less plausible to think that, in HISTORY EXAM, the subject does not (or cannot) know the target proposition. Assuming that would indeed lead to saying that in situations in which it is inappropriate for you to rely on p because you (justifiably) believe that you do not know that p, then you do not know that p. A (justified) belief that you do not know that p would be infallible. That is clearly an unwelcome result.

Premise 2: If it is not raining, I absolutely ought to avoid taking an umbrella.
Premise 3: There is a chance that it is not raining.
Lemma: Therefore, I should not take an umbrella.
Premise 4: But it is raining.
Conclusion: Therefore, I ought to take an umbrella.

This reasoning is similarly instable, but given the premises, the conclusion does not seem that reasonable. That is surprising, in particular, because in this reasoning, the conclusion logically follows from premises 1 and 4, while in the former reasoning, the conclusion does not logically follow from premises 2 and 4.

If these considerations are correct, we do not only have to explain the instability of pieces of reasoning using premises such as "p" and "it is possible that not p". We also have to explain why "it is possible that not-p" destabilizes "p", rather than the other way round.

This phenomenon of destabilization is a general phenomenon, happening also with premises invoking higher-order uncertainty. For example, Williamson (2005) considers a domino effect present in the following dialogue:

INSTABLE ASSERTIONS
Q1: Is q the case?
A: Yes.
Q2: Did you have warrant for your answer to Q1?
A: Yes.
Q3: Did you have warrant for your answer to Q2?
A: I don't know.

As Williamson writes:

At any point in such an interrogation, anything less than a positive answer seems to destabilize the previous positive answer, and therefore all the earlier positive answers in a domino effect. (Williamson, 2005)

Similarly, consider the following reasoning featuring higher-order uncertainty:

INSTABLE REASONING 3
Premise 1: If it is raining, I ought to take an umbrella.
Premise 2: If it is not raining, I absolutely ought to avoid taking an umbrella.
Premise 3: It is raining.
Lemma: Therefore, I ought to take an umbrella.
Premise 4: But I do not know (that I know…) whether I know that it is raining.
Conclusion: Therefore, I ought not to take an umbrella.

In this reasoning, premise 4 destabilizes premise 3 so that the conclusion appears more reasonable than its negation.

It is worth noting that, in itself, this phenomenon of destabilization is not relative to the practical context of the reasoning. In all contexts, a premise such as 4 destabilizes a premise such as 3. That does not mean that the practical context does not play a crucial role in determining whether a premise such as 4 can legitimately be used in the first place. For example, if the cost of error is insignificant and the possibility

of error is farfetched, relying on such a premise in reasoning would certainly be considered unreasonable. But once such a premise is in place, it automatically does its destabilizing work.

Let us take stock. An adequate epistemic condition for the epistemic norm of practical reasoning should be a condition apt to explain:

1. Why it is instable to premise "p" and "it is possible that not p".
2. Why premising "it is possible that not-p" always destabilizes the premise "p" rather than the other way round.
3. Why this phenomenon does not depend on the context of reasoning.

I shall now show that a certainty norm provides an epistemic condition the characteristics of which aptly explain these three features of reasoning.

Consider the following certainty norm:

(ECN) If A is facing a decision that depends on p, then it is epistemically permissible for A to rely on p in practical reasoning iff p is epistemically certain for A. (Beddor, 2020)

It is characteristic of certainty that p is certain (for S) if and only if it is not possible that not-p (for S).[9] If so, this norm easily explains the first feature. According to ECN, it cannot be both epistemically appropriate to premise "p" and "it is possible that not-p". For, on the one hand, if it is epistemically appropriate to premise "p", it is certain that p. If so, "it is possible that not-p" is obviously false. Therefore, "It is possible that p" cannot be an epistemically appropriate premise. On the other hand, if it is epistemically appropriate to premise "It is possible that not-p", it is certain that it is possible that not-p. Hence, "p is not certain" is certain and "p" cannot be an epistemically appropriate premise.

Regarding the second feature — why premising "it is possible that not-p" always destabilizes the premise "p" rather than the other way round —, the explanation is less straightforward. Let us start by noting that it can be true that p and that it is

---

[9] Proponents of epistemic certainty norms typically follow orthodoxy about epistemic modals and characterize epistemic certainty/uncertainty in terms of possibilities (in)compatible with the subject's evidence (see Kratzer, 1981). Henning (2021) has recently argued against such evidentialist norms that they are not transparent, for they "impose constraints on relying on p that are, in some sense, independent of whether p". Indeed, following philosophers such as DeRose (2009), Henning assumes that "epistemic norms should come with a secondary norm of reasonableness" requiring, on some views, to have a justified belief that one satisfies the primary norm. But if the primary norm imposes an evidential constraint (with respect to p) on a decision that is p-dependent (but which is not justified-belief-dependent), the secondary norm will impose a constraint on that decision that will be irrelevant (e. g., having a justified belief that one has a justified belief that p). In contrast, Henning argues, the transparency constraint can be respected if one combines an "Epistemic Must Norm" with a non-factualist account of epistemic modals (see, e.g., Yalcin, 2007). Let me note two things. First, it is rather contentious that the alleged "secondary norm" should be understood as a norm rather than as a mere regulation condition; and it is far from clear that regulation conditions need to be transparent in that sense (see Fassio, 2017 and Vollet, 2022). More importantly, were the advocates of the certainty norm to adopt the non-factualist account of epistemic modals favoured by Henning, it does not look like the arguments presented in the present paper would be weakened. Thanks to a reviewer for raising this issue.

possible that not p. But while "p" cannot be certain if "it is possible that not-p" is true, "it is possible that not-p" can be certain while "p" is true. I suggest that this can explain why "it is possible that p" always destabilizes "p", rather than the other way round. We just have to add the independently plausible assumption that pieces of reasoning are governed by a rule of minimal change (Harman, 1986). Suppose, indeed, that we must make a choice between "p" and "it is possible that not p". What would be the minimal change? Would it be to remove "p" rather than "it is possible that not p"? To remove "p" from the reasoning, while preserving "it is possible that not-p", we don't have to deny that p is true. We just have to accept that p is not an appropriate premise. In contrast, to remove "it is possible that not p", while preserving "p", we have to assume that "p" is certain, and therefore, that "it is possible that not p" is false. This implies a more radical change in our view. For it implies not only removing a premise because of its inappropriateness, it also implies removing it because of its falsity, which involves a change in our beliefs.

Finally, the proposed explanations of the two first features are based on basic characteristics of epistemic certainty. Since, according to the certainty norm, the same epistemic condition (namely, epistemic certainty) is required in all contexts, the third feature is exactly as we should expect: the target phenomenon does not depend on the context of reasoning.

5 Conclusion

There are cases suggesting that knowledge is not always sufficient for epistemically appropriate reliance. While some proponents of the knowledge norm of practical reasoning try to explain away these cases by appealing to the distinction between sufficient and insufficient reason, I have argued that they face a dilemma: either they avoid appealing to chances of error and they misconstrue the counterexamples, or they render mysterious why pieces of reasoning premising both "p" and "it might be that not p" look deeply instable. In contrast, it is clear that these cases do not threaten the certainty norm of practical reasoning. In addition, I have shown that the certainty norm can aptly explain intriguing features of instable pieces of reasoning. These considerations lend strong support to the claim that certainty, rather than knowledge, is the epistemic norm of practical reasoning.

References

Beddor, B. (2020). Certainty in action. *Philosophical Quarterly, 70*(281):711–737.
Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs,* 42, 167–189.
Brown, J. (2011). Fallibilism and the knowledge norm for assertion and practical reasoning. In Jessica Brown & Herman Cappelen (Eds.), *Assertion: New philosophical essays.* Oxford: Oxford University Press.
Crisp, T. M. (2005). Hawthorne on knowledge and practical reasoning. *Analysis,* 65(2), 138–140.
Dancy, J. (2004). *Ethics without principles.* Oxford University Press.

DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context*, Vol. 1. Oxford University Press.

Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.

Fassio, D. (2017). Is there an epistemic norm of practical reasoning? *Philosophical Studies,* 174, 2137–2166

Gao, J. (2019). Against the iterated knowledge account of high-stakes cases. *Episteme,* 16(1), 92–107

Gerken, M. (2011). Warrant and action. *Synthese,* 178, 529–547

Gerken, M., & Petersen, E. N. (2020). Epistemic norms of assertion and action. In Sanford Goldberg (Ed.), *The Oxford handbook of assertion.* Oxford: Oxford University Press.

Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge: MIT Press.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford University Press.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy,* 105, 571–590

Henning, T. (2021). An epistemic modal norm of practical reasoning. *Synthese,* 199(3–4), 6665–6686

Ichikawa, J. J. (2012). Knowledge norms and acting well. *Thought: A Journal of Philosophy,* 1(1), 49–55

Kelp, C., & Simion, M. (2017). Criticism and blame in action and assertion. *Journal of Philosophy,* 114(2), 76–93

Kratzer, A. (1981). The notional category of modality. In H.-J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts: new approaches in word semantics.* DeGruyte: Berlin and New York.

Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research,* 90(1), 72–95

Neta, R. (2009). Treating something as a reason for action. *Noûs,* 43(4), 684–699

Reed, B. (2010). A defense of stable invariantism. *Noûs,* 44(2), 224–244

Skorupski, J. (2010). *The domain of reasons*. Oxford University Press.

Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.

Vollet, J.-H. (2022). Epistemic excuses and the feeling of certainty. *Analysis,* 4, 663–672.

Vollet, J.-H. (2023). Antiluminosity, excuses and the sufficiency of knowledge for rational action. *Erkenntnis*. https://doi.org/10.1007/s10670-023-00682-7

Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly,* 55(219), 213–235

Williamson, T. (forthcoming). Justification, excuses, and sceptical scenarios. In F. Dorsch, & J. Dutant (Eds.), *The New Evil Demon*. Oxford University Press

Yalcin, S. (2007). Epistemic modals. *Mind, 116*, 983–1026.