

An Epistemic Modal Norm of Practical Reasoning

Tim Henning, University of Mainz

In a decision where much hangs on whether p , it would be irresponsible to simply guess that p and proceed. Thus, many people agree that there is an *epistemic norm of practical reasoning*. To justifiably rely on p in practical reasoning, one needs to be positioned well enough, epistemically, to do so.

One much-discussed candidate norm is the Knowledge Norm. Two standard formulations:

“One knows q iff q is an appropriate premise for one’s practical reasoning.” (Williamson 2005, 231)

“Treat the proposition that p as a reason for acting only if you know that p .” (Hawthorne and Stanley 2008, 577)

These formulations differ on a number of counts. The first, unlike the second, is a biconditional. When I speak of the Knowledge Norm here, I mean to refer to a conditional version, according to which knowledge is necessary in order for it to be proper to treat something as a practical reason. The formulations also differ in that one is about treating p as a premise in practical reasoning, the other about treating p as a reason for acting. I will assume that these pretty much come to the same thing. My own preferred formulation, by the way, will be about relying on p in practical reasoning.

The Knowledge Norm is not the only candidate, though. Other proposals abound, like a Justified-Belief-that-One-Knows Norm (Neta 2009), a Justified-True-Belief Norm (Littlejohn 2009 and 2012), a Warrant Norm (Gerken 2011), or a Sensitive-Belief Norm (Enoch, Fisher and Spectre 2012). Also, there is the related debate about the norm of assertion. While it is not clear that both domains require a common norm, some norms from the latter debate might be proposed as practical norms, too. One candidate would be a Safe-Belief Norm (Pritchard 2014). Finally, there is the Truth Norm (defended for assertion in Weiner 2005). While maybe it should not be called an epistemic norm, it may come with a “secondary” norm which is epistemic. (More on this later.)

This paper first criticizes all current proposals (except for the Truth Norm). It claims that they violate a plausible constraint that I label Transparency. I then propose an alternative epistemic norm. This norm captures much of the intuitive plausibility of other candidates, but it avoids the violation of Transparency. In this respect, it resembles the Truth Norm. Unlike the latter, however,

it is an epistemic norm in a narrow sense, and it does not count lucky guesses as proper on any level. The norm proposed contains, not a description of an epistemic condition of the reasoner or her evidence but an *epistemic modal*. It says: Rely on p in practical reasoning only if it must be that p .¹

Roadmap: §§ 1 and 2 describe two problems for existing proposals, the problems of Negative Bootstrapping and of Level Confusions. § 3 offers a general diagnosis, tracing the two problems to a violation of Transparency. § 4 introduces the alternative norm. § 5 discusses Non-Factualism about epistemic modals. § 6 clarifies the proposal. And §7 adds to the proposal an associated norm of “secondary propriety”, which again is a modal norm. §8 briefly discusses reasons which contain epistemic and probability operators. § 9 concludes with another brief look at the role of knowledge.

1. The Problem of Negative Bootstrapping

The first problem concerns those candidate epistemic norms which specify conditions that *entail belief* in the proposition at issue (or, indeed, belief in related propositions, see below). Arguably, these candidates include the Knowledge Norm, the Justified-Belief-One-Knows Norm, the Justified-True-Belief Norm, the Sensitive-Belief Norm, etc. These all entail that an irrational failure to believe can affect the balance of reasons that we can properly rely on, in ways that seem absurd.

Suppose my evidence strongly suggests that a good friend is in trouble, in which case I should go and look after him. Doing so would cause me some inconvenience, but nothing too serious. Suppose also that I cannot bear the thought that my friend is in trouble, so I am in denial. Despite the evidence I irrationally cling to the belief that all is well. Now, presumably, since I fail to believe that my friend is in trouble, I do not know it (or justifiably believe it, or believe I know it) either. So according to the above norms, I am not permitted to treat this proposition as a reason. We can imagine that the same goes for related propositions, such as that it is likely that my friend is in trouble, or even that there is a non-trivial chance that he is in trouble. These are all thoughts I could not bear, so I fail to believe them either. Again, the above norms entail that I am not in a position to treat any of these considerations as reasons. In consequence, it may well be that the balance of

¹ I will take the liberty to sometimes use letters like p and q in the grammatical role of singular terms for propositions (as in “rely on p ”) and to sometimes use them in the grammatical role of sentences (“it must be that p ”). I trust that no serious confusion will be caused by this, and I apologize to those whose aesthetic sense is offended.

those reasons that I *can* take into account tells against going to see my friend.² Thus I affect what reasons I should take into account, and so what my reasons support, by an *irrational failure to believe*.

This is a problem, however. Here is why. It is uncontroversial, I take it, that my failure to believe in accord with my strong evidence is an epistemic irrationality. But I also think that my decision not to look after my friend is a *practical* failure. (My evidence suggests that he is in severe trouble, and all I do is to try to ignore it as best I can!) Alas, the candidate epistemic norms make it hard to vindicate this latter judgment. Again, due to my failure to believe, I am allegedly not in a position to treat the proposition that my friend is in trouble (or the proposition that it is likely that he is, etc.) as a reason. And without such reasons to rely on, there may simply be nothing left that could direct me to the conclusion that I am to go and look after my friend. So how am I at fault?

This is a mirror image of what has come to be known as the bootstrapping problem in the debate on requirements of rationality. Many authors observe that a requirement of *enkrasia* seems to entail that by believing that there is reason for one to do something, one makes it the case that there is reason for one to do it (see Broome 1999, Kolodny 2005). Here, we find a similar absurdity: Just as one should not be able to create reasons by believing in them, one should not be able to get rid of them by irrationally failing to believe in them! This is the Problem of Negative Bootstrapping.

A word about the structure of the argument:³ The basic intuition is that all things considered, I *ought* to rely and act on the reason that my friend is in trouble. But epistemic norms of reasoning formulate permissions to rely on reasons, not requirements. So how does the argument go? It uses a principle that seems plausible: If a consideration is not epistemically accessible to a subject (in the sense that it is not epistemically permissible for her to rely on it), it cannot be true that she still ought to rely on it, all things considered. This principle is then used in a *modus tollens*. Again, it seems wrong for me not to rely on the reason that my friend is in trouble. Hence it cannot be true that this fact is not epistemically accessible. And therefore, norms that claim otherwise are problematic.

Let me repeat that this problem is not at all specific to the Knowledge Norm. It arises for many weaker proposals as well – in effect, for *any* norm that includes the component of belief. Notice that such a belief component can be an implicit part of some norms. For instance, consider norms that merely demand that one have a particular amount of evidence. On most accounts, *having* evidence requires certain psychological states or responses as well, whether it be knowledge (as in Williamson 2000) or something else. And these states, again, seem to entail belief. So once again this opens up the possibility of getting rid of practical reasons by irrational failures to believe.

² Of course, the proposition *that my friend is okay* is not among these reasons either. The idea is that I have other reasons to do other things, which I may rely on, and lack permission to rely on any particular views about my friend.

³ I thank an anonymous reviewer for prompting this clarification.

This discussion bears some similarities to J. Lackey's (2007) critical discussion of the knowledge norm of assertion. Lackey also argues that an unwarranted failure to believe should not make it any less proper to assert certain things. She holds that in the cases she presents, subjects ought to make certain assertions despite lack of belief in them – “selfless assertions,” as she calls them. Echoing Lackey, my point is this: In the above example, I should engage in selfless reasoning.

Maybe there are quick fixes? Maybe knowledge does not require belief (Radford 1966)? Or maybe I do believe that my friend is in trouble in the relevant sense (Gendler 2008)? For one thing, these would be surprisingly controversial commitments, attached to prominent theories of epistemic norms. For another, they would not take care of a second problem, to be discussed now. In fact, as I will argue, these ‘fixes’ for the first problem may render the second one more pressing.

A different kind of fix would, of course, be to turn to versions of an epistemic norm which are de-psychologized, i.e. which do not require any actual responses on the part of the reasoner. Examples might be norms that require a certain amount of available evidence (whether or not one actually possesses or responds to it) or some such. But these too fall prey to the second problem.

2. The Problem of Level Confusion

The second problem is this: Typical candidates for epistemic norms make us vulnerable to “philosophical” doubts, doubts that are intuitively irrelevant. To see this, we must introduce the distinction of “primary” and “secondary propriety” (to use the terminology of DeRose 2002, 180).

To begin, none of the conditions specified by typical epistemic norms are *luminous* (in the sense of Williamson 2000). None of them specify conditions such that, whenever a subject is in that condition, she is in a position to know that this is so. This raises questions: What about subjects who satisfy the condition specified by the norm but who faultlessly fail to realize that this is so? Conversely, what if a subject fails to satisfy the condition but reasonably takes herself to satisfy it?

Many authors (e.g., DeRose 2009, Hawthorne and Stanley 2008, Williamson 2000) introduce a “secondary” or “derived” norm to deal with these questions. This secondary norm specifies what we may expect from a subject who tries to guide herself by the primary norm, under what conditions such a subject can be excused for breaching the primary norm, when she deserves blame, etc. I propose the following formulation, using Williamson's (2000) term “reasonableness”:

Reasonableness

If a norm says that one ought to rely on *that p* only if condition *c* holds, one is reasonable (with regard to that norm) in relying on *that p* iff one has justification to believe that *c* holds.⁴

So if we take the Knowledge Norm, the idea would be that a subject who knows *p* may satisfy the Knowledge Norm but still be unreasonable in relying on *p*, because she is not justified in believing she knows *p*. In this case, she would deserve a certain kind of criticism or blame, even though she was correct to rely on *p*. Conversely, someone may be justified in believing she knows that *p* even though in fact she does not. In this case, she is not permitted to rely on *p* as a reason, and if she does, she needs an excuse. But there will be an excuse: she was justified in believing she knew *p*.⁵

I agree that epistemic norms should come with a secondary norm of reasonableness. Some authors (Lackey 2007, Kvanvig 2009) oppose the idea, claiming that there should be only *one* norm governing all of our responses. I discuss some of their objections below. As I will argue, these authors, too, have a point. Indeed, one salient feature of the epistemic norm I will propose is that the distinction collapses in certain cases (especially in first-person contexts) but not in other cases.

Now, imagine a variation of our case in which my response to my evidence about my friend is impeccable. I do believe he is in trouble, as I should. But a philosopher approaches me and asks: “I see that you are sure that your friend is in trouble. And I do not mean to doubt that he is. But are you sure *your state* satisfies the concept of *knowledge*?” In reply, I do not just need to produce evidence about my friend. I need higher-order evidence *about* it, showing that I am justified in

⁴ Many proponents of secondary norms (like DeRose 2009, Weiner 2005) require justified (or “reasonable”) belief. But this creates some unclarity. In their presentations, one is reasonable if one justifiably believes one satisfies the primary norm, and one is not reasonable if one believes one does *not* satisfy it. This, however, raises the question of how to deal with cases in which one lacks either belief. I think the most plausible way to generalize the norm to cover these cases makes justification central. A further advantage of this version of a secondary norm is that it makes fewer demands regarding the presence of higher order beliefs.

⁵ In what follows, I assume that a secondary norm is tied closely to our ordinary understanding of when agents are reasonable, or deserve blame etc. After all, the purpose of DeRose’s distinction and similar proposals is precisely to accommodate our ordinary intuitions about when agents do not assert or reason with full propriety. Admittedly, there has been some debate about the nature of secondary propriety, in particular concerning whether it is a technical notion or intended to capture intuitive judgments. As I see it, this debate arises from problems regarding the Knowledge Norm in particular. Many authors have argued that subjects who have a true, justified, but Gettiered belief that *p* do not need an excuse when they rely on *p* or assert *p* (see Douven 2006; Brown 2008, Kvanvig 2009, Schechter 2017). So whatever secondary propriety amounts to in these cases, it cannot be to provide an excuse (though see Littlejohn 2012). I think that this is a reason to doubt the Knowledge Norm, and not a reason to employ a technical and revisionary notion of secondary propriety. Thanks to an anonymous reviewer for alerting me to this issue.

believing that my first-order evidence is good evidence. Also, I will have to defend claims about *knowledge* – maybe that fallibilism is true and that brain-in-a-vat scenarios may properly be ignored.

When encountering these calls for justification, I may not waver in my first-order conviction that my friend is in trouble. But I may doubt whether my position vis-à-vis this claim is knowledge. In any case, the conclusion may be that I lack the justification to believe that I have knowledge.

Now, assuming that the epistemic norm comes with a secondary norm like Reasonableness, this should *directly affect* my willingness to rely on the proposition that my friend is in trouble. After all, to be *reasonable* in relying on the proposition *that he is in trouble*, I not only have to be in a certain epistemic position with regard to *this* proposition; I *also* need the justification to believe the higher-order proposition *that I am in the right epistemic position with regard to the proposition that my friend is in trouble*. Since epistemic positions are not luminous, these two conditions can come apart. Hence I might find myself staying home, *not* because of any uncertainty regarding my friend's need but because, say, I lack a compelling defense of fallibilism. According to the norms, this would be reasonable: I would be excused for not going to support my friend; I might deserve blame if I did.

In a classic paper (Alston 1980), W. Alston has observed that certain positions and arguments in epistemology involve or invite what he calls *level confusions*. The latter occur when we should be debating whether *p* but instead wind up debating a higher-level question, such as whether somebody knows that *p* or has evidence that *p*, etc. Given non-luminosity, propositions on different levels require different kinds and amounts of justifying evidence (importantly, even in the first-person case), hence level confusions can be harmful. My point is that given the available proposals of epistemic practical norms, reasonableness will make us prone to epistemic level confusions.

Hence I call this the Problem of Level Confusions. Let me return to a point I made earlier. In response to the Problem of Negative Bootstrapping, I said that controversial assumptions about knowledge or belief might fix it. But I added that these fixes, in addition to their controversial nature, may render the present problem (i.e. the Problem of Level Confusions) worse. We can now see why: If the epistemic norm requires a controversial account of core epistemic concepts, this is all the more food for “philosophical” doubt. For what reasonableness will require now is not just good higher-order evidence, an argument for fallibilism, etc. – but *also* a defense of a controversial account of knowledge (like Radford's) or of a controversial account of belief (like Gendler's), etc.

So return to the previous example, in which I irrationally refuse to believe that my friend is in trouble. Of course, adopting Gendler's or Radford's proposals would make it easier to defend the correct conclusion that I should go to the aid of my friend. In what sense might it make the

present problem, i.e. the Problem of Level Confusions, worse?⁶ It would turn the example into another case where a practical issue turns on intuitively irrelevant questions of philosophical doctrine. Our theory should not predict that in the decision whether to go to my friend, there is a relevant and legitimate doubt that requires an account like Radford's or Gendler's to be alleviated.

I also mentioned earlier that, in response to the Problem of Negative Bootstrapping one may turn towards de-psychologized constraints. E.g., one might say that the correct epistemic norm is really about evidence that is available (although not necessarily *had*, see above). It might say that one must treat p as a practical reason only if knowledge-level evidence for p is available – i.e., only if one is in a position to know p , whatever one's actual state is. Now, the problem with these proposals is that the elimination of belief does not eliminate the epistemological sophistication. And hence the Problem of Level Confusions persists. After all, reasonableness in acting on this norm will now require justified belief about which evidence is available, about what it means for evidence to be available, about when available evidence is knowledge-level, etc. Interesting as these questions may be, it is simply implausible that I am reasonable in staying home, with my conviction that my friend is in trouble fully and justifiably intact, just because I lack arguments in epistemology.

As I said at the beginning of this section, some authors eschew the distinction of primary and secondary propriety. Lackey, for example, claims that the distinction is “spurious.” On her view, if I do not satisfy an epistemic norm (of assertion, in her case) but reasonably believe that I do, I am blameless and have an excuse – but that does not have to mean that there must be another norm according to which the assertion was in some sense a proper assertion after all. With this, I disagree. Following Kvanvig (2010), I think that norms of this kind (i.e., norms of assertion and of practical reasoning) play one of their most important roles in the “egocentric predicament” of deciding what to say, and what to treat as a reason, oneself. Here, they must provide *guidance*. But if we cannot generally expect subjects to grasp whether they satisfy whatever epistemic norm we are proposing, we have to allow that these subjects guide themselves by other rules in answering these questions. And if these other rules play this legitimate guiding function, surely they deserve to be called norms of assertion and practical reasoning at least as much as the norms we propose.

Kvanvig, too, rejects the primary/secondary norm distinction. As indicated, he invokes the “egocentric predicament”, claiming that it does not make sense to guide oneself by both primary and secondary norms at once. His proposal is a norm that is, in a sense, luminous after all. He proposes a norm of assertion according to which asserting that p is epistemically proper only if the speaker's perspective provides *epistemic justification* for p . Epistemic justification is defined as the

⁶ Thanks to an anonymous reviewer for this question.

kind of justification required to turn true, non-Gettiered belief into knowledge. And Kvanvig claims that it obeys an iteration principle to the effect that *that p* is epistemically justified for a subject only if *that p is epistemically justified for that subject* is epistemically justified for her, and so on.

Unfortunately, this account runs into the same problem. In Kvanvig's view, a subject has epistemic justification for *p* only if her evidence for *p* and the rest of her total knowledge confirm, not just that *p* is justified to some high degree but that the justification enjoyed by *p* is *epistemic*, i.e., knowledge-level.⁷ But the question of when justification is knowledge-level is philosophically contested. Kvanvig, e.g., denies fallibilism and therefore argues that the subject's total evidence must be sufficient to confirm that all further inquiry will support *p*. He also claims that such epistemic justification is accessible in the strong sense that there can be no Williamson-style anti-luminosity argument against it. Accordingly, we will be justified in believing that we have epistemic justification *only if we are justified in a particular conception of knowledge-level justification*. So we will again be open to philosophical doubts, although in this case it is due to the suggested primary norm.⁸

I tentatively conclude that this way of avoiding the Problem of Level Confusions is blocked.

3. A General Diagnosis: Lack of Transparency

Let me attempt a general diagnosis: As an agent, my ultimate practical aim must be to go or stay depending on whether or not my friend is in trouble. Therefore, I will be willing to guide myself by an epistemic norm only if I am confident that it accords with this aim. My ambition to go and see my friend iff he is in trouble and my ambition to conform to epistemic norms should not *conflict*.

This, I think, is the common core of the Problems of Negative Bootstrapping and of Level Confusions. In the former case, the fact that I fail to form the relevant belief does not do anything to make it more or less likely that my friend is in trouble. In the latter case, the truth of fallibilism also seems independent of the question of whether my friend is in trouble. (At any rate, if fallibilism is supposed to be false, that should shake my confidence that I know that he is in trouble to a far greater extent than it shakes my confidence that he is in fact in trouble.) Therefore, the epistemic norms impose conditions that seem like irrelevant obstacles from the point of view of what matters.

⁷ Kvanvig insists on this because he thinks that his norm should provide an explanation of Moorean sentences like "P, but I do not know that P" (which figure prominently in defenses of the knowledge norm, see Williamson 2000).

⁸ Thanks to an anonymous reviewer for urging me to elaborate on the objection to Kvanvig.

A related point is pressed by Enoch, Fisher and Spectre in work on epistemic norms as standards of proof in criminal law. They discuss the proposal that statistical evidence may be inadmissible because it does not ground knowledge, and knowledge is the relevant epistemic norm:

“But why should the law of evidence care about knowledge or about epistemology more generally? [...] Let us emphasize that to insist that the law should after all care about knowledge is (pretty much) to be willing to pay a price in accuracy. Indeed, excluding statistical evidence amounts to excluding (what is often) good, genuinely probative evidence. And this means that the legal value of knowledge – if it has legal value, and if that value is what grounds the differential treatment of statistical and individual evidence – sometimes outweighs the value of accuracy, that, in other words, in order to make sure that courts base their ruling on knowledge, we are willing to tolerate more mistakes than we otherwise would have to, and indeed a higher probability of mistake on this or that specific case. This just seems utterly implausible.” (Enoch, Fisher, and Spectre 2012, 16)

This is not a paper on the philosophy of the criminal law, of course. But the worry I mean to press in this section is similar to the worry expressed in this quote. (Note that the authors make it clear that the issue is not just about “the long run”, but also about single cases.) Let me be more precise.

Take Hawthorne and Stanley’s definition of *p*-dependence:

“Let us say that a choice between options $x_1 \dots x_n$ is *p* dependent iff the most preferable of $x_1 \dots x_n$ conditional on the proposition that *p* is not the same as the most preferable of $x_1 \dots x_n$ conditional on the proposition that not-*p*.” (2008, 578)

Assuming that rational agents aim to minimize opportunity costs, my ambition in a *p*-dependent choice must therefore be to perform the act that is best, given that *p*, iff *p*. More officially, it is this:

Practical Aim: Perform the act that is best, given that *p*, iff *p*.

Now we may further assume that in the cases at hand, sound reasoning will lead the agent to perform the act that is best, given that *p*, iff the agent relies on *p*. (This would mean that no weaker premise would suffice to lead her to perform the relevant act.) Then my aim will also include this:

Accuracy: Rely on *p* iff *p*.

Clearly, this coincides with the Truth Norm. Thus, from the point of view of our practical aims, we must wish that in fact, whether we know it or not, we conform to the Truth Norm. Does this mean that the Truth Norm is correct? This does not follow. Compare: In discussions of practical rationality, many people deny that the correct norm tells us: “Do what has the best consequences!”,

even though this is in some sense our aim. Instead, many think that the correct norm tells us to do what is *likely* to lead to the best consequences, given our evidence or rational credence. Similarly, although our aim is to rely on p iff p , the correct norm may tell us to rely on p only if this is most likely to lead to the best consequences, given our evidence. In other words, the correct primary norm for relying on p may require that our evidence support p , or something the like. And like many others, I do think that some such proposal better captures our intuitions about relying on p .

I think Practical Aim and Accuracy suggest a condition of adequacy for epistemic norms. It is this: Epistemic norms of practical reasoning should not impose constraints on relying on p that are, in some sense, independent of whether p . They should not lead us to ask questions that do not pertain to the likelihood of p , lest they impose what we must see as irrelevant and harmful barriers. In the examples, this is what happens if I do not go to support my friend because I am uncertain of some epistemological claim, or because I recognize an irrational failure to believe on my part.

This can lead to the following condition of adequacy:

Transparency

A norm N accords with my aim (to rely on p iff p), only if according to N ,

- (i) any consideration that counts against relying on p also counts against p being true, and does the former because it does the latter, and
- (ii) any consideration that counts in favor of relying on p also counts in favor of p being true, and does the former because it does the latter.

If this holds, I will say that according to norm N , the question of whether to rely on p is transparent to the question of whether p . I will also call the epistemic norm N itself transparent in this case.

There are other ways to formulate the transparency idea. Here is one in terms of credence:

Transparency, Probabilistic Version:

Let N be an epistemic norm for relying on p in practical reasoning, and let P be the probability measure that gives the reasoner's rational credence. Let e be the condition that N imposes on relying on p . Then N is transparent iff for any consideration e , differences between $P(e \mid e)$ and $P(e)$ are explained by appeal to differences between $P(p \mid e)$ and $P(p)$.

In other words: Whatever (dis-)confirms *that p can be relied on* must do so *by* (dis-)confirming *that p* . I give this alternative formulation to make it clear that I am concerned with the general idea rather

than with the details of a particular formulation. (In fact, it is debatable whether my own proposed norm satisfies the Probabilistic Version. On some semantic views it does, on others it does not.)

Transparency is supported by the above examples. If I realize that I fail to believe what my evidence supports, this has strong impact on my credence in *that I know* without necessarily affecting my credence in *that my friend is in trouble*. Likewise, higher-order evidence and controversial epistemological hypotheses affect my credence in propositions about my epistemic state in a way that is stronger, and not fully explained by, their impact on my credence in *that my friend is in trouble*.

My presentation ties the condition Transparency to practical aims. It is only if a norm is transparent that we can apply it while keeping our eyes firmly on the issues that must matter to us as rational agents. Only a norm that is transparent will tell me: “Look, it does not matter whether you really manage to bring yourself to *believe* that your friend is in trouble. All that matters is whether he is. So you should use your evidence to establish, as well as you can, whether it is true that he is in trouble. Once this is settled, it does not matter whether you suffer from doxastic incontinence, or whether your evidence *also* allows you to establish higher-order claims about your evidence, etc.”

Why do I claim that epistemic norms should accord with practical aims? There are two things one might say here. One thing would be that the normative status of epistemic norms is ultimately instrumental. Epistemic norms concerning practical reasoning serve the aim of *conforming* to our practical reasons. If this idea is adopted, a norm that gets in the way of conforming to our reasons would not have any weight at all. Another thing one could say is this: Even if epistemic norms are not just instrumentally but intrinsically important, it would still be intolerable if we had to decide whether we aim for conformity with our practical reasons or for conformity with epistemic norms.

But isn't the only norm that is transparent in my sense the Truth Norm? No. Let me explain.

4. Introducing EMN

The Knowledge Norm derives much intuitive support from ordinary discourse. Criticism in terms of knowledge, such as: “If you did not know whether p , how could you act as if p ?” is very natural, and frequently encountered. But close attention to ordinary discourse reveals that there is a different way to formulate epistemic criticism, which is also very familiar. If someone is about to rely on p , a natural way to challenge her is to say that it might be that not p . E.g., if I go to the pub to meet a friend without having checked whether she will be there, it is natural to object by saying: “You should not just go there without calling. She might not be there!” My proposal takes this as a clue. I suggest an epistemic norm that demands that it is not the case that it might be that not- p :

Epistemic Must Norm (EMN): Rely on p only if it must be that p .

Of course, there is an orthodox semantic analysis of epistemic operators on which EMN does not really differ from the Knowledge Norm. According to this analysis, epistemic “must” is used to say that the content it embeds is entailed by what is known – and thus, according to typical assumptions about closure, itself known. This, I suspect, explains why the possibility of epistemic criticism in modal terms has not, so far, been considered worthy of investigation in its own right.

This conclusion would be premature. As I argue in the next section, other analyses of “might” and “must” are available. If correct, they can explain why my proposal is an improvement.

But first, let me be very clear about what the challenge is. Compare the following conditions:

Accuracy:	Rely on p	(if and) only if	p .
Knowledge Norm:	Rely on p	only if	you know that p .
Justified Belief Norm:	Rely on p	only if	you justifiably believe that p .
Sensitive Belief Norm:	Rely on p	only if	your belief that p is sensitive.
...			
EMN:	Rely on p	only if	it must be that p .

Evidently, the right-hand sides of the candidate epistemic norms impose additional constraints on relying on p , as compared to the right-hand side of Accuracy. And of course, this is not a bug but a feature. As I said, there seems to be something objectionable about acting on lucky guesses where much is at stake. But the problems I raised for candidates like the Knowledge Norm, the Justified Belief Norm, etc. was that the additional constraints imposed are *not transparent* – i.e., that they can be confirmed or disconfirmed by considerations that do not, in the same way, confirm or disconfirm what we should see as the only practically relevant criterion: the truth of the premise p .

My claim is that EMN does better in this regard. The right-hand side of EMN, while not logically equivalent to the right-hand side of Accuracy, does satisfy Transparency. How can this be? Well, to see how EMN satisfies Transparency, we need to know what “must” adds to the condition on the right-hand side. Clearly, there had better be *some* sense in which it does not add very much. And as it happens, this is exactly what much recent work on epistemic modals suggests.

5. Non-Factualism about Epistemic Modals

Recent work on epistemic modals has converged on claims that support my proposal. A sample:

“If a modal μ is used in an epistemic sense and combined with a sentence s , the resulting μ - s does not differ from s in content.” (Schnieder 2010, 603)

“When I say, ‘Bob might be in his office,’ I am talking about Bob and his office, not myself or the extent of my information.” (Yablo 2011, 271)

“[T]he question of whether $\diamond\phi$ is ‘transparent’, as it were, to the question of whether ϕ . [...] Believing that ϕ and believing that $\diamond\phi$ are states of mind supported by reasons of the same category.” (Yalcin 2011, 308)

However, to vindicate claims like these, the authors must part ways with the orthodox treatment of modals (mentioned earlier) – the treatment given, e.g., in Kratzer’s (1981) influential work. According to this view, $\Box p$, as uttered in a context c , is true in a world w iff for all worlds w' that are epistemically accessible from w (in a way determined by c), p holds true. To specify it formally:

$$\llbracket \Box p \rrbracket^{c,w} = 1 \text{ iff } \forall w' (wR_c w' \rightarrow w' \in p)$$

Similarly:

$$\llbracket \Diamond p \rrbracket^{c,w} = 1 \text{ iff } \exists w' (wR_c w' \wedge w' \in p)$$

On this view, a statement like “it must be that p ” makes a claim about an epistemic state of affairs. (Notice that the right-hand side is essentially equivalent to the analysis of knowledge ascriptions in Hintikka 1962.) It says that in w , what some relevant subject knows entails that p . Accordingly, this statement is true or false relative to a world depending on how things stand knowledge-wise (or evidence-wise) in that world. It does not directly concern how things stand p -wise in that world.

What alternatives could one offer to capture the fact that the subject matter of $\Box p$ and $\Diamond p$ is p , as the quotations above suggest? There is a variety of approaches. I merely mention one major

proposal, namely Yalcin’s (2007) so-called *domain semantics*.⁹ It differs from the orthodox account in that sentences are interpreted relative to not two but three parameters – a context c , world w , and information state s (where an information state is a set of worlds modelling a body of information). The semantics for epistemic “must” and “may” then takes the following form:

$$\llbracket \Box p \rrbracket^{c,s,w} = 1 \text{ iff } \forall w' \in s : \llbracket p \rrbracket^{c,s,w'} = 1$$

and

$$\llbracket \Diamond p \rrbracket^{c,s,w} = 1 \text{ iff } \exists w' \in s : \llbracket p \rrbracket^{c,s,w'} = 1.$$

The important thing to note is that the world w makes no appearance on the right-hand sides. That is: How things stand in w – knowledge-, evidence- or otherwise (!) – does not matter for the truth of $\Box p$ or $\Diamond p$ in w . So these modals are quantifiers over a domain of worlds which is *not* determined by facts in the evaluation-world, via any notion of accessibility. Instead, Yalcin treats s as an autonomous parameter, which is fixed independently of the w -parameter. That is, given some state of information, *all* worlds w will be worlds with regard to which $\Box p$ or $\Diamond p$ have the same value.

What does this mean? It means that there are no $\Box p$ -facts or $\Box p$ -worlds, nor $\Diamond p$ -facts or $\Diamond p$ -worlds, in the sense in which there are p -facts or p -worlds. In other words, there is no particular way for an individual world to be such that it, specifically, renders $\Box p$ or $\Diamond p$ true. In *still* other words, this means that the contents $\Box p$ and $\Diamond p$ are not propositions (at least in any usual sense).

What do these statements express, then? Yalcin adopts an idea from Veltman (1996): They perform a kind of test on a given information state. They deliver 1 if we find that the state is such that it contains only (or, for \Diamond , some) p -worlds and 0 otherwise. We might, accordingly, treat them as expressing properties or characteristic functions of information states, not of individual worlds.

How does this proposal explain the above intuitions about the subject matter of p and $\Box p$? To see this, consider the process of inquiry (modeled along the lines of Stalnaker 1984). When we try to establish whether p , we check whether there are any possible non- p worlds which, for all we can tell, may be actual. We sort through the possibilities we cannot rule out, with an eye to whether p is true in them. On Yalcin’s account, the same thing should happen when we inquire whether to accept a claim that $\Box p$. We inquire whether our information state has the property of having only p -elements. This inquiry is – exactly: a sorting through live possibilities with an eye to whether p ! The reason for this, again, is that on Yalcin’s view, $\Box p$ is not a claim *about* some information state.

⁹ As Yalcin notes, his account is indebted in many respects to the update semantics of Veltman (1996) (see also Willer 2013), which could have served as an example equally well.

So we do not have to evaluate claims about what our information state with regard to p is. Our information state is not part of the content but part of the index relative to which the content is evaluated. This is why the processes of inquiring whether $\Box p$ and of inquiring whether p coincide.

Contrast: According to the earlier orthodoxy, something quite different happens. We inquire whether there are any possible worlds which, for all we can tell, may be actual in which *it is not known that p* , or in which *the relevant evidence does not establish p* . So we go through the live possibilities with a very different set of criteria, and a very different question in mind. We sort through the worlds with an eye to what the facts about someone’s epistemic position in them are. (In a framework of accessibility relations, this means that we check the worlds in our information state with an eye to whether certain *other* worlds are p -worlds – namely those that are epistemically accessible from those former worlds. This is a different endeavor with possibly different results.)

In this way, Yalcin’s domain semantics can account for the intuitions voiced in the quotes above. However, some authors have objected that this may be too much of a success, so to speak. This is because logical systems that develop the basic intuitions of a domain semantics (as well as an update semantics) validate a very strong principle, according to which p logically *entails* $\Box p$. Many authors have criticized this feature. Schulz (2010a), for example, observes that it often seems to be rationally permissible to be rather confident that p while also being virtually certain that it is not the case that $\Box p$. Given that logical entailment should preserve rational certainty, this is a problem.

I am not entirely convinced that this problem is unsurmountable, in part because I think that epistemic “must” has uses that do not have quite the strength that the objection ascribes to it (see below). I cannot adjudicate between competing semantic accounts of epistemic modals here. I merely point out that critics like Schulz (2010b) have proposed alternative semantic theories which likewise preserve the non-factualist intuitions, and which thus support my proposal equally well.

On all these semantic views EMN satisfies Transparency in some version. They agree that any evidence that bears on whether $\Box p$ will have to be evidence that bears on whether p . Of course, if there is no proposition that $\Box p$, we should not say that evidence for p is *evidence for $\Box p$* . What we should say instead is that any consideration can have an epistemic impact on our conviction that $\Box p$ only in virtue of providing evidence for p .¹⁰ Also, on the domain semantics, $\Box p$ will not express a content for which a probability function is defined, so we cannot use the Probabilistic Version (as I said). If we adopt Schulz’s view, however, we can. On his view, we can meaningfully speak of the probability of $\Box p$. And indeed a consideration can affect the probability of $\Box p$ only by virtue of affecting the probability of p . In any case, the idea behind Transparency will surely be satisfied.

¹⁰ Thanks to an anonymous reviewer for prompting this clarificatory remark.

On all views, inquiring whether Bob must be in his office is not a different project from inquiring whether he is in his office. The same evidence matters in either case, and in the same way. And in neither case do we pay special attention to evidence about our evidence, or to theories of knowledge, etc. (Or, to put it more carefully: Cases in which EMN will require us to consider such higher-level questions will be precisely those cases, if any, in which we need to consider these questions even to determine the first-order question whether Bob is in his office. Consequentially, there will be no asymmetry in how higher levels matter to the inquiries whether p and whether $\Box p$.)

This, in a nutshell, is what makes it possible for the epistemic modal norm to be transparent, as defined above. Although the condition it imposes is not the same as in the case of the Truth Norm, the processes of inquiring whether the Epistemic Must Norm is satisfied and the process of inquiring whether the Truth Norm is satisfied will coincide, and so respond to the same reasons.

My proposed EMN is to be understood as using “must” (or \Box) in the sense that the semantic proposals just outlined aim to capture. Although these proposals differ in important respects, they also agree on certain basic questions – especially on those that relate to Transparency. Now, I do think that some such proposal will turn out to be more adequate as an account of natural language epistemic modals than the orthodox one – but, of course, I cannot argue for this here. At any rate, the norm I mean to suggest should be understood as having an operator of the kind just discussed.

6. Some Clarifications

Let me further characterize the EMN (which, again, tells us to rely on p only if $\Box p$) by contrasting it with other norms. One first important point is that the EMN differs from the Truth Norm. It is perfectly possible that p but that not $\Box p$, in which case it is not correct to rely on p . Note well that this is not a matter of secondary propriety, i.e. reasonableness. It is not that we are excused for not relying on p . It would not have been correct to do so – we would have needed an excuse for *that*.

As I have indicated earlier in this paper, I do consider this an important advantage of EMN over the Truth Norm. The crucial test case, to my mind, is the case of a lucky guess. As far as my intuitions go, for a subject to rely on a lucky guess in an important decision is not proper on any level. Defenders of the Truth Norm, however, must claim that the impropriety is merely of the secondary kind – that the subject deserves blame, but that she has still, unknowingly, reasoned as she should have. EMN yields the more plausible verdict that she *should not* have reasoned this way.

What about the Knowledge Norm? If $\Box p$ is true on our information, does that mean we know that p ? Not necessarily. In the framework of a domain semantics (or a dynamic semantics),

it all depends on the exact view we take on information states. It is, in fact, possible to stipulate that a subject's information state consists of what she knows. In this case (and assuming that knowledge is closed under entailment), $\Box p$ will indeed be true, relative to a subject's information, iff that subject knows that p . But even if this should be our choice, it is crucial to bear in mind that $\Box p$ still does not *state* or *report* the fact that the subject knows that p . As I keep insisting, $\Box p$ expresses a condition that, unlike a knowledge claim, is transparent to p (in the sense explained).

More importantly, however, we do not have to make this assumption about information states. In fact, in a basic domain or update semantics it is not even guaranteed that $\Box p$ is *factive*. If we wanted to guarantee this, we would have to formulate an additional constraint on information states. Using the language of the orthodox framework of accessibility relations, this condition is often called *Reflexivity*. Reflexivity says that for all information states s , the actual world is among the elements of s . (This, of course, amounts to a factive understanding of information: If we think of s as the intersection of all the propositions that constitute a body of information, Reflexivity amounts to the requirement that all of these propositions have to be true in the actual world.)

Should we embrace Reflexivity and thus the factivity of $\Box p$? This is a controversial question. As far as intuitions about the meaning of epistemic “must” in natural languages are concerned, opinions are sharply divided. Traditionally, many authors in philosophy seem to be drawn to a strong, factive reading of “must.” By contrast, many authors in linguistics insist that “must p ” is often *weaker* than “ p ”. Especially, they put emphasis on the fact that we seem to use “must” precisely in contexts where there seems to be some residual uncertainty. To take an example (from Mihoc et al. 2019): Suppose Jo is looking for Anna. She has just confirmed that it is 7 pm, and she knows that at 7 pm Anna is *typically* home. Here, Jo can perfectly well say: “Anna must be home now,” although she cannot flat-out assert that Anna *is* home now. The epistemic modal claim seems to quite clearly allow for some possibility of error. Consequently, Kratzer's (1981) influential account models epistemic modals as quantifiers over a domain that is defined, not just by the propositions that one knows (in the so-called modal base) but in addition by best fit with propositions about what is normally the case (in a stereotypical ordering source). There is no reason why update and domain semantics cannot introduce analogous restrictions on information states.

I think as far as natural language is concerned, the jury is still out on whether “must” is factive or not.¹¹ But I stipulate that whether or not it is ultimately the best way to understand “must” in English, in the EMN “must” or \Box should be stipulated to be non-factive. My main motivation for this proposal is simply that I find much to recommend in non-factive candidates for epistemic

¹¹ For a good overview of the debate and some experimental evidence supporting Kratzer's restricted quantification account, see Del Pinal and Wadlon (2019).

norms of practical reasoning. As many authors observe, where a subject has excellent evidence for p , it often seems that relying on p is not just excusable but the thing to do, even if p is false. So I opt for an understanding of EMN that is non-factive.¹² This makes for a plausible epistemic norm.¹³

EMN will call for further choices. An especially important one concerns the selection of the relevant information state. When we ask whether a subject satisfies EMN, do we evaluate the relevant modal claim relative to *her* information, or to *ours*, or some third thing? I will not commit to a specific proposal here. But I do think that an attractive solution will be to let context determine the relevant state. So it may be the case that in the context of the subject, it is correct to judge that she ought not to rely on p , while in the context of better informed eavesdroppers it is correct to judge that it would be proper for the subject to rely on p . Of course, there is a danger in this proposal. As MacFarlane (2014, 103 f) points out in a related context, no agent can be expected to guide herself by so many divergent norms at once. But this danger can be averted (as MacFarlane also notes) if we privilege certain contexts and information states for certain purposes. I will here do so by privileging the *agent's* information in the *secondary* norm, to be discussed below. The consequence is this: An agent may often be well aware that different points of view may differ on whether she should rely on p . This is not actually implausible, I think, as long as the agent can tell what is legitimately expected of her, what she will be blamed and held accountable for, etc. This will be fixed by the secondary norm, which requires that she guide herself by her own information.

7. EMN and a Modal Norm of Secondary Propriety

My proposal is not yet complete. I have not ensured that EMN will escape the second problem mentioned above – i.e., the Problem of Level Confusions. After all: If $\Box p$ semantically expresses a property of our information state, will Reasonableness not require higher-order evidence about properties of our information state? And will that not re-introduce the “philosophical” worries?

¹² Does this sit well with Accuracy? If our aim must be rely on p iff p , how can the adequate norm be non-factive? As I said, practical aims are not usually taken to translate directly into practical norms. What we expect of agents is to rely on p depending on whether their perspective suggests that this will satisfy Practical Aim – and that may very well be the case even if Practical Aim is not in fact satisfied.

¹³ Although it may seem otherwise at first glance, Schulz's theory grants us a similar freedom. In his theory, much depends on what it means to have a rational credence of 1. Many authors think that it is hardly ever rational to have such a credence. This cannot be what Schulz has in mind. In fact, he mentions that the relevant probability function will often be one that would result from updating with “negations of sentences we do not know to be false” (Schulz 2010b, 372). Given this, we again have the possibility to say that “must” allows for a possibility of error.

Yes. But I mean to propose something more. I propose that we use epistemic modals not just in the primary norm but also in the secondary norm. More specifically, I propose the following:

Epistemic Modal Reasonableness (EMR)

If a norm says that a subject ought to rely on *that p* only if condition *c* holds, the subject is reasonable (with regard to that norm) in relying on *that p* iff given the information of the subject, it must be that *c* holds.

For many primary epistemic norms, this pretty much coincides with the notion of Reasonableness described above. For instance: To be reasonable in the modal sense with regard to the Knowledge Norm, your information must be such that relative to it, it must be that you know the proposition in question. This will be so if your information establishes the epistemic fact that you know this. Roughly at least, this seems equivalent to the demand that you are justified in believing you know. Thus, EMR should enjoy much of the intuitive support of the previous norm of Reasonableness.

Things are dramatically different if our primary norm is EMN. The crucial point is this: The Epistemic Must Norm is the *fixed point* of the secondary norm EMR. This is because on all the semantic views I have mentioned above, iterations of \Box are logically vacuous – that is, the characteristic S4 axiom ($\Box p \rightarrow \Box \Box p$) is validated. In fact, typical versions of the semantics claim that \Box and \Diamond satisfy the axioms of the modal logic S5 (see Yalcin 2007, 994; Willer 2013, 12; Schulz 2010b, 381 ff); since I have opted against factivity, this is not true for the operator in EMN.

Since iterations of \Box are vacuous, the condition expressed by EMN is guaranteed to be luminous. If it must be that *p*, it must be that it must be that *p*. And, as I said, both conditions are associated with processes of inquiry that focus on whether *p*. To establish whether it must be that it must be that *p*, you do what you do when you establish whether it must be that *p*, which, again, is just what you do when you establish whether *p*. That means: To be reasonable in the light of EMN, you need not ask higher-level epistemological questions. You must only inquire whether *p*!

In other words: Given a modalized notion of reasonableness, EMN is primary and secondary norm at once. As we have seen, this is *not* because the secondary norm EMR is vacuous *per se*. In combination with other candidate norms, it imposes substantive constraints. But it is a notion of reasonableness which, when associated with the EMN, delivers the same criterion as the EMN.

Importantly, EMR is, unlike EMN, explicitly *tied to the subject's body of information*. As far as the primary norm, EMN, is concerned, I have allowed for contextual variability. But the very idea of a “secondary” norm is to take into account, specifically, the nature of the subject's own perspective.

For example: Suppose that Emily is on her way to a concert, and she is already late. Alas, she feels insecure about whether she turned off the stove. Should she head on to the concert, relying on the premise that it is off? In Emily's context, EMN permits this iff (given Emily's information) the stove must be off. We can imagine that she used the stove before she left, that she has not checked whether it is off, and that she has a track record of sometimes leaving it on. In this case, her evidence does not license the claim that it must be off. Unable to rely on that premise, she returns. At the same time, an eavesdropper watches her, unable to interfere and communicate with her. This eavesdropper is better informed. She can tell that there is no way that the stove is on. So she can say: "Oh no, poor Emily. She should have relied on proposition that the stove is off, because in fact it must be off. However, I do not blame her. It was reasonable to return; given *her* information, the stove might be on." I find this combination of judgments extremely plausible.

Note that when we judge EM-reasonableness, the prefix that determines the subject's information as the relevant information state will take wide scope over both occurrences of "must". So when we ask whether Emily is reasonable in relying on a proposition p , we ask whether it must be (relative to Emily's information) that it must be (relative to Emily's information) that p . Since the state is not allowed to "switch" midway in the interpretation, this iteration of modals is again vacuous. So in effect, my proposal says that when we contrast what Emily *should* treat as a reason, and what she would be *reasonable* to treat as a reason, we are contrasting two "must" claims, one "autocentric" (relativized to our own information) and one "exocentric" (relativized to Emily's).

Consequently, the distinction between primary and secondary propriety will matter only from a third-person point of view (as they say), i.e. from contexts with different amounts of information. Within the "egocentric" predicament, by contrast, the distinction collapses, since the subject will always ask what must be the case relative to her information. Accordingly, Kvanvig's worry does not arise. The reasoning subject will never face the task of guiding herself by two divergent norms.

8. Modal and Probabilistic Reasons

Often, reasons on which we act are themselves modalized (in an epistemic sense), e.g. when we decide to go downstairs because the keys might be on the kitchen sideboard. One corollary of my proposal concerns these cases: Any epistemic modal claim which is true (relative to the information state of the reasoner) automatically passes the norms, EMN and EMR. So if a reason is itself about what might or must be the case, both of my epistemic norms will coincide with the Truth Norm.

Reasoning under uncertainty is naturally conducted in terms of what might or must be the case. These, I suggest, provide ways for us to frame considerations that we can treat as epistemic fixed points, since they will, if true, automatically pass the primary and secondary epistemic norms.

One last thing I wish to discuss concerns *probabilities*. I am sympathetic to the view that much of our reasoning involves credences, i.e. probabilistic degrees of belief, rather than simple yes-or-no premise states. But many authors assume that probabilistic contents do not felicitously embed under epistemic modals. How can we apply EMN to probabilistic inputs to practical reasoning?

One option is to take inspiration from defenses of the Knowledge Norm. Stanley and Hawthorne suggest that credences pass their norm iff corresponding propositions about epistemic probabilities are known. But if I adopted this proposal, this would mean that in many cases, we are stuck with a need for higher-order evidence again. (Often we would have to check whether it must be that it is epistemically likely that p , or the like.) Transparency would be out of the window.

But there may be better options. Moss (2017) has argued that thoroughly probabilistic contents (and not just propositions about epistemic probabilities) do embed under epistemic modals, and she has developed a semantic theory to allow for such embeddings. As before, I can only give the briefest of sketches. The general idea is that epistemic modals are actually evaluated relative to a pair of an information state and a partition on this state. This partition gives alternative hypotheses, or a question. Epistemic modals do not quantify directly over worlds but over cells in the partition. And “It must be probable that p ” will mean that for all hypotheses (cells) in the partition, if we update our credence with that hypothesis the result will be a probability of $p > .5$.

Adopting some such proposal would make EMN and EMR apply even to probabilistic contents. And Transparency, it seems to me, would be respected. We accept “It must be that it is probable that p ” simply by being credally related to p in the right way, such that whatever assumption (out of a certain range) we make, we become or remain more than .5 sure that p is true.

9. Knowledge Again

I have argued that knowledge is not the norm of practical reasoning, and that the correct norm uses epistemic modals. Of course, I had to leave many questions unaddressed. I briefly mention one. I have granted that the Knowledge Norm derives support from ordinary discourse. Epistemic criticism is often framed in terms of knowledge. If my account is correct, why would this be so?

Though I cannot discuss this question at length, I will at least hint at a possible answer. As I said, orthodox semantic views assimilate epistemic modals to knowledge claims (see § 5 above).

But I think that for many ordinary uses, this direction should actually be reversed. Many instances of knowledge claims in ordinary discourse should be understood as very similar to epistemic modal claims (on their non-factualist construal), not *vice versa*. Here is what I have in mind: In other work (Henning 2018), I argue that certain attitude verbs (like “believe” and “want”) have *parenthetical* readings. That is, in many uses of sentences of the form “S believes that *p*”, the matrix clause “S believes that” semantically contributes a backgrounded side-remark, while the embedded clause “*p*” contributes the so-called at-issue content. The idea would be that “know” also has parenthetical uses in this sense. Empirical work in linguistics suggests precisely this (see, e.g., Simons 2007).

So my hypothesis would be the following: A criticism like “You do not even know whether your friend is at the pub” must be interpreted as “For all you know, your friend is not at the pub.” In this last construction, as in “Your friend might not be at the pub”, the at-issue content concerns the whereabouts of your friend, not your epistemic state (which is subject of a side-remark.) Parenthetical uses of this kind are, in important respects, to be treated like epistemic modal claims.

Whether this rough idea is borne out by the linguistic facts is a question for future work.¹⁴

References

- Alston, W. P. 1980, “Level Confusions in Epistemology”, *Midwest Studies in Philosophy* 5: 135–150.
- Broome, J. 1999, “Normative Requirements”, *Ratio* 12 (4):398–419.
- Brown, J. 2008, “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning”, *Noûs* 42: 167–189.
- Del Pinal, G. & Waldon, B. 2019, “Modals under Epistemic Tension”, *Natural Language Semantics* 27 (2):135-188.
- DeRose, K. 2009, *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press.
- Douven, I. 2006, “Assertion, Knowledge, and Rational Credibility”, *Philosophical Review* 115 (4):449-485.
- Enoch, D., Spectre, L., Fisher, T., 2012, “Statistical Evidence, Sensitivity, and the Legal Value of Knowledge”, *Philosophy and Public Affairs* 40: 197-224.
- Gendler, T. S. 2008 “Alief and Belief”, *Journal of Philosophy* 105 (10):634-663.
- Gerken, M. 2011, “Warrant and Action.” *Synthese* 178: 529-547.
- Hawthorne, J. and Stanley, J. 2008, “Knowledge and Action”, *Journal of Philosophy* 105: 571–590.
- Henning, T. 2018, *From a Rational Point of View*, Oxford: Oxford University Press.
- Hintikka, J. 1962, *Knowledge and Belief. An Investigation Into the Logic of the Two Notions*, Ithaca: Cornell University Press.
- Kolodny, N. 2005, “Why be Rational?”, *Mind* 114 (455):509-563.

¹⁴ For helpful comments on drafts and presentations, I thank Christian Nimtz, Ralf Busse, Nick Haverkamp, Timo Meier, Jan Barbanus, participants in the workshop *Dimensions of Rationality* at the University of Frankfurt, and two referees for this journal.

- Kratzer, A. 1981, "The Notional Category of Modality", in H.-J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, Berlin and New York: DeGruyter.
- Kvanvig, J. L. 2009, "Assertions, Knowledge, and Lotteries." In Patrick Greenough and Duncan Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press.
- Kvanvig, J. L. 2011, "Norms of Assertion", in Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*. Oxford: Oxford University Press.
- Lackey, J. 2007, "Norms of Assertion." *Noûs* 41: 594-626.
- Littlejohn, C. 2009 "Must we Act only on What we Know?", *Journal of Philosophy* 106 (8):463-473.
- Littlejohn, C. 2012 *Justification and the Truth-Connection*. Cambridge University Press.
- MacFarlane, J. 2014 *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- Mihoc, T., Bhadra, D. & Falaus, A. 2019 "Epistemic Modals, Deduction, and Factivity: New Insights from the Epistemic Future", *Proceedings of SALT* 29: 351-370.
- Moss, S. 2017, *Probabilistic Knowledge*, Oxford: Oxford University Press.
- Neta, R. 2009, "Treating Something as a Reason for Action", *Noûs* 43: 684–699.
- Pritchard, D. "Epistemic Luck, Safety, and Assertion," in Clayton Littlejohn & John Turri (eds.), *Epistemic Norms: New Essays on Action, Belief, and Assertion*, Oxford: Oxford University Press.
- Radford, C. 1966 "Knowledge – By Examples", *Analysis* 27 (1):1--11.
- Schechter, J. 2017 "No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion", In J. Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge-First: Approaches in Epistemology and Mind*. Oxford University Press. pp. 132-159.
- Schnieder, B. 2011, "Expressivism Concerning Epistemic Modals", *The Philosophical Quarterly* 60: 601-615.
- Schulz, M. 2010a, "Wondering What Might Be", *Philosophical Studies* 149: 367-386.
- Schulz, M. 2010b, "Epistemic Modals and Informational Consequence", *Synthese* 174: 385–395.
- Simons, M. 2007, "Observations on Embedding Verbs, Evidentiality, and Presupposition", *Lingua* 117: 1034-1056.
- Stalnaker, R. 1984, *Inquiry*. Cambridge University Press.
- Veltman, F. 1996, "Defaults in Update Semantics", *Journal of Philosophical Logic* 25: 221–261.
- Yalcin, S. 2007, "Epistemic Modals", *Mind* 116: 983–1026.
- Yalcin, S. 2011, "Nonfactualism about Epistemic Modality", in A. Egan and B. Weatherson (eds.), *Epistemic Modality*, Oxford University Press, 295–332.
- Weiner, M. 2005, "Must we Know what we Say?" *Philosophical Review* 114 (2):227-251.
- Williamson, T. 2000, *Knowledge and its Limits*, Oxford: Oxford University Press.
- Williamson, T. 2005, "Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge," *Philosophical Quarterly* 55:
- Yablo, S. 2011, "A Problem about Permission and Possibility", in A. Egan and B. Weatherson (eds.), *Epistemic Modality*, Oxford University Press, 270-29.