

Thinking and being sure

Jeremy Goodman & Ben Holguín

January 2, 2022

How is what we believe related to how we act? That depends on what we mean by ‘believe’. On the one hand, there is what we’re sure of: what our names are, where we were born, whether we’re currently looking at a screen, and so on. Surety, in this sense, is not uncommon – it does not imply Cartesian absolute certainty, from which no possible course of experience could dislodge us. But there are many things that we think that we are not sure of. For example, we might think that it will rain sometime this month, but not be sure that it will. Both what we’re sure of and what we think have important normative connections to action. But the connections are quite different. We’ll begin in section 1 by describing two surety norms, one on assertion and another on inquiry. In section 2 we appeal to cases of knowing without being sure to argue that those norms do not derive from parallel knowledge norms on assertion and inquiry. Section 3 considers norms framed in terms of what is certain, and section 4 considers norms on being unsure. In section 5 we turn to thinking, and explain its connections to guesses and forced choices. Section 6 explains related connections between thinking and decision making, by arguing for a particular link between intention formation and normative thought. In section 7 we argue that, in light of the forgoing, there is little reason to believe in a notion of “full belief”, intermediate in strength between the ordinary notions of thinking and being sure, but with some distinctive normative connections to action.

1 Asserting and inquiring

When we assert something, we represent ourselves as being sure of it.¹ This is why, in response to an assertion that p , ‘What makes you sure that p ?’ is a felicitous query, in a way that, say, ‘What makes you pleased that p ?’ is not. It is also why assertions of the form ‘ p , although I’m not sure that p ’ sound terrible, in a way that, say, ‘ p , although I’m not pleased that p ’ need not: you represent yourself as sure that p while also denying that you are sure that p , and so cannot be speaking truly (another norm on assertion we will return to) while also being sure of what you say.²

¹See Unger (1975), Stanley (2008), Beddor (2020).

²Some will deny that there is a surety norm on assertion on the grounds that such a norm would make proper assertion unrealistically demanding. The idea is that we are not (or

When we assert something, we also represent ourselves as knowing it.³ This is why, in response to an assertion that p , ‘How do you know that p ?’ is a felicitous query, and why assertions of the form ‘ p , although I don’t know that p ’ sound terrible. So there is both a surety norm and a knowledge norm on assertion. The question is how these norms are related.

Norms give rise to other norms in at least two ways. First, suppose there is a norm of the form: “You should ϕ only if you are F ”. If being G is a necessary condition on being F , then there is another *descriptively parasitic* norm “You should ϕ only if you are G ”. (If one cannot be F without being G , then one cannot satisfy the norm to ϕ only if one is F without also satisfying the norm to ϕ only if one is G .) Second, suppose that in addition to the norm “You should ϕ only if you are F ” there is also a norm “You should be F only if you are G ”. Then there is a *normatively parasitic* norm “You should ϕ only if you are G ”. (If one ϕ s without being G , then one must either be in violation of the norm to ϕ only if one is F or in violation of the norm to be F only if one is G .)

With these two notions of norm parasitism in place, there are a number of hypotheses we might entertain about how different norms on assertion are related. According to Unger and Williamson, the norm “You should assert p only if you know p ” is basic, and gives rise to a descriptively parasitic norm “You should assert p only if p ” (since knowledge is factive) and also to a descriptively parasitic norm “You should assert p only if you are sure that p ”, since Unger and (as we read him) Williamson hold that being sure that p is a necessary condition on knowing that p .

Note that in saying that there is a *basic* knowledge norm on assertion, these authors are not denying that there are parasitic knowledge norms on assertion too. Indeed, Williamson is committed to such parasitic norms, since he accepts the norm “You should be sure that p only if you know that p ”. From the surety norm on assertion, which Williamson thinks is descriptively parasitic on the basic knowledge norm, and this (perhaps basic) knowledge norm on surety, it follows that there is a doubly parasitic knowledge norm on assertion, in addition to the basic one he endorses.

In light of this fact, one might question the need for a basic knowledge norm on assertion in the first place. For if one instead takes there to be a basic surety norm on assertion (“Assert p only if you are sure that p ”) and a basic knowledge norm on being sure (“Be sure that p only if you know that p ”), then these will generate a (normatively parasitic) knowledge norm on assertion. This is our

at least should not be) sure of very much – much less than it is appropriate to assert. We think this thought is simply mistaken, conflating semi-technical notions of absolute certainty or credence 1 with the attitude ordinarily ascribed using ‘be sure’ in English. For note how common it is to describe oneself or others as being sure of things, as in ‘I’m sure this won’t work’ or ‘I’m not sure they’ll win, but I’m sure they’ll try’. Moreover, if surety were so rare, then denials of surety would seem trivial. But they often seem perfectly informative, as in ‘I’m not sure who will be at the party’ or ‘Jones isn’t sure he’ll make it home in time’.

Dorst and Mandelkern (forthcomingb) argue against the surety norm on assertion from a different direction, by arguing that merely reasonably thinking that p suffices for properly asserting p .

³See Moore (1942), Unger (1975), and Williamson (2000, ch. 11).

preferred view, which we will expand on in the present section and argue for in the following one.

To illustrate this proposal, let's consider how it explains the infelicity of '*p*, but I don't know that *p*', and how this explanation differs from its explanation of the infelicity of '*p*, but I'm not sure that *p*'. The reason why an assertion of '*p*, but I don't know that *p*' is infelicitous is that it expresses a proposition the speaker doesn't know. (If they knew it they'd know its first conjunct, in which case its second conjunct would be false, and the conjunction untrue, and hence not known.)⁴ This means that in asserting '*p*, but I don't know that *p*', either you're unsure of what you've asserted – thereby violating the surety norm on assertion – or you're sure of something you don't know – thereby violating the knowledge norm on being sure. Either way, you've done something improper. This inevitable impropriety explains the assertion's infelicity.

What about '*p*, but I am not sure that *p*'? Suppose you were to assert it. Then either you aren't sure of what you assert – thereby violating the surety norm on assertion – or you are sure of what you assert, in which case you are presumably sure of its first conjunct, and are hence asserting something false, and hence something you don't know – violating the knowledge norm on surety. Either way, you've done something improper. This inevitable impropriety explains the assertion's infelicity.

Williamson (p.c.) objects to any account of the knowledge norm on assertion that make it merely normatively parasitic on other norms. His objection is that, although such accounts correctly predict that when one asserts without knowing one is in violation of *some* norm, they fail to predict that when asserts without knowing there is something wrong with one's assertion *as such*. That is to say: if the surety norm on assertion is basic while the knowledge norm is parasitic, then since it is not impossible to be sure both that *p* and you're not sure that *p*, an assertion of '*p*, but I'm not sure that *p*' could be entirely proper *qua* *assertion*. Its felt impropriety would not be due to the fact that you've *said* anything you're not supposed to have said; rather its felt impropriety would be due to the fact that you've revealed yourself to be in a normatively defective state of mind – viz., a state of being sure of something you don't know.

We are unmoved by this objection, mainly because the judgment that there is something wrong with such assertions “as such” strikes us as tenuous, and certainly nowhere near as robust as the conversational judgments appealed to above. But more to the point, consider assertions like ‘I am sure that *p*, though all of my evidence suggests that not-*p*’. At least when considered schematically and without context, such assertions strike us as bad in essentially the same way that assertions of '*p*, but I am not sure that *p*' do. Yet ‘I am sure that *p*, though all of my evidence suggests that not-*p*' expresses a proposition that a person

⁴Here we suppress subtleties to do with the possibility that ‘know’ is context-sensitive. Suppose in our mouths ‘know’ expresses knowledge, but in some other speaker's mouth it expresses a different relation knowledge* that is more demanding than knowledge, and the speaker both knows *p* and knows they don't know* *p*. Then they could speak from knowledge in using '*p*, but I don't know that *p*'. See Worsnip (2017) for a helpful review of recent work on how to think about knowledge norms on assertion if there is such context-sensitivity.

could know about themselves, and hence, according to Williamson, could be asserted without violating any basic norm on assertion. Whatever Williamson says about the anomalousness of such assertions (e.g., that only a bizarre person could make them in conformity with the basic norms on assertion), we expect a parallel explanation will be available for ‘ p , but I’m not sure that p ’ assuming a basic surety norm on assertion.

Of course, this all is merely defensive maneuvering: we have not yet said anything that tells against Williamson’s claim that the knowledge norm on assertion is basic and the surety norm parasitic. Indeed, that picture is quite elegant. Our argument against it will be given in the next section.

Let us now turn to norms on inquiry, which have been less systematically explored than norms on assertion. A recent exception is Friedman (2019),⁵ who defends the norm “You should inquire into whether p only if you do not believe that p ”. She uses the example of inspector Morse, waking up knife in hand, covered in blood, next to a corpse, with no memory of the night before. Friedman argues that, if Morse believes that he committed a murder, but goes about his business investigating the crime anyway, then his inquiry is a sham. By contrast, as long as he fails to believe that he committed a murder, even if he has his suspicions, his inquiry need not be a sham.

We take Friedman’s work to suggest an important norm on inquiry: namely that you should inquire into whether p only if you are not sure that p . For suppose Morse thinks he committed a murder, but isn’t sure that he committed a murder. The fact that he is not sure seems sufficient to prevent his inquiry from being a sham, despite the fact that he thinks he did it.

Whether or not this this marks a disagreement with Friedman depends on what she intends by ‘believe’: if it is the state of mind ordinarily ascribed by ‘be sure’, then there is no disagreement; otherwise, there is. In section 7 we will consider how to think about ‘believe’ as it is used in contemporary epistemology. Until then we will focus on our preferred norm on inquiry: “Inquire into whether p only if you aren’t sure that p ”.

As in the case of assertion, it is worth considering what knowledge norms there may be on inquiry, and how these may be related to surety norms. Consider the norm “Inquire into whether p only if you don’t know that p ”. If there is such a norm, and if there is also a knowledge norm on surety, then the surety norm on not inquiring will be normatively parasitic on the knowledge norm on non-inquiry. For if one inquired into whether p while being sure that p , then one would either be violating the knowledge norm on surety or violating the non-knowledge norm on inquiry.

As with norms on assertion, this knowledge-centric account of the norms on inquiry is quite elegant. In the next section we will argue that it is nevertheless mistaken, and that surety-based norms are preferable with respect to both assertion and inquiry.

⁵See also Beddor (ms) and Woodard (ms).

2 The primacy of surety norms

Our argument for the primacy of surety norms on assertion and inquiry will turn on the fact that, *pace* Unger and (our reading of) Williamson, being sure that p is not a necessary condition on knowing that p .⁶ If this is right, then the aforementioned knowledge norms do not give rise to parasitic surety norms. Yet, in cases where one knows without being sure, assertion is not proper, which is not predicted by a basic knowledge norm on assertion, and inquiry need not be improper, contrary to what is predicted by a basic ignorance norm on inquiry.

There can be cases of knowing that p without being sure that p because there can be cases of remembering that p without being sure that p and cases of perceiving that p without being sure that p , and both remembering that p and perceiving that p are ways of knowing that p .⁷ We will consider both cases in turn.

Suppose Jones reads, and thereby comes to know, that the Battle of Hastings was fought in 1066. We may suppose that, at that moment, he is also sure that the battle was fought then. Years later, he has long since forgotten reading the book. Someone asks him if he knows when Battle of Hastings was fought. He replies ‘I’m not sure, but I think it was in 1066’. Although Jones is no longer *sure* that the battle was fought in 1066, he still *thinks* that it was, and his thinking this is the result of a memory trace preserved from the time he originally learned it. In such a case, we think Jones knows: he hasn’t forgotten that the battle was fought then, so he still remembers that it was fought then, so he knows that it was fought then.

We should note that, in saying this, we are not claiming anything as strong as Radford (1966), who influentially argued that remembering that p is compatible with being sure that one does not know that p , and indeed that one has never learned anything about whether p . Our argument requires nothing so extreme. Jones is pretty confident that the Battle of Hastings was fought in 1066. He just isn’t sure that it was.

It would be improper for Jones to flatly assert ‘The Battle of Hastings was fought in 1066’. For if it were proper, it would also be proper for him to assert ‘The Battle of Hastings was fought in 1066, but I’m not sure that it was’ (since he is clearly in a position to assert the second conjunct), and such sentences continue to sound terrible. This impropriety is not explained by a knowledge norm on assertion, since Jones does know that the Battle was fought then and that he is not sure that it was. But it is explained by a surety norm on assertion, since Jones isn’t sure of this conjunction.

⁶Beddor (2020) makes similar arguments in favor of the primacy of a certainty norm on assertion. We will consider his proposal in the next section.

⁷Jason Stanley (2008) has given the following independent argument for the possibility of knowing without being sure. He notes that assertions of the form ‘I know that p , although I’m not sure that p ’ sound notably worse than assertions of the form ‘They know that p , but they aren’t sure that p ’. If surety were a necessary condition for knowledge, then both assertions would sound terrible, since they would be asserting something impossible. But if surety is not necessary for knowledge, then the first-person/third-person contrast can be assimilated to the contrast between ‘ p , but I’m not sure that p ’ and ‘ p , but they’re not sure that p ’.

The same goes for inquiry. It would not be improper for Jones to inquire into the question of when the Battle of Hastings was fought. It would be perfectly reasonable for him to check Wikipedia – not merely to convince someone else, but also to convince himself. A knowledge norm on inquiry incorrectly predicts that this would be improper. A surety norm does not.⁸

Similar cases can arise in perception. Smith is looking at a red wall in ordinary light. But she has some suspicion that it might be a white wall in red light. As a result, she is not sure that it is red. Still, her visual system is functioning normally. It is therefore plausible that Smith sees that the wall is red. Moreover, it is plausible that seeing that something is the case is a way of knowing that it is the case. This is a consequence, for example, of Williamson’s (2000, ch. 1) influential thesis that knowledge is the most general factive mental state. But more direct motivations are available as well. If Smith sees that the wall is red, then surely she can tell that the wall is red. Claiming otherwise seems terrible: ‘Smith sees that the wall is red, but she can’t tell that the wall is red’ sounds like a contradiction. Moreover, if she can tell that the wall is red, then surely she knows that the wall is red. Again, claiming otherwise seems terrible: ‘Smith can tell that the wall is red, but she does not know that the wall is red’ also sounds like a contradiction. This suggests that seeing that the wall is red must be a way of knowing that it is red.

So we have another case of knowing without being sure. As before, assertion would be improper: Smith cannot properly assert that the wall is red. Likewise, inquiry would be proper: there would be nothing improper about her checking the lighting.

Our judgments about cases of memory and perception thus suggest that the basic norm on assertion is not a the knowledge norm – “Assert that p only if you know that p ” – but a surety norm: “Assert that p only if you are sure that p ”. Similarly, the basic norm on inquiry is not an ignorance norm – “Inquire into whether p only if you don’t know whether p ” – but a being unsure norm: “Inquire into p only if you’re not sure whether p ”.

Of course, many epistemologists have thought that cases of the kind we have described are counterexamples to the claim that remembering that p and perceiving that p are ways of knowing that p . We suspect that many of them have been drawn to this conclusion because they look at unsure rememberers and perceivers and think “That person isn’t behaving like someone who knows; nor

⁸Beddor (ms) uses a similar case (his ‘Ancient History’ example) to argue that the ‘aim’ of inquiry is not knowledge but rather being ‘absolutely certain’/‘having credence 1’. While his overall view is in many ways close to ours, there are some important differences stemming from the fact that agents can be sure of propositions without being absolutely certain of them and/or having credence 1 in them. For example: suppose you inquire into whether a coin is fair or double headed by flipping it repeatedly. Eventually, you come to know and be sure that it is double headed (cf. Bacon (2014), Dorr et al. (2014)), at which point further inquiry would be improper. But assuming your credences evolve by Bayesian conditioning on the observed outcomes of the tosses, you will never come to have credence 1 that it is double headed. Even if we allow your credences to evolve in some other way, it strikes us as unnatural to distinguish the moment when you first come to be sure from the moment when further inquiry first becomes improper.

should they – they couldn’t assert the proposition in question, and indeed it would be entirely reasonable for them to inquire into its truth.” This objection is a real challenge for those, like Williamson, who think that knowledge and ignorance are the basic norms on assertion and inquiry. But they are not problems for those, like us, who think that being sure and being unsure are the norms on assertion and inquiry, since the subjects in question are unsure. In this way, by achieving a degree of normative distance between knowledge and action, we render more plausible the main descriptive doctrine of “knowledge-first” epistemology, according to which all factive mental stative attitudes are kinds of knowledge.

We have argued that a surety norm on assertion is more predictive than a knowledge norm and that an ignorance norm on inquiry makes incorrect predictions that an unsurety norm does not. Still, this leaves open a number of difficult questions in the theory of the norms on assertion and inquiry.

Take sentences of the form ‘ p , but for all I know I don’t know that p ’. The infelicity of such sentences is not explained by a knowledge norm on assertion (at least not by itself). It would, however, be explained by an iterated-knowledge norm on assertion, according to which one should assert that p only if one knows that one knows that p . With sentences like these in mind, Cohen and Comesaña (2013), Greco (2015a,b), and Dorst (2019b) claim that there is an iterated-knowledge norm on assertion and that it is *descriptively* parasitic on the (non-iterated) knowledge norm. This is because they accept the KK thesis, which says that knowing implies knowing that you know. Indeed, they take the existence of an iterated-knowledge norm on assertion to be powerful evidence in favor of KK.

Now, one can believe in an iterated-knowledge norm without thinking that it is descriptively parasitic on the non-iterated norm in this way. For example, one could instead postulate a basic iterated-knowledge norm on assertion. Alternatively, one could posit a different mechanism whereby the iterated-knowledge norm might arise from other more basic norms.⁹ Either way, the important point for present purposes is that an iterated-knowledge norm on assertion, whatever its source, still cannot account for all the data that motivate a basic surety norm on assertion. This is because the cases that demonstrate the possibility

⁹Here is a sketch of such a mechanism. We sometimes criticize people for being reckless about a norm even in circumstances where they happen to conform to it (cf. Williamson (2005), Benton (2013)). We might codify this tendency by saying that, if there is a norm “You should ϕ only if you are F ”, then there is an *epistemically parasitic* norm “You should ϕ only if you know that you are F ”. In the case of the knowledge norm on assertion, the idea would be that there is something problematically risky about asserting a proposition that for all one knows is in violation of the norm of assertion, just as there is something problematically risky about serving drinks that for all one knows are in violation of the norm “Don’t serve poisonous drinks”. On this view, the iterated-knowledge norm on assertion would be epistemically parasitic on the (non-iterated) knowledge norm. (If we reject KK, then for such a view to be plausible it will probably have to be that the force of the epistemically derivative norms diminishes the more times epistemic parasitism must be invoked to derive them.) Note that epistemic parasitism could itself be normatively parasitic on a more basic phenomenon of *surety parasitism* (according to which, if there is a norm “ ϕ only if you are F ”, then there is derivative norm “ ϕ only if you are sure you are F ”).

of knowledge without surety can be modified to demonstrate the possibility of iterated knowledge without surety.

Suppose in the perception case that Smith has the background knowledge that her visual system is functioning normally, and that she is either looking at a red wall in ordinary light or a white wall in red light. Suppose she also knows that, if the wall is red in ordinary light and her visual system is working normally, then she knows that the wall is red. (After all, if this is something *we* can know about Smith on general philosophical grounds, then surely it is something she can know about herself in the same way.) Putting together (i) her background knowledge about her eyesight and circumstances, (ii) her background knowledge about the epistemology of perception, and (iii) her perceptual knowledge about the color of the wall, Smith can know that she knows that the wall is red. Yet she will continue to be unsure whether it is red.

Similar considerations apply in the memory case. Jones knows (by remembering) that the Battle was in 1066, knows (by introspection) that he believes it was fought in 1066, and is thus in a position to know (by deduction) that he truly believes the battle was fought in 1066. He also knows (by inference to the best explanation) that if he has a true belief about the date of the battle, it is because of a memory trace from having once learned it. And so he knows (by reflecting on all this) that a memory trace of this kind would constitute remembering, and hence knowing, that the Battle was fought in 1066. Putting the pieces together, he knows that he knows that the Battle was fought in 1066, though he is not sure that it was fought then.

In both the perception and memory cases, despite the fact that the agent knows that they know a certain proposition, they are not in a position to assert that proposition, and inquiring into its truth would be perfectly in order. So iterated-knowledge norms are no substitute for surety norms.

3 On certainty

There has been very little discussion in contemporary epistemology framed in terms of what we can “be sure” of. There are exceptions, such as Ayer’s (1956) analysis of “knowing as having the right to be sure”, but much more discussion has been framed in terms of what we can “be certain” of. Our view is that ‘*S* is sure that *p*’ means the same as ‘*S* is certain that *p*’, with the latter perhaps being slightly more formal and perhaps having a greater propensity to shift the contextual standards of what is required for surety/certainty in a more demanding direction. This allows us to bring the present discussion into dialogue with those who have written about norms related to being certain, principally Unger, Stanley, and Beddor.

Unlike ‘be sure’, ‘be certain’ also has an epistemic use, ‘It is certain that *p*’. On this use, ‘certain’ behaves more like the epistemic modal ‘must’ than like a propositional attitude such as ‘think’ or ‘know’. It is worth noting that assertions of the form ‘*p*, but it isn’t certain that *p*’ (and ‘*p*, but it might not be the case that *p*’) seem to be infelicitous in the same way that assertions of the

form ‘ p , but I’m not sure/certain that p ’ are. One might wonder whether this is evidence that there is an epistemic certainty norm on assertion, of the form “Assert p only if it is certain(/must be the case) that p ”.

While there may be an important norm in the vicinity, this norm is subject to obvious counterexamples. Suppose we know that you have a coin in one hand but it’s not certain which hand: it might be in your left and it might be in your right. Of course, you know which of your two hands the coin is in. Suppose you’ve told someone else. Then you’ve either said that it’s in your left hand or said that it’s in your right hand, despite it not being certain that it is in your left hand nor certain that it is in your right hand. Your assertion wasn’t thereby improper. The reason, of course, is that although it might be in either hand, the sentence ‘It might be in either hand’ would be false in your mouth, at least on its most natural interpretation. Similarly, ‘It’s neither certain that it’s in my left hand nor certain that it’s in my right hand’ would be false in your mouth in the context of a conversation in which you’ve just revealed the location of the coin to a third party. Like epistemic modals, epistemic certainty is context-sensitive in ways that tend to be anchored to the epistemic situation of the speaker. The norm floated at the end of the last paragraph fails to account for this fact.

Mindful of this context-sensitivity, and following Stanley (2008) and Beddor (2020), we’ll say that p is *certain for* a person just in case p is in the extension of ‘it is certain that’ in their mouth. This allows us to formulate the following epistemic certainty norm on assertion: “Assert that p only if it is certain for you that p ”. Stanley and Beddor endorse this norm as the best explanation of the infelicity of ‘ p , but it isn’t certain that p ’, and we agree with them. As with the knowledge norm on assertion, the question is how this norm relates other norms on assertion – in particular, to the norm “Assert that p only if you are sure/certain that p .”

The simplest way of accounting for the epistemic certainty norm on assertion – at least given what we have said so far – would be to treat it as descriptively parasitic on the knowledge norm on assertion, by holding that anything you know is epistemically certain for you. It is important to observe that people like Jones and Smith who know without being sure are not obviously counterexamples to this thesis. They are not sure/certain of what they know, but that does not mean that, e.g., ‘It is certain(/must be the case) that the Battle of Hastings was fought in 1066’ would be false in Jones’s mouth. Of course, it wouldn’t be something he’s in a position to assert, but that’s because it isn’t something he’s certain of. (As we’ll discuss shortly, Stanley and Beddor agree that there are cases where someone isn’t certain that p despite p being certain for them.)

A different strategy for unifying the knowledge, surety, epistemic certainty norms on assertion would be to propose a single basic norm on assertion: one should assert that p only if one *knows for sure* that p . Since knowing for sure plausibly entail knowing, being sure, and epistemic certainty, the three corresponding norms on assertion will then be descriptively parasitic on the knowing-for-sure norm. This proposal might seem to disreputably rely on an idiomatic notion of ‘knowing for sure/certain’, but as Beddor (2020) notes, constructions of this form are surprisingly cross-linguistically robust. Moreover, the notion

needn't be taken as primitive – for example, it is naturally understood as having a state of surety/(subjective) certainty that amounts to knowledge.^{10,11}

Stanley and Beddor take a different tack. According to them, the primary norm on assertion is epistemic certainty rather than surety. We have a number of reservations about this proposal.

An initial worry is that, as noted above, knowledge may suffice for epistemic certainty. If it does, then the proposal fails to correctly predict the impropriety of asserting things you know but aren't sure of.

A second worry is that epistemic certainty may not suffice for knowledge. If it doesn't, then the proposal fails to correctly predict the impropriety of asserting things you don't know. One reason epistemic certainty may not suffice for knowledge is that the two notions arguably have different logical behavior. People don't know every logical consequence of the things they know. By contrast, epistemic certainty is arguably closed under logical consequence (in the sense that any proposition entailed by propositions that are certain for a person will itself be certain for that person). Indeed, Beddor argues for this claim roughly as follows: unembedded uses of 'it is certain that' and 'must' are interchangeable; 'must' expresses the property of being true in every world consistent with a given body of evidence; so epistemic certainty has the same normal modal logic as 'must' does. If he is right, then the truth of Fermat's Last Theorem was certain for Peano even though he didn't know it. The epistemic certainty norm on assertion then fails to correctly predict the fact that Peano shouldn't have asserted it.

Similarly, epistemic certainty seems not to be sufficient for being sure/certain. This non-entailment is explicitly endorsed by Stanley and Beddor. How then do they propose to explain the surety norm on assertion? We will focus on Beddor's discussion, which strikes as the more promising of the two.¹² He claims that

¹⁰This is arguably more demanding than both knowing p and being sure that p . For example, Jones might become sure that the battle was fought in 1066 on the basis of reading tea leaves while continuing to know it was fought then by remembering that it was, but the tea-leaf-induced state of surety won't amount to knowledge, so Jones won't know for sure. If this is right, then the knowledge norm on surety should probably be strengthened to a knowing-for-sure norm.

¹¹Note that this proposal conflicts with the one defended in the last section, that the knowledge norm on assertion is normatively parasitic on the surety norm. We are open to both views. What we are opposed to is the idea that surety norms are parasitic on mere knowledge norms.

¹²Here is Stanley's account:

Subjective certainty, unlike epistemic certainty, is under the rational control of an agent. Rational agents who seek to adhere to the norm of epistemic certainty would manifest their adherence by only asserting propositions of which they were subjectively certain. Instead of governing the act of assertion, the norm of subjective certainty would emerge from rational requirements on an agent's adherence to the norm of epistemic certainty. (p. 52)

The idea seems to be that in trying to conform to an epistemic certainty norm on assertion, rational people will end up conforming to a surety norm on assertion. Maybe this is true, but if so it needs to be explained. In particular, it is unclear what role epistemic certainty is playing here – would being rational and merely trying to assert only what is *true* also entail only asserting what one is sure of?

the surety norm on assertion is normatively parasitic on the epistemic certainty norm on assertion, in virtue of the following surety norm on epistemic certainty: “Be sure that p if p is epistemically certain for you”. But even if Beddor is right that there is something non-ideal about failing to be sure of something that is certain for you, such a norm is not strong enough to explain the surety norm on assertion, since we don’t treat failing to conform to it in anything like the way we treat failing to be sure of what one asserts.

To see this, consider the following example. Holmes has just shared all of his evidence with Watson. Watson is still not sure who committed the crime. But it is certain for him, since if Watson were to ask Holmes “Does this mean it’s certain that the butler did it?”, Holmes could reply “Yes” and speak truly. However, suppose that instead of asking Holmes, Watson for some reason instead leaves the conversation and starts telling people that the butler did it, without being sure that the butler did it. We judge these assertions to be improper, on account of their being insincere (in the sense that Watson represents himself of being sure of that which he isn’t). They involve an additional and more significant impropriety than Watson’s mere failing to be sure of what is certain for him. Beddor’s proposal cannot explain this fact, since it locates the impropriety of Watson’s assertion in his mere failure to be sure of what is certain for him. Cases like these strongly suggest that the primary norm on assertion has to do with being sure rather than with epistemic certainty.¹³

Before moving on, we want to respond to an argument Stanley (2008) gives against the epistemic certainty norm on assertion being (merely) normatively parasitic on the surety norm on assertion, together with an epistemic certainty norm on being sure (“Be sure that p only if p is certain for you”). Like Williamson’s objection to merely normatively parasitic knowledge norms on assertion, Stanley rejects this proposal on the grounds that normatively parasitic epistemic certainty norms are too weak to explain our judgments about the infelicity of asserting things that aren’t certain for you. He first observes that utterances of the form ‘I believe that p , although I don’t know that p ’ are not typically infelicitous. He then appeals to a knowledge norm on “full belief” to argue that, given a knowledge norm on assertion, these assertions would have to be infelicitous if assertions that violated normatively parasitic norms were thereby infelicitous. Since these assertions are not infelicitous, Stanley concludes that parasitic norm-violations are insufficient to explain infelicity judgments, and hence that the above proposal fails to explain the infelicity of assertions like ‘ p , but it’s not certain that p ’.

We agree with Stanley that it is important to be attentive not only to the existence of norms but also to their strength. Indeed, we relied on this idea in our

¹³One could reply on behalf of Beddor that we should want a more refined notion of epistemic certainty, one on which non-obvious consequences of epistemic certainties can fail to be epistemically certain. While we are not resistant to such a proposal, the way we would want to develop it would make surety necessary for epistemic certainty, in effect collapsing the present proposal and the knowing-for-sure proposal discussed above. A *prima facie* challenge for this proposal explaining cases where we know we are unsure whether p , yet still sincerely and felicitously pose questions of the form ‘Is it certain that p ?’ (as we imagined Watson doing with Holmes).

criticisms of Beddor above. But in this case we think Stanley’s argument relies on an improper diagnosis of the data. The best explanation of why utterances of the form ‘I believe that p , although I don’t know that p ’ tend to be felicitous is not that we may felicitously violate parasitic norms. It is that there isn’t a knowledge norm on the attitude expressed by ‘believe’ in ordinary English. We will return to this point in detail in section 7.

4 Being unsure and suspending judgment

In this section we’ll expand on the discussion of the attitude of *being unsure* and the norms on inquiry that involve it. We will also highlight connections between being unsure and some recent work by Jane Friedman (2013, 2017, 2019) on suspension of judgment.

The unsurety norm on inquiry says that one should inquire into whether p only if one is unsure whether p . But what does it take to be unsure whether p ?

It isn’t enough to fail to be sure either way (i.e., to be neither sure that p nor sure that not- p). Spoons are neither sure that hot dogs are sandwiches nor sure that they aren’t, but spoons aren’t unsure whether hot dogs are sandwiches. Nor is being unsure just a matter of being something with a mind that isn’t sure either way. Julius Caesar was neither sure that hot dogs are sandwiches nor sure that they aren’t, but still he wasn’t unsure whether hot dogs are sandwiches.

We propose that to be unsure whether p is to *understand* the question whether p while neither being sure that p nor being sure that not- p .¹⁴ We won’t have much to say about the admittedly vague notion of “understanding a question”, though we intend to use it in such a way that we can correctly say that Caesar did not understand the question of whether hot dogs are sandwiches. But we will not be assuming that understanding a questions requires ever actively considering it.¹⁵ As we’re understanding ‘understanding’, Hania understands the question whether Biden and Trump’s phone numbers end in the same digit, despite having never considered it. So we think it would be correct to say that Hania is unsure whether Biden and Trump’s phone numbers end in the same digit, since she is neither sure that they do nor sure that they don’t.

We do not deny that it can sound a bit odd to say ‘Hania is unsure whether Biden and Trump’s phone numbers end in the same digit’. But to our ears it sounds no odder than saying ‘Hania isn’t sure whether Biden and Trump’s phone numbers end in the same digit’. Assuming that to be sure whether p is to either be sure that p or sure that not- p (and Hania is neither), this oddity must have a pragmatic explanation, which can then be applied equally to being unsure.

¹⁴Cf. Goodman and Lederman (2021), who argue that identifying being unsure with understanding without being sure best explains the incompatibility of being unsure and being sure even in cases of Hesperus/Phosphorus-style identity confusion.

¹⁵In this respect it is like a common way of understanding the notion of “awareness” in epistemic logic; see Schipper (2015).

Indeed, we would like to go further. Suppose Hania has never considered the question of whether her phone number and her birthday end in the same digit. It sounds odd to say ‘Hania is sure that her phone number and birthday end in the same digit’. But supposing they in fact do share the same final digit and that Hania knows both numbers by heart, we think this utterance expresses a truth. In our view, being sure is a sufficiently dispositional matter that the fact that Hania’s phone number and birthday end in the same digit should be counted among the things Hania is sure of (and thus understands), despite the fact that she has never explicitly considered the question. In saying the same about being unsure, we are saying that it is a member of a family of similarly dispositional mental states.

Let us now return to norms on inquiry. As Friedman’s inspector Morse example illustrates, there is something defective about inquiring into whether p while failing to be unsure whether p .¹⁶ This is true even for Friedman’s (2017) expansive understanding of inquiry into whether p , which includes not just the activities of evidence gathering and conscious reasoning about p , but also the attitudes of being curious whether p or wondering whether p . Again, if Morse is sure that he himself committed the murder, then he should not (and perhaps even cannot) be curious whether he did, or idly wonder whether he did. It might be reasonable for Morse to continue to gather evidence relevant to the question of who committed the murder as part of an inquiry into another related question, such as how widespread the evidence of the murder is, or how justified he is in being sure that he is the murderer. But these would not be inquiries into the question of who committed the murder (at least not in the operative sense of ‘inquiry’).

Given the intuitive connections between being unsure and inquiry (and likewise being sure and assertion), we are surprised by how little philosophers have had to say about being unsure *per se*. More commonly epistemologists talk about attitudes like “being agnostic about” or “uncommitted on” a proposition, or “withholding” or “suspending judgment” on a question.¹⁷ We will now consider some recent claims Friedman has made about these attitudes – which, following her lead, we will lump together and call ‘suspending’ – and assess which if any are true about being unsure.

One prominent claim of Friedman’s (2013) is that suspension is an attitude “in its own right”. By this she means (at least) that suspension is not merely the absence of belief and disbelief, nor merely a kind of higher-order attitude, such as a belief about what one fails to believe. Being unsure satisfies both of these conditions: it implies having mental states (namely understanding), and so can’t be a mere absence of belief (or surety); and it doesn’t require having any higher-order attitudes about one’s mental states.¹⁸

¹⁶Here and throughout we assume that one can only inquire into questions one understands, and thus that: if one is inquiring into whether p , then if one isn’t unsure whether p , then one is either sure that p or sure that not- p .

¹⁷For a recent and especially clear example of this pattern, see McGrath (2021).

¹⁸One might worry that, on our view, being unsure is an attitude in its own right only in a very thin sense, since the only positive attitude it entails is mere understanding, and, against

Friedman (2019) has also argued that one should not inquire into questions unless one suspends judgment about them, which she contrasts with having beliefs about their answers. Insofar as what she intends by ‘belief’ is weaker than being sure – an issue we will discuss in section 7 – we are not committed to this claim. In fact we reject it, since we think that anything falling short of surety can suffice for reasonable inquiry. Here we find Friedman’s (2017) expansive conception of inquiry helpful: even if one is confident enough that p for further evidence gathering or deliberation to be ill-advised, one may still reasonably remain curious whether p , or idly wonder whether p , as long as one remains unsure whether p .

More ambitiously, Friedman (2017) has suggested that suspension is a limiting case of inquiry: anyone who suspends on a question is, in some minimal sense, inquiring into that question. We think the parallel claim about being unsure is clearly false: there are plenty of pointless questions we understand but have never considered and would never have any interest in if we did, despite being unsure of their answer. We are not inquiring into whether there are an even number of stars in the galaxy, even in the expansive sense that encompasses curiosity and idle wondering. However, since this last suggestion of Friedman’s is not part of the orthodox way of thinking about suspension, we don’t think it is much evidence that epistemologists writing about suspension are concerned with an attitude more demanding than merely being unsure.¹⁹

At the heart of Friedman’s picture of suspension is a three-way taxonomy of non-probabilistic doxastic attitudes: someone who understands the question of whether p either believes p , disbelieves p , or suspends judgment on p . We reject this familiar trichotomous perspective. While there is an important three-way distinction between being sure that p , being sure that not- p , and being unsure whether p , the latter admits of further categorical distinctions. In particular, being unsure whether p is compatible with nevertheless *thinking* that p (or thinking that not- p , or having no opinion). Let us now turn to these attitudes.

5 Thinking and guessing

Thinking that p doesn’t require being sure that p . For example, someone might think that it will rain in the next month without being sure that it will. Recently Moss (2019) has challenged this claim, arguing that although these descriptions

the backdrop of understanding, being unsure makes no positive contribution to one’s mental life in the way that other propositional attitudes do. In reply: while this may be so, we intend our account of being unsure (in terms of understanding and being sure) to be neutral on which of these attitudes is more basic than the others. Perhaps being sure and being unsure are equally basic, and understanding should be understood disjunctively, as being either unsure or sure. (Compare being unhappy, which is arguably as basic as being happy.) We will not explore this question of priority further here; see Williamson (2000, ch. 3) for more on how to make sense of which propositional attitudes are more basic than others.

¹⁹We are also not opposed to the idea that not much needs to be added to being unsure in order for it to count as inquiry in a very weak but not unnatural sense. Perhaps considering a question one is unsure about, even in the absence of any “deep” curiosity or speculation, is a limiting case of inquiry.

are natural, they are strictly speaking false: in describing someone unsure of whether it will rain as nevertheless thinking that it will rain, we are merely speaking loosely. We will not here directly rebut this position. While we understand how someone in the grip of the trichotomous picture mentioned at the end of the last section could reason themselves into accepting it, to our mind the correct reaction is simply to reject the trichotomous picture.²⁰

Still, for all we have said so far, cases of thinking without being sure are rare.²¹ Perhaps they only arise in special circumstances that make being sure unusually demanding – say, when we are thinking about the open future – or thinking unusually easy – say, when it comes along for free via a memory trace (as with Jones from section 2). But this is not our view. We think that even for the kinds of subject matters where we are often sure of many things and where what we think is driven by reflective deliberation, we often think much more than we are sure of. Here we are in agreement with a number of recent authors who have defended this descriptive claim: e.g., Hawthorne et al. (2016), Dorst (2019a), Rothschild (2020), Dorst and Mandelkern (forthcoming), and Holguín (forthcoming).

To illustrate, suppose you’ve recently taken a difficult multiple-choice exam. The instructor is reviewing some particular question, for which your answer was (c). Consider two questions the instructor might ask: (i) ‘What made you sure the answer was (c)?’; (ii) ‘What made you think the answer was (c)?’. There is a striking contrast here. (i) sounds very strange: it’s a hard test, so the instructor should expect that students are regularly choosing between answers without being sure which is correct. (ii), by contrast, sounds perfectly normal. There is nothing overly presumptuous about assuming that a student who answered (c) thought the answer to the relevant question was (c). Indeed, if the instructor were to ask ‘Why did you put (c)?’, an answer along the lines of ‘Because I thought it was the correct answer’ would seem unhelpful, if not outright petulant. In normal circumstances the instructor takes for granted that the student put the answer they did because they thought it was the right answer. They wanted to know *why* the student thought it was the right answer, not *whether* the student thought it was right the answer.

Now, it isn’t impossible for the question ‘What made you think the answer

²⁰Less radically, Nagel (forthcoming) and Williamson (forthcoming) claim that first-personal ‘I think that p ’-utterances should typically be understood as a kind of hedged assertion of p , rather than as an assertion about one’s mental state. While we don’t deny that these utterances function as hedged assertions, we think this is not only compatible with *but would be explained by* their being self-ascriptions of a mental state: after all, if one asserts that they think that p rather than p itself, this implicates that they aren’t in a position to assert p , and hence that they only think it. Moreover, the inference from ‘ S said “I think that p ”, and they were neither lying nor mistaken about what they think’ to ‘ S thinks that p ’ seems impeccable in a way that is difficult to explain on the assumption that people who say ‘I think that p ’ aren’t typically asserting that they think that p . In any event, most of the examples below involve non-first-personal ‘think’-ascriptions.

²¹We assume that thinking that p implies understanding whether p and not being sure that not- p , so cases of thinking without being sure are cases of thinking while being unsure. We will also assume that being sure implies thinking: witness the absurdity of ‘They might not think that p , but they’re sure that p .’

was (c)?' to have a false presupposition. The student could felicitously reply 'I didn't *think* it was (c) – I was running out of time and filled in the answer at random.' But there seems to be a default assumption that people's answers on multiple choice tests, when they have time to deliberate and are trying to answer correctly, will also be propositions that they think are the correct answers to the questions they are answering.

There is a similar pattern in prospective rather than retrospective cases. Suppose you are advising someone on how to use a voucher at the racetrack, and they don't know anything about the conditions under which the bet they are placing will pay out. We might reasonably advise them: 'Bet on the horse you think will win'. By contrast, 'Bet on the horse you are sure will win' is bizarre advice, since following it would require being sure which horse will win, which one shouldn't be.

So one can properly think that p even when one fails to know (or be sure) that p . Indeed, one can properly think that p even when one *knows* that one fails to know that p . In this respect thinking is different from being sure.

Thinking also seems to be highly *question-sensitive*.²² For example, suppose you think it's 40% likely that Djokovic will win Wimbledon, and don't think anyone else is more than 5% likely to win. Now compare the range of permissible answers to the question 'Who do you think will win Wimbledon?' versus 'Do you think Djokovic will win Wimbledon?' (or even better, 'Who do you think will win Wimbledon: Djokovic or the field?'). To our ears, it is perfectly natural to answer the first question with 'I think Djokovic will win' and to answer the second question with 'No, I think someone else will', at least if only one of the questions is asked.²³ In answering this way, it does not follow that you have an incoherent doxastic state. Nor does it follow that you maintain coherence only by changing your mind between the two questions about whether Djokovic will win Wimbledon. You just happen to think that Djokovic will win relative to the former question but not relative to the latter. Similar points go through for multiple-choice exams ('What do you think the answer to this question is?' versus 'Do you think the answer to this question is (c), or something else?') and horse races ('Who do you think will win?' versus 'Do you think the winner will be the favorite or one of the underdogs?').

With this in mind, we'll say that a person is *opinionated about* a question Q just in case, for some p , they think that p is the true answer to Q . We've argued that the epistemic requirements on having an opinion about a question are significantly weaker than the epistemic requirements on being sure of one of its answers. But are there any epistemic constraints on one's opinions? We think there are. In particular, we think one shouldn't think that p is the answer to Q if there is another answer p^* that one thinks is more likely than p .²⁴ But

²²See Holguín (forthcoming) and Dorst and Mandelkern (forthcominga) for further discussion of the question-sensitivity of thinking.

²³Asking the questions in immediate succession tends to force the second answer to conform to the first.

²⁴Throughout we use 'think likely' and related constructions in a way that should be understood as invoking something like philosophers' notion of credence/subjective probability.

as long as p is your *best guess* to Q (in the sense of being an answer to Q that you think is at least as likely as any other), then it is *okay* to think that p is the answer to Q (in the sense that it doesn't violate any norms on thinking).

Putting the pieces together, we endorse the following normative principle on thinking:²⁵

Thinking

- (i) Think that p is the true answer to Q only if p your best guess to Q .
- (ii) If it's okay that p be your best guess to Q , then it's okay to think that p is the true answer to Q .

Our notion of being 'okay' is intended to be the permissive dual of the kind requirement associated with conditional imperatives like the norms on assertion and inquiry discussed earlier, as well as (i) above.²⁶

Note that (ii) implies that there is no knowledge norm on thinking, as there is with being sure. We think this is a good result. Again, there seems to be nothing inherently wrong with thinking that it will rain without knowing it, whereas there does clearly seem to be something wrong with being sure it will rain without knowing it. Indeed, on our view there is not even a *truth* norm on thinking that p , since p doesn't have to be true in order for it to be an okay best guess.²⁷ (For ease of exposition, below we will assume that people's best guesses are okay unless otherwise noted, and we will suppress explicit question-relativity when the intended question is clear from context.)

Thinking is a normative principle. It does not imply that, if p is one's (unique) best guess about Q , then one thereby thinks p is the answer to Q . Moreover, given that (ii) is a merely permissive claim, for all we have said no one is ever required to think anything. More to the point, for all we have said there is nothing wrong with failing to think that any particular answer to an exam question is correct, or that any particular athlete will win Wimbledon. If a student denies that they thought the answer they gave was correct while explaining that they gave it because they thought it was at least as likely as any other, or if a tennis fan grants that Djokovic is the favorite but denies thinking that he will win, we don't think they need be doing anything improper. Having a best guess implies neither that you have an opinion nor that you ought to.²⁸

²⁵This is a special case of the more general 'cogency' norm from Holguín (forthcoming). For ease of exposition, we ignore a number of important subtleties. To give one salient example: we are assuming the relevant answers are always complete answers. Allowing the answers to be partial introduces a range of complications.

²⁶We prefer this jargon to "rationally permissible" because we wish to sidestep the debate about whether rationality is normative (cf. Kolodny (2005) and Broome (2013)) while ensuring that (ii) has implications about the *absence* of competing norms (as discussed presently).

²⁷This is not to deny that there is any sense in which thinking might in some sense "aim at truth" (cf. Velleman (2000) and Wedgwood (2002), on truth as the aim of belief); see Dorst and Mandelkern (forthcominga) for discussion, building on Levi (1967).

²⁸See Williamson (forthcoming). This point is not threatened by the fact that one cannot generally say 'I don't think that p '. That fact is explained by the so-called "neg-raising" of 'think', whereby ' S doesn't think that p ' is heard as equivalent to ' S thinks that not- p '; see Rothschild (2020).

That said, there are certain practical contexts where agents typically do form opinions about questions in accordance with their best guesses to those questions – namely guessing contexts. Here we have in mind situations of forced-choice under conditions of uncertainty, as in multiple-choice exams, trivia questions, game shows, and the like. Our view is that a *typical* way of guessing an answer to a question is to start by forming an opinion as to the question’s answer (i.e. to come to think that one of its answers is true), and then to offer one’s opinion as one’s guess. That is:

Guessing

People typically guess answers to questions by becoming opinionated about the questions and then guessing the answers they think are true.

We think the conversational patterns discussed earlier provide strong support for this claim. Again, it is perfectly natural to ask an examinee or a game-show contestant for their guess by asking them what they think the answer to a given question is. Similarly, upon hearing a person make a guess, it is perfectly natural to ask them why they thought the content of their guess was true. This isn’t to say that people’s (not-insincere) guesses are invariably things they think: for example, under time pressure one might pick an answer at random without thinking it is true. We are only claiming that normal people’s guesses are typically also things they think are true.

6 Thinking and deciding

We have argued that being unsure whether p is no bar to reasonably thinking that p . We have also argued that, in general, it is permissible, upon considering a question, to come to think that one of its answers is true, provided this answer is one that you think is no less likely than any other. But so far we have only implicated this way of coming to think that p in explaining a limited and artificial range of actions – namely, guesses. In this section we will argue that this way of becoming opinionated about questions is in fact a central feature of decision making under conditions of uncertainty.

Our starting point is the idea that deliberation is a kind of normative inquiry, where we decide what to do by thinking about what we *should* do. Put this way, the idea hardly seems controversial, let alone novel. Our goal is to make it more precise, to illustrate why it is a non-trivial psychological hypothesis that depends for its plausibility on being framed in terms of what people think is true, rather than something more demanding like what they are sure is true.

Here is the claim we aim to defend:

Deciding

People typically make decisions by becoming opinionated about what they should do, and then intending to do what they think they should do.

This claim is parallel to Guessing, which we defended in the previous section. It admits the same kind of exceptions: just as we sometimes guess under time

pressure, or unreflectively, thereby guessing something that is not what we think is the answer to the question we happen to be answering, so too people sometimes make decisions under time pressure, or unreflectively, thereby intending to do something without thinking that it is what they should do.²⁹ Perhaps symmetric choice situations, like deciding how to act in a game of rock paper scissors, present another exception: you decide to do something arguably without thinking you should do it, because you think there isn't any action that you positively *should* do.³⁰

Deciding is a descriptive generalization, not a normative one. It should not be confused with the claim that, if you are rational in thinking that you should do something, then you are rationally permitted (or required) to intend to do it.

²⁹We will mostly use 'decide' as synonymous with 'come to intend', but there are subtle differences. For example, if you try to do something that you aren't sure you'll succeed in doing (like making a three-point shot in basketball), it is natural to say you intended to make the shot but not that you decided to make the shot. (We owe this observation to BLINDED.) Deciding to do something, unlike merely intending to do it, seems to involve being sure that you will do it. While this distinction has important implications for cognitivism about intention and related views, it is orthogonal to the present discussion, which is concerned with uncertainty about what you should do but not uncertainty about what you will do.

³⁰This isn't the only way of thinking about symmetric choice situations. Consider someone who knows that a coin is ever so slightly biased towards heads. Are they always going to guess that it will land heads when tossing the coin to determine who will go first in a game? Not necessarily. If they don't, how should we explain their guessing behavior? It doesn't seem implausible to us that, despite what they know about the bias of the coin, nevertheless they still think (irrationally) that the coin will land tails on occasions where they guess it will. If this kind of thinking can guide behavior when it runs counter the known biased chances, surely it can also guide behavior in unbiased cases too. So perhaps in symmetric choice situations people follow Deciding by irrationally (but pretty harmlessly) thinking, at least momentarily, that they should do what they then decide to do.

What about people who have a stable disposition to guess heads if they know the coin is even slightly biased – what should we say about them in symmetric choice situations? One might be tempted to say that, in such cases, their decision is based not on thinking they *should* play rock (for example), but is instead based on thinking they *may* play rock. (Thanks to BLINDED for this suggestion.) But we aren't so sure. Suppose John explains why he decided to take two desserts at a buffet by saying 'because I could'. This strikes us as enthymematic: in light of the fact that he could (in the sense of being permitted) take two desserts, and the fact that he wanted them, he came to think that I should take two desserts, and decided accordingly, in keeping with Deciding. And while rock-paper-scissors decisions don't feel as obviously enthymematic to us, a similar diagnosis might apply. For note that in many situations there are costs to continued deliberations, in light of which one will do best in the long run to decide to do what one thinks is the best of the options one has considered, rather than consider further options, at least if one is satisfied that the option is good enough. In such cases, the fact that an option you think is good enough has come to mind arguably makes the option not merely permissible but something one positively should do, given the inherent costs and diminishing returns of further deliberation. So perhaps people have a tendency to consider a range of options, and, once enough of them are on the table, come to think of one of the options that it is what they should do, and then decide to do it in keeping with Deciding. For some fascinating recent work on the role of considering a limited range of options that come to mind in decision making, see Morris et al. (ms).

Yet another way of reconciling Deciding with symmetric choice situations is to deny that these cases are "decisions" at all (in the relevant sense). Perhaps the agent only "decides" to throw one of rock, paper, or scissors, and further, non-deliberative process selects one of these actions for execution. We aren't partial to this idea ourselves, but see Ullmann-Margalit and Morgenbesser (1977) for an influential proposal along these lines.

(In other words, it is not a *decision theory*: it doesn't give sufficient conditions on rational decisions in terms of one's rational opinions.) It also should not be confused with a 'wide-scope' rational requirement, according to which a person is fully rational only if everything they think they should do is something that they intend to do. (In other words, it should not be confused with Broome's (2013) much discussed 'Enkratic Requirement', which he takes to be definitive of the operative "deliberative" notion of 'ought'.) The normativity in Deciding is in the content of what one thinks; the principle says nothing about what people should think or decide.³¹

The general idea that deliberation is guided by the question 'What should I do?', and then settled by answering that question, is one that we hope is intuitive enough and in any event familiar. As for the specific principle Deciding, it can be motivated in the same way that we motivated Guessing. Suppose you decide to ϕ . Someone might then question your decision by asking 'What makes you think that ϕ is what you should do?'. While it isn't impossible for you to protest against the presupposition of the question – you might claim, for example, that you made the decision under time pressure, or impulsively, without thinking that it was what you should do – such protestations are atypical. By contrast, suppose someone instead questioned your decision by asking 'What makes you sure that ϕ is what you should do?'. This question sounds presumptuous in a way that the earlier 'thinks' question does not. The natural way to protest against its presupposition would be to say something like 'I'm not sure that it's what I should do, but it's what I think I should do and I have to do something', thereby acknowledging that you at least think you should do it. The principle Deciding naturally explains why the 'think' query is felicitous. The fact that a parallel 'sure' query sounds presumptuous suggests that, by contrast, it is not atypical to decide to do things that one is not sure one should do. Moreover, since by Thinking it is okay to become opinionated about any question, including the question of what you should do, conforming to Deciding never requires one to become improperly opinionated.³²

³¹The 'should' in Deciding should be understood in the "deliberative" sense – see Thomason (1981), Wedgwood (2007), Schroeder (2011), and Broome (2013) for more on what distinguishes this from other interpretations of 'should'. (The literature refers to it as the deliberative 'ought', but we here assume that 'ought to' and 'should' are synonyms in English.) One relevant point for our discussion below is that uncertainty about what one (deliberatively) should do is different from uncertainty about which of one's available actions would have the best consequences if one did it. This is for reasons to do with Parfit's famous miners puzzle: faced with the decision of whether to flood either shaft, agents can be sure that they should flood neither shaft while also being sure that flooding neither shaft is *not* what would have the best consequences. See Kolodny and MacFarlane (2010) and Cariani et al. (2013) for further discussion.

³²Here are two other considerations in favor of Deciding. The first has to do with how we seek advice; cf. Schroeder (2011). Suppose you're having trouble making a decision and you go to a friend for help. The most natural way to solicit advice is to ask 'What should I do?'. If your friend replies that they aren't sure, you might still follow up 'Okay, but what do you *think* I should do?'. If your friend says they think you should do one thing, but then you decide not to do it, then, intuitively, you haven't followed their advice, and your friend might protest by asking why you asked them for advice in the first place. Deciding helps to make sense of what is going on here.

If Deciding is true (and the parallel principle formulated in terms of being sure rather than in terms of thinking is false), then there is an important contrast between practical deliberation and theoretical inquiry. As discussed in section 2, inquiry into a question is generally permissible so long as the inquirer is not sure of any answer to the question they’re inquiring into. So it takes surety to close inquiry. By contrast, merely thinking that one should do something typically suffices for intending to do it. In this sense, it only takes thinking to close practical deliberation. While some might see a tension here, we think it is a virtue of our framework that it allows us to distinguish the way in which intention-formation is the aim of practical deliberation from the way that surety is the end of inquiry. As Bratman (1987) and others have emphasized, we need to be able to form intentions so that we can act in the absence of further deliberation. But having formed an intention to do something does not mean that it is inappropriate to continue inquiring about whether it is what you should do. In many real-world situations we have a limited amount of time to deliberate before acting. A typical and cognitively healthy way of approaching such deliberations is to quickly make a best guess about what you should do and then, in conformity with Deciding, form an intention to do it, so that you will be able to act somewhat intelligently when the time comes. Then, assuming the decision is important enough, it will often be worth inquiring further into what you should do for as much time as it takes to either come to be sure of what you should do, or to recognize that further deliberations aren’t going to change your opinion.

7 Whither “belief”?

So far we have focused on the attitudes of thinking and being sure, remaining neutral on how these are related to belief, as well as to other attitudes like suspecting, being confident, and being nearly certain. We would like to close with a picture of how we see the landscape of doxastic attitudes. A full elaboration and defense of this picture is beyond the scope of this paper; see *BLINDED*. But in assessing the claims about thinking and being sure defended here, we think it is helpful to see in outline how those attitudes can be situated with respect to related ones, and how they have a distinguished position among them.

We should start by addressing the elephant in the room. In our view ‘believe’

The second consideration in favor of Deciding begins with the truism that we deliberate about decisions by considering what to do. How should we think about the meaning of infinitival questions like ‘what to do’ as complements to psychological verbs? The standard treatment in linguistics is that infinitival constructions like ‘A is considering what to do’ have the logical form ‘A is considering what PRO X do’ where PRO is a phonologically null pronoun that is anaphoric on the subject A, and X stands for some kind of modal auxiliary, such as ‘should’ or ‘will’. This argument isn’t decisive, since it doesn’t establish that the modal auxiliary in question is in fact the deliberative ‘should’. But that is a natural hypothesis, especially considering the fact that other options seem strange. For example, a common theme in action theory is that decision making is not a matter of deliberating about what one is going to do (at least, not in the way that one might try to predict another person’s actions). So we think this argument lends some further support to Deciding.

in ordinary English is synonymous with ‘think’. That is to say: to believe that p is to think that p . We have nothing new to say on this point that hasn’t already been said by (e.g.) Hawthorne et al. (2016) and Rothschild (2020), and we refer the reader to their arguments.

What about “full belief” or “outright belief” as these terms are used in contemporary epistemology? In our view, it is usually most productive to see such discussion as concerned with the ordinary notion of being sure. This is mainly because full belief tends to be characterized as the attitude we take towards propositions we are disposed to sincerely assert and not to inquire into, and, if we are right, then being sure is governed by parallel norms, while thinking is not. We will now defend this hypothesis against four objections.

One objection is that epistemologists are competent English speakers and take themselves to mean by ‘belief’ what ordinary speakers mean. We are unmoved by this objection. Philosophers have a long history of taking ordinary notions and imbuing them with technical meanings – see, e.g., ‘truth’, ‘meaning’, ‘reference’, ‘exist’, ‘real’, ‘inference’, etc. Second, philosophers explicitly use jargoned qualifiers like ‘full’ and ‘outright’ to clarify the operative notion; they are not merely deferring to ordinary usage. Finally, the error we are attributing to epistemologists who equate full belief with the ordinary meaning of ‘believe’ is not that of inventing a notion and claiming a kind of importance for it that one might be suspicious of if English hadn’t had the good sense to lexicalize it. For English *has* had the good sense to lexicalize the important notion in question – namely, with ‘be sure’.

A different objection is that full belief is not as demanding as being sure. But the idea that being sure is extremely demanding is, we think, driven by confusing it with something else, such as Cartesian certainty (i.e., not merely being sure, but also not disposed to change your mind in response to future counter-evidence) or subjective probability 1 (which, if the operative notion of probability patterns with known objective chances, may not apply to any ordinary propositions about the future we are willing to assert or treat inquiry into as settled). We have shown that the ordinary notion of being sure implies neither of these things.

A third objection is that the epistemologist’s notion of full belief is necessary for knowledge, but being sure isn’t. We find this objection unpersuasive too. We don’t deny that belief in the ordinary sense (i.e., thinking) is necessary for knowledge. But the same cannot be stipulated for the quasi-technical term ‘full belief’. And given that one can know a proposition without being in a position to assert it and without being prohibited from inquiring into its truth, it’s hard to see how an attitude that satisfies the theoretical roles of full belief could be a necessary condition on knowledge.

A fourth objection concedes that full belief is more demanding than the attitude expressed by the ordinary use of ‘belief’, but questions why we should think it is as demanding as being sure. Perhaps it is something in between. In reply to this objection, we want to begin by agreeing that this is a perfectly coherent hypothesis. In particular, we think there is a linear ordering of attitudes, between thinking and being sure, ordered according to how demanding they are.

Explaining the details of how we are thinking about what this demandingness ordering amounts to would take us too far afield. But the sorts of examples of intermediate attitudes we have in mind are things like *being confident* that p and *being pretty sure* that p . If these examples are any indication, then attitudes intermediate in strength between thinking and being sure are not fit to play the role of full belief. After all, provided that one isn't sure that it will rain, none of being pretty sure, confident, or even *very* confident that it will rain prevents one from reasonably inquiring into whether it will rain. Nor do any of these attitudes seem to suffice for being in a position to assert that it will rain.

This leads us to a positive characterization of being sure on the spectrum of ever increasing demandingness: *being sure is the least demanding doxastic attitude subject to a knowledge norm*. As we saw in section 1, a knowledge norm on being sure is needed to explain the derivative knowledge norm on assertion via the surety norm on assertion. This knowledge norm on being sure also explains the infelicity of 'I'm sure that p , but I don't know that p ', which is also not explained by a surety norm alone. These sentences are instructive because parallel assertions of the form 'I am pretty sure/confident/very confident that p , but I don't know that p ' are not infelicitous in the same way.

This proposal – that being sure is the least demanding doxastic attitude subject to a knowledge norm – is reminiscent of Williamson's (2000, ch. 1) characterization of knowledge as the most general factive mental state. It is not an analysis of being sure in terms of other attitudes, but rather a way of highlighting the significance of being sure in terms of its unique position in relation to other attitudes. It thereby allows a natural response to a potential worry about the epistemological significance of being sure – namely that either being sure is overly demanding, because to be sure is to be as sure as one can possibly be, or it is uninteresting, because it corresponds to some arbitrary cutoff in a continuum of possible levels of being sure.

Let us now turn to the other direction. Is thinking that p a distinguished attitude among the doxastic attitudes, either in terms of knowledge or in some other way? We think it is distinguished in being the weakest attitude on the spectrum that includes it, being confident, being very confident, being pretty sure, being sure, being very sure, and so on.³³ Again, the details of this picture are beyond the scope of this paper. But very roughly, all of these attitudes add something to thinking – namely that the proposition in question has a high enough *plausibility* (for which having the corresponding subjective probability is necessary but not sufficient). So thinking that p is the weakest such attitude. This fits with the fact that, although 'confident' and 'sure' are gradeable expressions, 'think' and 'believe' are not (the philosophers' neologism 'degrees of belief' notwithstanding).

Is thinking distinguished in terms of its relation to knowledge in any way? It is plausibly necessary for knowledge.³⁴ And it can constitute knowledge in the

³³For this reason we agree with Hawthorne et al. (2016) that suspecting that p is no weaker than thinking that p .

³⁴This claim may need to be restricted to certain kinds of knowledge; e.g., one might think that perceiving that p is a way of knowing that p compatible with consistently thinking that

absence of being sure. But for all we have said here, it rarely does so, or at least only does so when the special conditions exemplified in cases of perceiving or remembering without being sure obtain. Crucially, the permissibility of guessing by making up ones mind shows that thinking does not, in any interesting sense, aim at knowledge, since mere guesses, even if true, are paradigms of beliefs that fail to amount to knowledge.³⁵ We see it as an interesting question for further work whether there is any attitude intermediate between thinking and being sure that is necessary for knowledge.

It bears emphasis that the sense in which thinking is a fairly undemanding attitude is primarily normative rather than descriptive. One is always *permitted* to be such that, for some answer to a given question, one thinks that it is the true answer to that question (provided that answer is the one thinks is most likely to be true). But for all we have said no level of subjective probability implies thinking (cf. Holguín (forthcoming) and Williamson (forthcoming)). Perhaps one can have subjective probability 1 that a dart won't land at a certain point on a continuum without thinking (relative to the question of whether it will) that it won't. On the other extreme, though, we have placed no subjective-probability requirements on thinking. Perhaps one can have subjective probability 0 that a dart will land at a certain point while thinking (relative to the question of which point it will land on) that it will land there, this being a best guess.

The picture that emerges is one where subjective probabilities crucially constrain the binary attitudes of thinking and being sure. While these attitudes fall on a spectrum, and the attitudes on this spectrum may be associated with probabilities in a certain way, it is not in the familiar Lockean way, where these attitudes demand merely clearing a certain threshold of subjective probability. Indeed, subjective probability thresholds are orthogonal to this family of attitudes, in the sense that the weakest such attitude (thinking) is for all we have said neither incompatible with minimal subjective probability nor entailed by maximal subjective probability. Similarly, being sure and other attitudes on the spectrum are not entailed by maximal subjective probability, since they require thinking. What distinguishes being sure from the other attitudes in this family is its epistemic role: it aims at knowledge, and epistemologists theorizing about "belief" have, perhaps unwittingly, been aiming at it all along.

References

A. J. Ayer. *The Problem of Knowledge*. New York: St Martin's Press, 1956.

not-*p*. Note also that the question sensitivity of thinking would then induce a corresponding question-sensitivity in knowledge. Goodman and Salow (2021) develop a framework for thinking about question sensitivity of knowledge and belief marching in step; while we think their notion of 'belief' is best understood as being sure, their framework deploys probabilistic orderings of questions' answers in a way that would harmonize easily with the theory of thinking in Holguín (forthcoming). See BLINDED.

³⁵Though see Holguín (2021) for arguments to the contrary.

- Andrew Bacon. Giving your knowledge half a chance. *Philosophical Studies*, (2):1–25, 2014. doi: 10.1007/s11098-013-0276-6.
- Bob Beddor. New work for certainty. *Philosophers' Imprint*, 20(8), 2020.
- Bob Beddor. Fallibilism and the aim of inquiry. ms.
- Matthew A. Benton. Dubious objections from iterated conjunctions. *Philosophical Studies*, 162(2):355–358, 2013.
- Michael Bratman. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press, 1987.
- John Broome. *Rationality Through Reasoning*. Oxford: Wiley-Blackwell, 2013.
- Fabrizio Cariani, Magdalena Kaufmann, and Stefan Kaufmann. Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy*, 36(3):225–259, 2013.
- Stewart Cohen and Juan Comesaña. Williamson on Gettier cases and epistemic logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):15–29, 2013.
- Cian Dorr, Jeremy Goodman, and John Hawthorne. Knowing against the odds. *Philosophical Studies*, 170(2):277–287, 2014. doi: 10.1007/s11098-013-0212-9.
- Kevin Dorst. Lockeans maximize expected accuracy. *Mind*, 128(509):175–211, 2019a.
- Kevin Dorst. Abominable KK failures. *Mind*, 128(512):1227–1259, 2019b.
- Kevin Dorst and Matthew Mandelkern. Good guesses. *Philosophy and Phenomenological Research*, forthcominga.
- Kevin Dorst and Matthew Mandelkern. Assertion is weak. *Philosophers' Imprint*, forthcomingb.
- Jane Friedman. Suspended judgment. *Philosophical Studies*, 162(2):165–181, 2013.
- Jane Friedman. Why suspend judging? *Noûs*, 51(2):302–326, 2017.
- Jane Friedman. Inquiry and belief. *Noûs*, 53(2):296–315, 2019. doi: 10.1111/nous.12222.
- Jeremy Goodman and Harvey Lederman. Perspectivism. *Noûs*, 55(3):623–648, 2021.
- Jeremy Goodman and Bernhard Salow. Knowledge from probability. In Joseph Y. Halpern and Andres Perea, editors, *Theoretical Aspects Theoretical Aspects of Rationality and Knowledge 2021 (TARK 2021)*, pages 171–186, 2021.

- Daniel Greco. Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3):656–673, 2015a.
- Daniel Greco. Iteration principles in epistemology I: Arguments for. *Philosophy Compass*, 10(11):754–764, 2015b.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173(5):1393–1404, 2016.
- Ben Holguín. Knowledge by constraint. *Philosophical Perspectives*, 25:275–302, 2021.
- Ben Holguín. Thinking, guessing, and believing. *Philosophers’ Imprint*, forthcoming.
- Niko Kolodny. Why be rational? *Mind*, 114(455):509–563, 2005.
- Niko Kolodny and John MacFarlane. Ifs and oughts. *Journal of Philosophy*, 107(3):115–143, 2010.
- Isaac Levi. *Gambling with truth: An essay on induction and the aims of science*. Cambridge, MA: MIT Press, 1967.
- Matthew McGrath. Being neutral: Agnosticism, inquiry and the suspension of judgment. *Nous*, 55(2):463–484, 2021. doi: 10.1111/nous.12323.
- G. E. Moore. A reply to my critics. In Paul Arthur Schilpp, editor, *The Philosophy of G. E. Moore*. Evanston, IL: Northwestern University Press, 1942.
- Adam Morris, Jonathan Phillips, Karen Huang, and Fiery Cushman. Generating options and choosing between them rely on distinct forms of value representation, ms. URL <https://psyarxiv.com/2ayw3/>.
- Sarah Moss. Full belief and loose speech. *Philosophical Review*, 128(3):255–291, 2019.
- Jennifer Nagel. The psychological dimension of the lottery paradox. In Igor Douven, editor, *The Lottery Paradox*. Cambridge University Press, forthcoming.
- Colin Radford. Knowledge—by examples. *Analysis*, 27(1):1–11, 1966.
- Daniel Rothschild. What it takes to believe. *Philosophical Studies*, 177(5):1345–1362, 2020.
- Burkhard K. Schipper. Awareness. In Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Epistemic Logic*, pages 77–146. College Publications, 2015.
- Mark Schroeder. Ought, agents, and actions. *Philosophical Review*, 120(1):1–41, 2011.

- Jason Stanley. Knowledge and certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- Richmond H. Thomason. Deontic logic as founded on tense logic. In Risto Hilpinen, editor, *New Studies in Deontic Logic*, pages 165–176. 1981.
- Edna Ullmann-Margalit and Sidney Morgenbesser. Picking and choosing. *Social Research: An International Quarterly*, 44(4):757–785, 1977.
- Peter Unger. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press, 1975.
- David Velleman. On the aim of belief. In David Velleman, editor, *The Possibility of Practical Reason*, pages 244–81. Oxford: Oxford University Press, 2000.
- Ralph Wedgwood. The aim of belief. *Philosophical Perspectives*, 16:267–97, 2002.
- Ralph Wedgwood. *The Nature of Normativity*. Oxford University Press, 2007.
- Timothy Williamson. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.
- Timothy Williamson. Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *Philosophical Quarterly*, 55(219):213–235, 2005.
- Timothy Williamson. Knowledge, credence, and the strength of belief. In Amy Flowerree and Baron Reed, editors, *Towards an Expansive Epistemology: Norms, Action, and the Social World*. London: Routledge, forthcoming.
- Elise Woodard. Why double-check? ms.
- Alex Worsnip. Contextualism and knowledge norms. In Jonathan Jenkins Ichikawa, editor, *The Routledge Handbook of Epistemic Contextualism*, pages 177–189. Routledge, 2017.