



Probabilistic semantics for epistemic modals: Normality assumptions, conditional epistemic spaces and the strength of *must* and *might*

Guillermo Del Pinal¹

Accepted: 23 July 2021 / Published online: 12 October 2021
© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

The epistemic modal auxiliaries *must* and *might* are vehicles for expressing the force with which a proposition follows from some body of evidence or information. Standard approaches model these operators using quantificational modal logic, but probabilistic approaches are becoming increasingly influential. According to a traditional view, *must* is a maximally strong epistemic operator and *might* is a bare possibility one. A competing account—popular amongst proponents of a probabilistic turn—says that, given a body of evidence, *must* ϕ entails that $Pr(\phi)$ is high but non-maximal and *might* ϕ that $Pr(\phi)$ is significantly greater than 0. Drawing on several observations concerning the behavior of *must*, *might* and similar epistemic operators in evidential contexts, deductive inferences, downplaying and retractions scenarios, and expressions of epistemic tension, I argue that those two influential accounts have systematic descriptive shortcomings. To better make sense of their complex behavior, I propose instead a broadly Kratzerian account according to which *must* ϕ entails that $Pr(\phi) = 1$ and *might* ϕ that $Pr(\phi) > 0$, given a body of evidence and a set of normality assumptions about the world. From this perspective, *must* and *might* are vehicles for expressing a common mode of reasoning whereby we draw inferences from specific bits of evidence against a rich set of background assumptions—some of which we represent as defeasible—which capture our general expectations about the world. I will show that the predictions of this Kratzerian account can be substantially refined once it is combined with a specific yet independently motivated ‘grammatical’ approach to the computation of scalar implicatures. Finally, I discuss some implications of these results for more general discussions concerning the empirical and theoretical motivation to adopt a probabilistic semantic framework.

Keywords Epistemic modals · Probabilistic semantics · Must · Might · Evidentials · Scalar implicatures · Probabilistic reasoning · Common-sense reasoning

✉ Guillermo Del Pinal
delpinal@illinois.edu

¹ University of Illinois Urbana-Champaign, Champaign, USA

1 Introduction

Epistemic modals such as *must* and *might* are vehicles for expressing what follows, and with what force, from some specific body of evidence or information. Standard accounts model epistemics using tools from quantificational modal logic, but there is increasing interest in the hypothesis that natural languages interface with, or have access to, a kind of (natural) probabilistic logic. This probabilistic turn has led to attractive accounts of adjectives such as *probably* and *likely*, and various other modal operators, but its extension to the epistemic auxiliaries remains controversial.¹ In addition, while much work has tried to motivate the general move to a probabilistic framework, comparatively less has examined specific accounts of the strength of epistemic *must* and *might*. Yet given the expressive power of probabilistic frameworks, they can be used to model a wide range of specific semantic accounts. So unless we can substantially constrain the possibilities, the move to a probabilistic framework, as such, will provide no special insight into the fascinating interface between our linguistic and our general cognitive capacities to reason from (usually limited) information. This paper motivates several empirical constraints on semantic and probabilistic models of *must* and *might*, and uses them to discriminate amongst three reasonable hypotheses concerning their strength and shed light on central properties of the specific mode of reasoning from information which those epistemic operators seem to target.

Focusing on their epistemic readings, the modal auxiliaries *must* and *might* are part of a complex maze of acceptability patterns. In this paper, I focus on variations of some widely-discussed patterns which, taken together, present a difficult descriptive challenge to any account of these operators. To get a sense for the challenge, consider the uses of *must* in (1)–(4). ‘Deductive conclusions’ and ‘epistemic downplaying’, illustrated in (1) and (2), point to uses of *must* in which it seems to have very strong, perhaps maximal, epistemic force:

- (1) **Deductive conclusions** Elli is looking for her watch. She’s certain it’s in safe-box A or B. Elli says, *My watch is in box A or B. It’s not in A. So it must be in B.* Elli’s *must*-claim in this conclusion seems perfectly adequate, and conveys something stronger than, say, *So it’s almost certain that the watch is in B.* Why is there a contrast between the *must* and the *almost certain* conclusions?
- (2) **Epistemic downplaying** Jasmine checked two websites: both say that it is raining. Jasmine tells Mary, *There’s no point going outside. It must be raining.* It turns out, however, that it isn’t raining. Mary reproaches Jasmine, *You were wrong! It isn’t raining.* Jasmine replies: *I wasn’t wrong! I only said that it MUST be raining.* Why does Jasmine’s reply feel odd or unjustified?

In contrast, ‘evidential uses’ and ‘epistemic tensions’, illustrated in (3) and (4), capture uses of *must* in which it seems to convey some degree of epistemic weakness:

- (3) **Evidential uses** Margot is looking out the window and sees that it is pouring rain. Margot tells Sam, *It must be raining outside.* Why does this assertion sound odd? Suppose instead that Margot and Sam are epistemologists, and

¹ See, e.g., Swanson (2016), Yalcin (2010), Carr (2015), Moss (2018), Cariani (2016), Lassiter (2017).

their exchange takes place in a seminar on skepticism. Why does the same assertion now feel acceptable?

- (4) **Epistemic tensions** Lisa just read on her smartphone that it is raining, but the stakes are high. Lisa tells James, *It must be raining, but it's possible/there's a chance that it isn't*. Why are these epistemic tensions relatively acceptable? Compare, in the same context, the epistemic tension, *It must be raining outside, but it might not be*. Why is this odd?

In short, while deductive conclusions and epistemic downplaying suggest that *must* is a maximally strong epistemic operator, evidential uses and epistemic tensions suggest that it is compatible with a limited degree of uncertainty. In what follows, I will refine and extend these kinds of patterns and use them to discriminate amongst three *prima facie* reasonable accounts of the strength of epistemic *must* and *might*. To make the comparisons between theories perspicuous, I will formulate each account using a uniform probabilistic semantic framework.

The first account is a direct implementation of the traditional modal logic-inspired view according to which *must* is a simple necessity and *might* a bare possibility epistemic operator (von Fintel and Gillies 2010, 2021). In a probabilistic framework, this amounts to the view that, given some body of evidence, *must* ϕ entails that $Pr(\phi) = 1$ while *might* ϕ entails that $Pr(\phi) > 0$. This account deals nicely with basic examples of deductive conclusions (1), epistemic downplaying (2) and, when suitably modified, also with some basic evidential uses (3). However, it is challenged by various kinds of epistemic tensions (4) which suggest that *must* is not a maximal and *might* is not a minimal epistemic operator (Sect. 3). For this and related reasons, I will argue that *must* is better modeled as non-veridical ('human') necessity and *might* as practically relevant possibility, as originally emphasized by Kratzer (1991, 2012).

There are two natural ways of implementing Kratzer's insight in a probabilistic framework. One is to hold that *must* ϕ entails that, given a body of evidence, $Pr(\phi) > \theta_{\text{must}}$, where θ_{must} is a high but less than 1 threshold, while *might* ϕ entails that $Pr(\phi) > 1 - \theta_{\text{must}}$. The suggestion, then, is to lower the threshold for *must* and increase the one for *might* relative to the maximal/minimal thresholds recommended by the traditional approach. This threshold-based account—which tends to be favored by theorists sympathetic to probabilistic approaches to epistemic modals (e.g., Swanson 2006; Lassiter 2016, 2017)—deals nicely with acceptability patterns which suggest that *must* is not a maximally strong and *might* a minimally weak epistemic operator, such as evidential uses (3) and epistemic tensions (4). Still, I will argue that it has a serious shortcoming (Sect. 4). In various contexts—incl. (but not limited to) deductive conclusions (1) and epistemic downplaying (2) (cf. von Fintel and Gillies 2010, 2021)—it incorrectly predicts that *must*-claims should pattern, in terms of acceptability, with overt claims of high but non-maximal certainty or likelihood.

The second implementation of Kratzer's insight, unlike the first one, uses a maximal threshold ($= 1$) for *must* and minimal one (> 0) for *might*, yet weakens *must* and strengthens *might* via the stipulation that both include a default operation which conditionalizes the probability function on a set of contextually relevant 'normality assumptions', conceived of as reasonable yet (usually) defeasible background assumptions or expectations about the world. The goal of this paper is to show that this

conditional non-maximal/minimal account is empirically superior, relative to desiderata (1)–(4) and related variants, to both the traditional maximal/minimal account and the threshold-based probabilistic non-maximal/minimal account (Sects. 3–4). In addition, I will show that, when combined with an independently motivated ‘grammatical’ theory of scalar implicatures, the conditional account can deal with various challenging extensions of our target desiderata (Sect. 5).²

I should mention three clarifications about the scope and limits of this investigation. First, the aim of this paper is not to defend probabilistic over quantificational models of epistemics in general. Accordingly, I focus on cases that help discriminate between specific probabilistic accounts, rather than on cases that motivate the move from quantificational to probabilistic frameworks.³ Second, this paper is about the ‘strength’ of epistemic *must* and *might*. For the most part, I will remain neutral on questions about the nature of their modals bases, such as whether they are really epistemic, or whose knowledge/beliefs they represent (Sect. 2.1).⁴ Third, many influential probabilistic accounts of epistemics have been defended in conjunction with revisionary semantic frameworks, such as expressivism and dynamic semantics. In what follows, I adopt instead a truth-conditional implementation, mainly to present and discuss the competing accounts in a familiar and unified framework (Sect. 2.1). Still, most of the novel observations and results of this paper can inform parallel debates about the ‘strength’ of epistemics in other semantic frameworks.

2 Probabilistic framework and competing accounts of *must* and *might*

This section sets the stage for our discussion by introducing a basic probabilistic semantic framework, some general background assumptions about epistemic modals, and each of the competing accounts of epistemic *must* and *might*, focusing on their core commitments and predictions.

² Current theories of *must* and *might* in the same broad family as the conditional account include Kratzer (1981, 1991), Roberts (2015), Moss (2015, 2018) and Del Pinal and Waldon (2019). Relative to those accounts, this paper attempts to make four novel contributions. (i) Present an argument against maximal/minimal accounts which includes novel data on (embedded) epistemic tensions. (ii) Present cases that can be used to discriminate between the threshold-based and the conditional accounts, ultimately supporting the latter. (iii) Develop a novel account of the interaction between epistemic auxiliaries and covert exhaustification operators which improves the predictions of the conditional account relative to our target desiderata. And (iv) provide empirical support for the unique components/stipulations of the conditional account.

³ For discussions of the motivation to move from quantificational to probabilistic frameworks, see Swanson (2006, 2011, 2016), Portner (2009), Yalcin (2010), Lassiter (2015, 2016, 2017), Moss (2015, 2018), Cariani (2016), Santorio and Romoli (2017), among others.

⁴ For discussions on the nature of ‘epistemic’ modal bases, see DeRose (1991), Egan et al. (2004), Stephenson (2007), Yalcin (2007), Portner (2009), Hacquard (2010), MacFarlane (2011), Dowell (2011), Roberts (2015), among others.

2.1 Probabilistic framework and background assumptions

To implement a truth-conditional probabilistic semantics, assume that interpretations are relativized to functions e from worlds to epistemic probability spaces (Yalcin 2010). An epistemic space is a pair $\langle E, Pr \rangle$ of a set of worlds E and a function Pr . E is a subset of the space of possible worlds W which corresponds to a set of worlds epistemically accessible from the evaluation world.⁵ Pr is a function which assigns to each subset of W a number in $[0, 1]$ satisfying: (i) $Pr(E) = 1$ and (ii) if p and q are disjoint, $Pr(p \cup q) = Pr(p) + Pr(q)$. By relativizing interpretations to epistemic probability spaces, we can spell out the semantics of epistemic terms using Pr . In this framework, $E_{e(w)} = \bigcap f_e(w)$ for each world $w \in W$, where f_e is an epistemic conversational background such that $f_e(w)$ represents what is known in w . Together with assumption (i), this ensures that the probability density is in the space determined by a pure epistemic modal base and that we can model maximal and minimal epistemic operators.⁶ This implementation is flexible with respect to whose evidence is represented by modalized assertions, but in the cases we will focus on, it will usually include the speaker’s evidence.

Suppose that O is a probabilistic propositional (epistemic) operator. Schematically, the entry for O will look as in (5). This says that $O(\phi)$ holds, in w , if the probability of the prejacent ϕ in the relevant probabilistic space $e(w)$ is greater than or equal to the contextually determined threshold for O , where $\theta_c^O \in [0, 1]$.

$$(5) \quad \llbracket O(\phi) \rrbracket^{c,w,e} = 1 \text{ iff } Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e} = 1\}) \geq \theta_c^O$$

Using entries like (5), we can model epistemic modals of different strengths via specific constraints on their contextually determined thresholds. However, we will add one more parameter to our interpretation function which, following Kratzer (1981, 2012), I will argue is used by some (perhaps all) epistemic modals. As a first approximation, this parameter corresponds to a ‘stereotypical conversational background’ g which picks out, at each world w , a set of ‘normality’ assumptions which capture general background expectations about the world (once refined, these functions should also be sensitive to suitable standards as determined by specific discourse contexts). A schematic entry for a probabilistic operator, O^* , that uses both relevant evidence and normality assumptions is presented in (6). Note that $e^{g(w)}$ is obtained from e and conversational background g as follows: if $e(w) = \langle E, Pr \rangle$, then $e^{g(w)}(w) := \langle E^{g(w)}, Pr^{g(w)} \rangle$, where $E^{g(w)}$ is defined as $\bigcap g(w) \cap E$ and $Pr^{g(w)}$ is defined as Pr conditionalized on $\bigcap g(w)$.⁷

⁵ For simplicity, we assume that the set of all possible worlds W is finite.

⁶ To model maximal and minimal epistemics, an additional assumption is needed. For the finite spaces under consideration, this amounts to a kind of principle of generality amongst the worlds compatible with the evidence, namely, for each $w \in W$ and $w' \in E_{e(w)}$, $Pr_{e(w)}(w') > 0$. This ensures that, when modeling epistemics, the corresponding probabilistic operators don’t (implicitly) exclude epistemic possibilities on non-epistemic grounds such as what is simply believed or assumed.

⁷ This is based on a simple adaptation of Yalcin’s (2010) probabilistic version of Kratzer’s semantics for conditionals. The difference is that what is in this case incorporated into the probabilistic space is not an overt antecedent but a contextually provided set of normality assumptions about the world, as determined

$$(6) \quad \llbracket O^*(\phi) \rrbracket^{c,w,e,g} = 1 \text{ iff } Pr_{e^g(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^g(w),g} = 1\}) \geq \theta_c^{O^*}$$

This convention allows us to distinguish ‘pure’ epistemic probabilistic operators from epistemic operators that also incorporate, by default, sets of relevant normality assumptions about the world. Whether this distinction matters for natural languages in general, or for modal auxiliaries in particular, is an open question. Yet it is a distinction that, at this point, we want our framework to capture.

Using this framework, we can now introduce some background assumptions about particular modal operators. In what follows, we will examine cases that involve interactions and comparisons between *must* and *might* and adjectives such as *certain* and *possible* (among others). To use these cases to constrain theories of the auxiliaries, we need to make some reasonable initial assumptions—which can eventually be slightly relaxed—concerning the modal force of these epistemic adjectives. These are that *certain* is a maximally strong epistemic operator, as captured in (7-a), and that *possible* is a strictly minimal epistemic operator, as captured in (7-b).

(7) Strong *certain* + weak *possible*:

- a. $\llbracket \textit{certain } \phi \rrbracket^{c,w,e,g} = 1 \text{ iff } Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) = 1$
- b. $\llbracket \textit{possible } \phi \rrbracket^{c,w,e,g} = 1 \text{ iff } Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) > 0$

In both (7-a) and (7-b), the probability of the prejacent is evaluated relative to the available evidence, without conditionalizing on normality assumptions—i.e., we are modeling these terms as ‘pure’ epistemic operators. In addition, since in this framework $E_{e(w)} = \bigcap f_e(w)$, where $\bigcap f_e(w)$ picks out a realistic epistemic modal base such that $w \in \bigcap f_e(w)$, (7) amounts to using our probabilistic framework to model *certain* as an epistemic necessity operator and *possible* as a bare possibility operator. To be sure, (7) is not intended as a full semantic account of *certain* and *possible*. Still, there is suggestive evidence, some presented in what follows, that these terms do have the logical strength captured in (7) (see Lassiter 2016, 2017; Santorio and Romoli 2017; Del Pinal and Waldon 2019).⁸

Given this basic framework and background assumptions, let us now describe the three competing probabilistic accounts of epistemic *must* and *might*.

2.2 Maximal/minimal account

The maximal/minimal account is based on the traditional view that *must* is a maximally strong epistemic necessity operator and *might* is a bare possibility one (von Stechow and Gillies 2010, 2021). To capture this in a probabilistic framework we can simply assume

Footnote 7 continued

by g . To keep entries like (5) simple, we will assume that, for each $w \in W$, $g(w)$ picks out a consistent and non-empty set of normality assumptions.

⁸ As discussed in Sect. 2.2, one consequence of the entries in (7) is that *certain* and *possible* use the same kind of modal base as *must* and *might*. Strictly, we only need to make this assumption for uses of these adjectives in simple unembedded matrix positions like *It is certain/possible that p*. In addition, this assumption is compatible with the view that the epistemic adverbs *certainly* and *possibly* tend to default to more ‘subjective’ modal bases (Lyons 1977; Nuyts 2001) and/or are weaker than their adjectival cousins (see Lassiter 2016).

that *must* is (at least) as strong as *certain* and that *might* is (at least) as weak as *possible* (cf. Rudin 2016). This proposal is captured in (8). Proponents of this view need not hold that there is a perfect correspondence between *must* and *certain*, and/or between *might* and *possible*. For example, the auxiliaries and the adjectives might differ in their presuppositions, or along any other non-truth conditional dimension (cf. Barker 2009). Accordingly, what is captured by (8-c), strictly speaking, are some of the core Strawson-entailments which characterize this account, under the assumption that the target epistemic terms are assigned a uniform modal base.

- (8) Maximal/minimal account of *must* and *might*:
- a. $\llbracket \textit{must } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) = 1$
 - b. $\llbracket \textit{might } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) > 0$
 - c. Strawson-entailments:
 - (i) $\textit{must } \phi \models \phi$
 - (ii) $\textit{must } \phi \models \textit{certain } \phi$
 - (iii) $\textit{possible } \phi \models \textit{might } \phi$
 - (iv) $\textit{must } \phi \models \neg \textit{possible } \neg \phi$
 - (v) $\textit{must } \phi \models \textit{knows } \phi$

According to the account in (8), then, *must* has maximal and *might* has minimal epistemic strength. Since $E_{e(w)}$ is a realistic modal space, *must* is modeled as a veridical operator, and *might* as a bare possibility one. Now, most theorists who defend probabilistic accounts for *must* and *might* don't defend this maximal/minimal account, but there are three reasons why we should examine its prospects. First, this account deals nicely with some of our initial desiderata. For example, since it models *must* as maximally strong, it can explain why we can use *must* in deductive conclusions, as captured in (1), and why it is hard to downplay an assertion of *must* ϕ when ϕ turns out to be false, as captured in (2). Second, relative to the other desiderata in (1)–(4), this account is not as easy to dismiss as is sometimes assumed. In its simplest version, it has trouble explaining the observation, captured in (3), that it is odd to assert *must* ϕ when there is direct evidence for ϕ . Yet von Fintel and Gillies (2010) argue that basic evidential patterns can be explained by assuming, not that *must* is semantically non-maximal, but rather that it presupposes that the evidence for ϕ should be 'indirect'. Third, even if this account fails, determining precisely why it does helps ensure that, when evaluating other accounts, we consider the full range of data that motivated abandoning (8).⁹

2.3 Threshold-based non-maximal/minimal account

As mentioned earlier, given a probabilistic framework there are two natural ways of modeling *must* as a non-maximal and *might* as a non-minimal epistemic operator. The first and arguably more popular approach is simply to lower the probability threshold for *must* and increase the one for *might*. This proposal can be implemented as in (9), which is a simple version of a widely adopted probabilistic account of the auxiliaries (Swanson 2006; Lassiter 2016, 2017). On this view, *must* ϕ entails that the probability

⁹ Amongst philosophers, the view that *must* is a maximally strong and *might* a bare possibility operator is quite popular and often used as a default hypothesis (see e.g., Yalcin 2007, 2010; Dorr and Hawthorne 2013; Stalnaker 2014; Mandelkern 2016, 2019).

of ϕ , $Pr(\phi)$, is above some high but non-maximal contextually determined threshold θ_c^{must} , as in (9-a), where we assume that, for any context c , $\theta_c^{must} < \theta_c^{certain} = 1$. It follows that *must* ϕ is compatible with the possibility that $\neg\phi$, as long as $Pr(\neg\phi) < 1 - \theta_c^{must}$, which in turn entails that *must* ϕ is not veridical—i.e., doesn't entail ϕ . In addition, by comparing the entry for *certain* in (7-a) and the one of *must* in (9-a) we can see that, according to this account, *certain* ϕ asymmetrically entails *must* ϕ . To maintain the standard dualities, we hold that *might* ϕ entails that $Pr(\phi) > 1 - \theta_c^{must}$. As a result, *might*, defined as in (9-b), is stronger than *possible*, defined as in (7-b). The core Strawson entailments of this account are captured in (9-c).

- (9) Threshold-based non-maximal/minimal account of *must* and *might*:
- a. $\llbracket must \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) > \theta_c^{must}$
 - b. $\llbracket might \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e,g} = 1\}) > 1 - \theta_c^{must}$
 - c. Strawson-entailments:
 - (i) *must* $\phi \not\models \phi$ (iv) *possible* $\phi \not\models might \phi$
 - (ii) *must* $\phi \not\models certain \phi$ (v) *might* $\phi \models possible \phi$
 - (iii) *certain* $\phi \models must \phi$ (vi) *must* $\phi \not\models \neg possible \neg\phi$

This threshold-based account can deal with the kinds of cases that are problematic for the maximal/minimal account, such as the epistemic tensions illustrated in (4). The reason for this is that *must* ϕ is strictly compatible with the possibility that $\neg\phi$, while we can still assume that *must* and *might* are duals, with the result that *might* is stronger than a bare possibility operator. The challenge, however, is to square this view with uses of *must* that seem maximally strong, as in deductive conclusions (1) and downplaying scenarios (2) (and parallel cases in which *might* seems to function as a bare possibility operator).¹⁰

2.4 Conditional non-maximal/minimal account

The conditional non-maximal/minimal account predicts the same entailment patterns between *must*, *might*, *certain* and *possible* as the threshold-based account in (9). Yet the implementation is quite different. Specifically, this account follows the traditional maximal/minimal account in (8) in holding that *must* and *might* involve probabilities 1 and > 0 respectively. To achieve the required logical strength and relations, however, I propose that the auxiliaries, unlike their adjectival cousins, use an epistemic probability space that is conditionalized on contextually relevant sets of normality assumptions. From this perspective, *must* and *might* don't use just a pure epistemic space—they use a space conditionalized on (possibly non-veridical) normality assumptions, as captured in (10-a) and (10-b). Recall that $e^g(w)$ is obtained from e and a stereotypical conversational background g as follows: if $e(w) = \langle E, Pr \rangle$, then

¹⁰ Since Kratzer (1981), non-maximal/minimal accounts of *must* and *might* have had a significant following amongst semanticists who use quantificational frameworks (see, e.g., Roberts 2015; Giannakidou and Mari 2016). Amongst philosophers, Dorr and Hawthorne (2013) allow for 'constrained' (non-pure) epistemic readings of the auxiliaries, and Willer (2013) develops a dynamic framework in which *might* can be modeled as a 'live possibility' operator—which is stronger than just 'bare possibility' operator. Using Willer's framework, one can easily formulate a non-maximal dynamic entry for *must*.

$e^{g(w)}(w) := \langle E^{g(w)}, Pr^{g(w)} \rangle$, where $E^{g(w)}$ is defined as $\bigcap g(w) \cap E$ and $Pr^{g(w)}$ is defined as Pr conditionalized on $\bigcap g(w)$.

- (10) Conditional non-maximal/minimal account of *must* and *might*:
- a. $\llbracket \textit{must } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e^{g(w)}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{g(w)},g} = 1\}) = 1$
 - b. $\llbracket \textit{might } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e^{g(w)}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{g(w)},g} = 1\}) > 0$
 - c. Strawson-entailments:
 - (i) *must* $\phi \not\models \phi$
 - (ii) *must* $\phi \not\models \textit{certain } \phi$
 - (iii) *certain* $\phi \models \textit{must } \phi$
 - (iv) *possible* $\phi \not\models \textit{might } \phi$
 - (v) *might* $\phi \models \textit{possible } \phi$
 - (vi) *must* $\phi \not\models \neg \textit{possible } \neg \phi$

The conditional account in (10) generates the same pattern of Strawson-entailments between *certain*, *must*, *might*, and *possible* as the threshold-based account in (9). So how can we empirically separate these non-maximal/minimal accounts?

To begin to see how these two accounts can be differentiated, let us focus on their respective entries for *must*. On the threshold-based account in (9-a), *must* ϕ can be paraphrased as ‘ ϕ is almost certain given the contextually relevant evidence’. On the conditional account in (10-a), *must* ϕ can be paraphrased as ‘ ϕ is certain given the relevant evidence *and* some reasonable assumptions about the world’. This difference can be captured schematically as in (11)–(12). On the threshold-based account, *must* is weakened by lowering its threshold relative to that of *certain*, as captured in (11). In contrast, on the conditional account, *must* is weakened because the probability of its prejacent is determined based on the relevant evidence and a set of normality assumptions that is a superset of the set of normality assumptions, if any, used to determine the probability of the prejacent of *certain*, as captured in (12).

- (11) Threshold-based non-maximal/minimal account:
- a. $0 = \theta^{\textit{possible}} < \theta^{\textit{might}} < \theta^{\textit{must}} < \theta^{\textit{certain}} = 1$
 - b. for all $w \in W$, $g_{\textit{poss/cert}}(w) = g_{\textit{might/must}}(w)$
- (12) Conditional non-maximal/minimal account:
- a. $0 = \theta^{\textit{possible}} = \theta^{\textit{might}} < \theta^{\textit{must}} = \theta^{\textit{certain}} = 1$
 - b. for all $w \in W$, $g_{\textit{poss/cert}}(w) \subseteq g_{\textit{might/must}}(w)$

To see why this theoretical difference leads to empirical differences, let us zoom in on the notion of ‘normality assumptions’, as I understand it here. These are background assumptions that interlocutors take for granted when using evidence to draw inferences within particular domains. They can include ‘trivial’ assumptions such as basic (domain general) principles of reasoning, but also substantive expectations about the world, such as that the basic laws of physics won’t suddenly change and even that if Google says a store is open at t , that store is open at t . In everyday contexts, the sets of background assumptions used to draw inferences from bits of evidence tend to be quite rich, and include assumptions which, as Kratzer emphasizes, are represented as defeasible.¹¹ To capture this notion of normality assumptions, I propose the following minimal doxastic constraint (focusing here on unembedded modalized

¹¹ For example, consider a typical every day inference from a set of specific and salient information, such as when S infers on the basis of looking in Google maps that some restaurant is open. This is a typical

expressions¹²): (i) the normally assumptions that can be used by a speaker S in context c should have the status of mutually held beliefs by the relevant interlocutors in c (i.e., they should be taken for granted in the conversation/deliberation context); and (ii) this does *not* require that S (or the interlocutors) be committed to *knowing* those assumptions—specifically, some normality assumptions can be explicitly represented as defeasible.¹³

At this point, we can derive a key difference between certain doxastic entailments of the threshold-based and of the conditional non-maximal/minimal accounts. Given the conditional account and the doxastic constraint on g , an assertion of *must* ϕ by S will entail that $B_s(\phi)$. To see this, take for concreteness a knowledge norm of assertion (it is easy to check that weaker norms will also work, as long as they require full belief). We can then represent S as committed to $K_s(\text{must } \phi)$. From $K_s(\text{must } \phi)$ we can infer that S believes (i) the relevant evidence, (ii) each of the propositions in the relevant set of normality assumptions, and (iii) that ϕ follows from (has probability 1 given) (i) and (ii). Premise (ii) follows from our doxastic constraint on normality assumptions, and premise (iii) from the semantics of *must* under the conditional account. From (i)–(iii) it follows that by asserting *must* ϕ , S is committed to $B_s(\phi)$.¹⁴

Footnote 11 continued

situation which supports *must*-claims like *McDonald's must be open—I just checked Google Maps*. The salient evidence includes, roughly, ‘ S remembers checking the schedule for the target McDonald’s restaurant on Google Maps’, ‘Google Maps said that McDonald’s is open at the relevant time’, and so on. Yet note that the target inference follows from that specific evidence only given some reasonable (yet strictly defeasible) general background assumptions about the world such as ‘If Google Maps says that r is open at t , then r is open at t ’, ‘If one has an episodic memory as if p at t , then p happened at t ’, and so on.

¹² In some embedded uses of *must*—esp. under propositional attitudes—the doxastic constraint should be anchored to the relevant subject/s. The details will depend on one’s view of the interaction between epistemic modals and propositional attitudes. For discussion, see Sect. 3.3.

¹³ The view that normality assumptions have a doxastic but not an epistemic constraint requires a doxastic and epistemic logic that allows agents to (coherently) reflectively believe propositions that they do not believe they know. Accordingly, the background logics should include (i) the distribution axiom for K_i and B_i , (ii) veridicity only for K_i , so that $K_i(\phi)$ asymmetrically entails $B_i(\phi)$, and (iii) B_i is weak in the sense that $B_i(\phi \wedge \diamond \neg \phi)$ should be consistent, where \diamond stands for simple possibility over a modal base anchored to i . There are various logics for the K_i and B_i -operators that respect (i)–(iii). For example, van Benthem and Smets (2015) model B_i as a universal quantifier over the ‘most plausible’ worlds of epistemic spaces. To cohere with our view, we can assume that if g_p picks out the premises for a plausibility ordering and g for a stereotypical ordering source, then for each $w \in W$, $g(w) \subseteq g_p(w)$, i.e., plausibility orderings can include more information. From this perspective, although an assertion by S of *must* ϕ entails that $B_s(\phi)$, S can consistently acknowledge that it is strictly possible that $\neg \phi$. As we will see in Sect. 3, this ensures that the conditional account can deal with epistemic tensions in which *must* ϕ can be conjoined with the bare possibility that $\neg \phi$. In short, a relatively weak semantics for B_i allows us to hold that speakers/interlocutors have to believe, in that sense, the normality assumptions they use even though some of those background assumptions are explicitly/implicitly represented as defeasible. (For additional evidence that *believe* is weak in roughly this sense, see Hawthorne et al. 2016; Rothschild 2020).

¹⁴ On this implementation of the conditional account—i.e., given our conception of normality assumptions— $B_s(\phi)$ doesn’t entail $K_s(\text{must } \phi)$ or $B_s(\text{must } \phi)$. This is because not just any of S ’s beliefs count as normality assumptions in a given context c —only those that are part of the common ground in c , i.e. that have the status of mutually held beliefs. S may hold various idiosyncratic beliefs that S correctly thinks are not part of the common ground—and can’t be used as background assumptions when reasoning with others about what follows from specific bits of evidence. This prediction is born out in cases like the following. John is anxiously waiting for Peter to arrive at the dinner party. John tends to trust his gut feelings, although he is not so deluded as to think that others would treat his intuitions as reliable indicators.

(13) Doxastic entailments of the conditional account:

a. $K_s(\textit{must } \phi) \models B_s(\phi)$

In contrast, given the threshold-based account, $K_s(\textit{must } \phi)$ doesn't entail that $B_s(\phi)$; it only entails that S believes that $Pr(\phi) > \theta^{\textit{must}}$. This is because S 's holding that $Pr(\phi)$ is greater than an (arbitrarily) high but <1 threshold doesn't guarantee that $B_s(\phi)$.¹⁵ For example, consider a fair lottery case of the sort used to argue against simple 'Lockean' theories of belief (Harman 1986). Suppose John holds one ticket of a fair lottery, knows that he has a very low chance of winning, and although the official results have been drawn, doesn't yet know what they are. John can then assert that he is 'almost certain' he lost, and believe that it is extremely likely that he lost, without also believing that he lost.

(14) Doxastic entailments of the threshold-based account:

a. $K_s(\textit{must } \phi) \not\models B_s(\phi)$
 $K_s(\textit{must } \phi) \models B_s(Pr(\phi) > \theta^{\textit{must}})$
 where $\theta^{\textit{must}}$ can be very high but < 1 .

In Sect. 4, I present various acceptability patterns which challenge the hypothesis that $\theta^{\textit{must}}$ expresses a high but < 1 threshold. I will argue that, in general, *must*-claims just don't seem to pattern with claims that unambiguously express, quantitatively or qualitatively, high but non-maximal probability or certainty. In contrast, the view that *must* ϕ expresses something like conditional certainty, and entails that the speaker believes ϕ , issues in the right predictions for each of these challenging patterns. If this argument is successful, and we wish to maintain the standard dualities, we should in turn favor the conditional over the threshold-based strategy for increasing the strength of *might*.

2.5 Interim summary

In this section, I used a simple probabilistic framework to present three prima facie reasonable hypotheses concerning the strength epistemic *must* and *might*, and spell out, in each case, those of their core entailments and predictions relevant to desiderata (1)–(4). My goal for the rest of this paper is to defend the conditional non-maximal/minimal account. I will do this in two steps. The first step (Sect. 3) focuses on cases that undermine the maximal/minimal account and support—but do not yet discriminate between—the non-maximal/minimal accounts. The second step (Sect. 4) focuses on cases that undermine the threshold-based account and support the conditional non-maximal/minimal account. If this argument succeeds, I will have shown that the conditional account is the only one of the three candidates that can adequately capture

Footnote 14 continued

Someone rings the doorbell, John gets the feeling, and says *I believe it's John, but I wouldn't go so far as to say that it must be it definitely is him*. In scenarios like this, these kinds of epistemic tensions seem quite acceptable.

¹⁵ One could reject this consequence by adopting a simple Lockean theory of belief, such that $\theta^{\textit{believe}} \leq \theta^{\textit{must}}$, but most would agree that this is a costly move.

the uses of *must* and *might* in epistemic tensions, deductive conclusions, epistemic downplaying, and evidential contexts.

3 Maximal/minimal vs. non-maximal/minimal accounts

This section compares the maximal/minimal account and the non-maximal/minimal accounts of *must* and *might*, focusing on their predictions for various kinds of epistemic tensions. I will argue that the observed acceptability patterns support three generalizations: (i) *certain* ϕ asymmetrically entails *must* ϕ , (ii) *must* ϕ is strictly compatible with the bare possibility that $\neg\phi$, and (iii) *might* ϕ asymmetrically entails *possible* ϕ . The key cases control for indirectness, hence are problematic even for maximal/minimal accounts that incorporate an indirectness presupposition. To support (i)–(iii), we will consider intuitive contrasts (Sect. 3.1), experimental data (Sect. 3.2), and embedded epistemic tensions (Sect. 3.3).

3.1 Epistemic tensions: basic cases

According to the maximal/minimal account in (8), *must* is at least as strong an epistemic operator as *certain*. In light of this, consider the examples in (15) (see Lassiter 2017, ch. 6). The oddness of (15-a) suggests that asserting *certain* ϕ commits one to *must* ϕ . This is compatible with (8) and our background assumptions about *certain*. However, the acceptability of (15-b) suggests that we can coherently assert *must* ϕ while being explicitly non-committed to the claim that ϕ is certain:

- (15) *Hearing the distinctive engine sound of Mary's old diesel Volkswagen van...*
- a. ??I'm certain that's Mary's car, though it doesn't have to be.
 - b. That must be Mary's car, though I'm not certain that it is.

Giannakidou and Mari (2016) discuss similar cases where *must* ϕ seems to be compatible with lack of full certainty in ϕ , such as (16-a). The comparative oddness of the minimal pair involving *knows*, in (16-b), *certain*, in (16-c), and the unmodalized bare assertion, in (16-d), suggests—against the knowledge and veridicality predictions in (8-c) of the maximal/minimal account—that *must* ϕ does not entail *knows* ϕ or ϕ .

- (16)
 - a. Ariadne must be sick, but I am not entirely sure.
 - b. #I know Ariadne is sick, but I am not entirely sure.
 - c. #I'm certain that Ariadne is sick, but I am not entirely sure.
 - d. #Ariadne is sick, but I am not entirely sure.

The acceptability patterns in (15) and (16) suggest that *certain* ϕ asymmetrically entails *must* ϕ . To further probe this generalization, note that, as Copley (2004) and Lassiter (2017) point out, constructions of the form *p in fact/indeed q* are acceptable if *q* asymmetrically entails *p*, but are odd if *p* entails *q*:

- (17)
 - a. Tom ate most of my cookies—in fact/indeed, he ate all of them.
 - b. #Tom ate all of my cookies—in fact/indeed, he ate most of them.

Consider then the contrast between (18-a) and (18-b). While expressions of the form *must* ϕ , *in fact/indeed certain* ϕ are acceptable, those of the form *certain* ϕ , *in fact/indeed must* ϕ feel more degraded (cf. Lassiter 2017, ch. 6). The overall pattern suggests, again, that *certain* ϕ asymmetrically entails *must* ϕ :

- (18) *Hearing the distinctive engine sound of Mary's very old diesel Volkswagen van...*
- a. It must be Mary's car that's approaching—in fact/indeed, it's certain that it is.
 - b. ??It's certain that it's Mary's car that's approaching—in fact/indeed, it must be.

The *p in fact/indeed q* construction can also be used to examine whether *might* is at least as weak as *possible*. Given the context in (19), consider the contrast between (19-a) and (19-b). While expressions of the form *possible* ϕ , *in fact/indeed might* ϕ seem acceptable and convey relevant information (that it is a 'live' possibility that Bill attends the party), those of the form *might* ϕ , *in fact/indeed possible* ϕ seem odd or at least comparatively less acceptable. This suggests, contra the maximal/minimal account, that *might* ϕ asymmetrically entails *possible* ϕ .

- (19) *Jamil invites James to his New Year's party. James says that he wants to come but under no circumstances wants to run into Bill.*
- a. Jamil: It's possible that Bill comes; in fact/indeed, he might come.
 - b. Jamil: #Bill might come; in fact/indeed, it's possible that he comes.

Summing up, the patterns in (15)–(18) are in tension with the view that epistemic *must* is a maximally strong epistemic operator, and in particular an operator that is at least as strong as *certain*, and those in (19) are in tension with the view that *might* is as weak as a bare possibility operator. So those acceptability patterns amount to suggestive evidence against the maximal/minimal account of *must* and *might*. In contrast, those patterns are expected given the entailments predicted by the non-maximal/minimal accounts.¹⁶

3.2 Epistemic tensions: experimental evidence

Lassiter (2016, 2017) presents corpus data in which *must* ϕ expressions are conjoined with expressions and contexts that (explicitly) entail that it is strictly possible that $\neg\phi$. Taken at face value, these cases are in tension with S-entailment (8-c)-(iv) of the maximal/minimal account. Representative examples are presented in (20):

¹⁶ The *p in fact/indeed q* test works as intended only if there is independent reason to hold that *p* and *q* are in logical relations with each other (cf. Yalcin 2016, pp. 236–237). Accordingly, I am not using this test to argue that, say, $\langle \textit{must}, \textit{certain} \rangle$ and $\langle \textit{might}, \textit{possible} \rangle$ stand in logical relations with each other. Rather, I am assuming that they do so stand, at least relative to the contexts used above, and use this test to examine hypotheses about what those relations could be.

- (20) a. This is a very early, very correct Mustang that has been in a private collection for a long time . . . the speed[meter] shows 38,000 miles and it must be 138,000, but I don't know for sure.
- b. I have an injected TB42 turbo and like the current setup. There is an extra injector located in the pipping from the throttle body. Must be an old DTS setup but I'm not certain. Why would they have added this extra injector?

Should supporters of strong *must* be worried by this kind of corpus data? Why not reply, with von Fintel and Gillies (2021), that 'people often say the weirdest things'? More to the point, one could argue that examples like (20-a)–(20-b) often involve mid-discourse changes of mind, expansions/contractions of the modal horizon, and/or uses of *must* with pragmatic slack (see von Fintel and Gillies 2010, 2021; Klecha 2014).

To test this kind of response on behalf of the maximal/minimal account, Del Pinal and Waldon (2019) performed a series of experiments which generated acceptability judgments for various epistemic tensions modeled after those in (20). The aim was to compare various kinds of epistemic tensions, given fixed background contexts and minimal pairs. In this way, one can determine the relative strength of different epistemic terms, while controlling for the potential confounds mentioned above. The stimuli were obtained by combining a vignette scheme like (21) with one of the experimental conditions ('epistemic tensions') in (21-a)–(21-e):

- (21) I just bought a vintage bicycle at a garage sale in my neighborhood. It will need some work, but it's in decent shape. EPISTEMIC TENSION HERE. The previous owner didn't know the name of the manufacturer.
- a. 'must $\phi \wedge$ possible $\neg\phi$ ':
The bike must be from the 60s, but it's possible that it isn't.
- b. 'must $\phi \wedge$ might $\neg\phi$ ':
The bike must be from the 60s, but it might not be.
- c. 'certain $\phi \wedge$ possible $\neg\phi$ ':
It's certain that the bike is from the 60s, but it's possible that it isn't.
- d. 'certain $\phi \wedge$ might $\neg\phi$ ':
It's certain that the bike is from the 60s, but it might not be.
- e. 'must $\phi \wedge \neg\phi$ ':
The bike must be from the 60s, but it isn't.

The overall pattern of results was the following. Tensions of the form 'must $\phi \wedge$ possible $\neg\phi$ ' were rated significantly more acceptable than 'must $\phi \wedge$ might $\neg\phi$ '. The latter tensions, in turn, were rated as indistinguishable from plain baseline contradictions of the form 'must $\phi \wedge \neg\phi$ '. This result, summarized in (22-a), undermines prediction (8-c)-(iii) of the maximal/minimal account, suggesting instead that *might* ϕ is stronger than *possible* ϕ . In addition, tensions of the form 'must $\phi \wedge$ possible $\neg\phi$ ' were rated as significantly more acceptable than those of the form 'certain $\phi \wedge$ possible $\neg\phi$ ', which in turn behaved like the baseline contradictions. This result, summarized in (22-b), undermines prediction (8-c)-(ii) of the maximal/minimal account, and suggests instead that *certain* ϕ is stronger than *must* ϕ . Importantly, the same pattern of

results was observed when the conjuncts of each epistemic tension appeared in the order shown in (21-a)–(21-e) (e.g., ‘must $\phi \wedge$ possible $\neg\phi$ ’) and when they appeared in the reverse order (e.g., ‘possible $\neg\phi \wedge$ must ϕ ’).

(22) Results of Del Pinal and Waldon (2019)

(‘ $x > y$ ’ stands for ‘ x was rated as significantly more acceptable than y ’):

- a. must $\phi \wedge$ possible $\neg\phi >$ must $\phi \wedge$ might $\neg\phi$
- b. must $\phi \wedge$ possible $\neg\phi >$ certain $\phi \wedge$ possible $\neg\phi$

Given the experimental setup and controls, we can now ask whether participants/interlocutors find expressions of the form ‘must $\phi \wedge$ possible $\neg\phi$ ’ acceptable due to the involvement of ‘pragmatic factors’, broadly construed, such as that speakers use *must* with pragmatic slack, and/or undergo a mid-discourse change of mind by expanding/contracting their modal horizon (von Fintel and Gillies 2021; Klecha 2014). Suppose, following the maximal/minimal account, that *must* is at least as strong as *certain*, and *might* at least as weak as *possible*. Crucially, the pragmatic factors being proposed are, from the participants’ perspective, a kind of general toolbox for charitable interpretation, used so as to increase the coherence of speakers’ assertions. Accordingly, such factors should increase the degree of acceptability not only of epistemic tensions of the form ‘must $\phi \wedge$ possible $\neg\phi$ ’ but also of those of the form ‘must $\phi \wedge$ might $\neg\phi$ ’ and ‘certain $\phi \wedge$ possible $\neg\phi$ ’. Hence any potential effect of those general pragmatic factors, independent of its magnitude, can’t explain the substantial *differences* in the perceived acceptability of minimal pair tensions with *must* and *certain* (e.g., ‘must $\phi \wedge$ possible $\neg\phi$ ’ vs. ‘certain $\phi \wedge$ possible $\neg\phi$ ’), and minimal pair tensions with *might* and *possible* (e.g., ‘must $\phi \wedge$ might $\neg\phi$ ’ vs. ‘must $\phi \wedge$ possible $\neg\phi$ ’). In contrast, the results come out as expected if we hold, in accordance with the non-maximal/minimal accounts, that *certain* ϕ asymmetrically entails *must* ϕ , *must* ϕ is non-veridical (i.e., is strictly compatible with the bare possibility that $\neg\phi$), and *might* ϕ asymmetrically entails *possible* ϕ .

To try to salvage at least part of the pragmatic suggestion, proponents of the maximal/minimal account could argue that epistemic auxiliaries and adjectives differ with respect to their tolerance for pragmatic slack. Indeed, Lasersohn (1999) discussed examples of truth-conditionally equivalent terms that seem to differ with respect to their tolerance for slack. Now, to try to explain the results in (22), one would need to stipulate, specifically, that *must* is more slack tolerant than *certain*. This would help explain why epistemic tensions of the form ‘must $\phi \wedge$ possible $\neg\phi$ ’ are rated as significantly more acceptable than those of the form ‘certain $\phi \wedge$ possible $\neg\phi$ ’. Yet two observations challenge the adequacy of this stipulation. First, in ordinary discourse, *certain* is often combined with various kinds of slack regulators (e.g., ‘totally/absolutely certain’), suggesting that, in its bare/unmodified form, it allows quite a bit of slack. Second, even when enhanced with that stipulation, the maximal/minimal account can’t predict the full comparative results in (22). For example, we would then expect that tensions of the form ‘must $\phi \wedge$ might $\neg\phi$ ’ should feel more acceptable than tensions of the form ‘certain $\phi \wedge$ possible $\neg\phi$ ’. For the target stipulation entails that the former, but not the latter, could be made coherent by apply-

ing slack. Yet tensions of both forms were rated low in terms of acceptability and, in particular, as unacceptable as plain contradictions (see Del Pinal and Waldon 2019).¹⁷

It is worth emphasizing that I am not denying that, in the course of assigning interpretations to utterances, interlocutors often use slack, pragmatic enrichments, contract or expand the relevant modal horizon, and so on. My claim is rather that appealing to those sorts of pragmatic factors doesn't help the maximal/minimal account explain the comparative acceptability results in (22). In contrast, those results come out as expected given the non-maximal/minimal accounts.

3.3 Embedded epistemic tensions

Yalcin (2007) argues that 'relational' views of epistemics such as Kratzer (1981, 1991) predict that (23-a) should have a reading roughly paraphrasable as (23-b). Yet unlike (23-b), (23-a) feels resiliently odd or incoherent.

- (23) a. #Mary supposes/imagines that it's raining and it might not be.
 b. Mary supposes/imagines that it's raining and she doesn't know that it is.

This contrast has been used to try to motivate more or less radical departures from standard accounts of epistemics and propositional attitudes.¹⁸ Our focus here is on debates specifically about the strength of *must* and *might*, and will appeal to these and similar cases of embedded epistemic tensions to make a different point. I will argue that even if the maximal/minimal and non-maximal/minimal accounts are paired with a plausible account of the interaction between propositional attitudes and modal operators that can explain the contrast in (23), only the non-maximal/minimal accounts make the correct predictions for a range of novel variations of similar embedded epistemic tensions.

As Anand and Hacquard (2013) point out, one way to explain the contrast in (23) is to hold that propositional attitudes affect modal spaces in their scope. Following that lead, I will adopt a semantic account of the oddness of (23-a)—adapting a recent suggestion by Ninan (2018). Consider the entry in (24-a), where $e^{S_x}(w)$ stands for a probabilistic space that captures what x supposes in w , and $e^{S_x^w}(\cdot)$ stands for a 'rigidified' suppositional probabilistic space: for any w' , $e^{S_x^w}(w') = e^{S_x}(w)$. Given (24-a) and any of our candidate entries for *might* (i.e., (8-b), (9-b), or (10-b)), *suppose* (ϕ and *might* $\neg\phi$) comes out as incoherent, i.e., as trivially false. To see this, consider the truth-conditions in (24-b). For the first conjunct to come out true, $Pr_{e^{S_x}(w)}$ has to assign probability 1 to the set of ϕ worlds. Crucially, *suppose* rigidifies the probabilistic space in its scope, which is the one used by the embedded *might*. It follows that, for all worlds w' accessible from w , $Pr_{e^{S_x}(w)} = Pr_{e^{S_x^w}(w')}$. This entails that, in $Pr_{e^{S_x^w}(w')}$, the $\neg\phi$ worlds are assigned probability 0, so the second conjunct comes out false.

¹⁷ Similarly, we would predict that epistemic tensions of the form 'must $\phi \wedge$ might $\neg\phi$ ' should be rated as at least as acceptable as tensions of the form 'must $\phi \wedge$ possible $\neg\phi$ '. For if use of slack over *must* is what rescues the latter from strict incoherence, that same mechanism should also rescue the former from the same fate. Yet tensions of the form 'must $\phi \wedge$ possible $\neg\phi$ ' were rated as substantially more acceptable than those of the form 'must $\phi \wedge$ might $\neg\phi$ '.

¹⁸ For discussion, see Anand and Hacquard (2013), Dorr and Hawthorne (2013), Roberts (2015), Gianakidou and Mari (2016), Ninan (2018).

- (24) a. $\llbracket x \text{ supposes } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e^{S_x}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{S_x},g} = 1\}) = 1$
- b. $\llbracket x \text{ supposes } (\phi \text{ and might } \neg\phi) \rrbracket^{c,w,e,g} = 1$ iff
 - $Pr_{e^{S_x}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{S_x},g} = 1\}) = 1 \wedge$
 - $Pr_{e^{S_x}(w)}(\{w'' : \llbracket \text{might } \neg\phi \rrbracket^{c,w'',e^{S_x},g} = 1\}) = 1$

In short, given the revised entry for *suppose* in (24-a), which shifts the probabilistic space in its scope so as to match what is supposed in the evaluation world, we get a purely semantic explanation for the oddness of (23-b), namely, that it is trivially false.¹⁹ This holds for similar attitudes (e.g., *thinks*) and any reasonable probabilistic account of *might*. That is, our competing accounts all predict that (23-a), given the analysis in (24-b), is trivially false, hence should feel odd or incoherent.²⁰ Accordingly, contrasts like (23-a)–(23-b)—the ones usually discussed in the literature—do not directly inform debates about the *strength* of the epistemic auxiliaries.

Yet the predictions of the maximal/minimal and the non-maximal/minimal accounts can be distinguished for certain variants of the standard cases of embedded epistemic tensions. Specifically, we should examine any potential contrasts in acceptability between expressions of the forms in (25-a)–(25-c):²¹

- (25) a. $x \text{ supposes/thinks } (\phi \wedge \text{might } \neg\phi)$
- b. $x \text{ supposes/thinks } (\text{must } \phi \wedge \text{might } \neg\phi)$
- c. $x \text{ supposes/thinks } (\text{must } \phi \wedge \text{possible } \neg\phi)$

¹⁹ To get a semantic account of the oddness of (23-a), why do we appeal to ‘rigidified’ probabilistic spaces? Consider the entry in (i), where $e^{S_x}(w)$ stands for a probabilistic space that captures what x supposes in w , and in which we don’t further rigidify the probabilistic space used by the modal in the scope of the attitude. Given (ia) and any reasonable entry for *might*, $x \text{ supposes } (\phi \text{ and might } \neg\phi)$ comes out as consistent. This can be seen from the truth-conditions in (ib), which are satisfied in the following situation: x supposes that ϕ in w , so the first conjunct comes out as true, and in addition, x supposes that, in each world compatible with what x supposes in w , x is agnostic about ϕ , so the second conjunct also comes out as true.

- (i) a. $\llbracket x \text{ supposes } \phi \rrbracket^{c,w,e,g} = 1$ iff $Pr_{e^{S_x}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{S_x},g} = 1\}) = 1$
- b. $\llbracket x \text{ supposes } (\phi \text{ and might } \neg\phi) \rrbracket^{c,w,e,g} = 1$ iff
 - $Pr_{e^{S_x}(w)}(\{w' : \llbracket \phi \rrbracket^{c,w',e^{S_x},g} = 1\}) = 1 \wedge$
 - $Pr_{e^{S_x}(w)}(\{w'' : \llbracket \text{might } \neg\phi \rrbracket^{c,w'',e^{S_x},g} = 1\}) = 1$

The reason why, given these assumptions, $x \text{ supposes } (\phi \text{ and might } \neg\phi)$ has a consistent reading is this: although a semantic effect of *suppose* is to shift e to e^{S_x} , still $Pr_{e^{S_x}(w)}$ can determine a different probability measure at the evaluation world w and at any world w' compatible with what is supposed at w . As a result, if we combine *any* of the entries for *might* in Sect. 2 with an entry for *suppose* as in (ia), the oddness of (23-a) would have to be given a non-semantic explanation (for attempts, see Roberts 2015; Dorr and Hawthorne 2013).

²⁰ This account correctly predicts that expressions like $x \text{ supposes } (\phi \text{ and } x \text{ doesn't know } \phi)$ can have coherent readings. This follows from the stipulation that propositional attitudes like *suppose/imagines/knows* shift the modal space over which they are defined (hence they can also do this when embedded under other attitudes). For a related discussion, see Anand and Hacquard (2013).

²¹ As Anand and Hacquard (2013) argue, some propositional attitudes—e.g., *hope* and *doubt*—seem to admit possibility but not (weak) necessity epistemic modals. However, attitudes like *imagine/suppose/think/believe* seem to admit both kinds of epistemic modals. For example, *John supposes that it must be raining* is acceptable (and arguably subtly different in meaning compared to *John supposes that it is raining*).

Assuming the previous ‘shifty’ account of attitudes like *suppose/thinks*, all of our competing accounts predict that instances of (25-a) should come out as incoherent, hence should feel odd. The same applies to instances of (25-b), since they all treat *must* and *might* as duals. Our competing accounts differ, however, in their predictions for instances of (25-c). Accounts which treat *must* as a maximally strong, veridical operator (cf. Yalcin 2007; von Fintel and Gillies 2010; Anand and Hacquard 2013; Ninan 2018), predict that instances of (25-c) are also incoherent, hence should feel roughly as odd as (25-a) and (25-b). In contrast, non-maximal/minimal accounts predict that such attitudes can coherently admit conjunctions of *must* ϕ with the bare possibility that $\neg\phi$ (e.g., when one supposes that one’s evidence and normality assumptions entail ϕ , but that one’s evidence on its own doesn’t strictly entail ϕ). On this view, instances of (25-c), unlike (25-a)–(25-b), can be strictly coherent, hence should feel felicitous or at least significantly less odd.

Non-maximal/minimal accounts of *must* and *might* make the right predictions in these kinds of cases. This is easiest to see if we focus on examples in which both embedded conjuncts can be naturally seen as addressing a question under discussion. Consider the context in (26). While the embedded epistemic tension with *must* $\phi \wedge$ *might* $\neg\phi$, in (26-b), is distinctly odd, the corresponding tension with *must* $\phi \wedge$ *possible/slight chance* $\neg\phi$, in (26-a), feels markedly better. In addition, the embedded tensions with the bare prejacent, $\phi \wedge$ *possible/slight chance* $\neg\phi$, in (26-c), are also distinctly worse than (26-a). The same applies to (26-d), although there is perhaps a slight improvement in this case.

- (26) *The available evidence is in: it strongly suggests that Cain is the murderer. In this country, the judge has to make a verdict and issue a sentence.*
- a. The judge thinks that although Cain must be the murderer, it strictly possible/there is slight chance that he isn’t. So she won’t go for the maximum sentence.
 - b. #The judge thinks that although Cain must be the murderer, he might not be. So she won’t go for the maximum sentence.
 - c. #The judge thinks that although Cain is the murderer, it is strictly possible/there is a slight chance that he isn’t. So she won’t go for the maximum sentence.
 - d. ??The judge thinks that although it is certain that Cain is the murderer, it strictly possible/there is slight chance that he isn’t. So she won’t go for the maximum sentence.

When considering events like court decisions, it is natural to acknowledge that one has to appeal, not just to evidence, but also to (defeasible) normality assumptions about the world, which can introduce uncertainty even in the best cases. These kinds of examples suggest, contra the maximal/minimal yet in accordance with the non-maximal/minimal accounts, that expressions of the form *x supposes/thinks* (*must* ϕ and *possible* $\neg\phi$) are strictly acceptable.

A similar point can be made with a different kind of construction, where we embed just one of the conjuncts, with the goal of trying to increase the informativity of the corresponding epistemic tension. As Anand and Hacquard (2013) argue, *x hopes* ϕ

entails that ϕ is at least a bare possibility for x and that x prefers ϕ to $\neg\phi$. From this perspective, epistemic tensions of the form x *hopes* $\phi \wedge$ *must* $\neg\phi$, in contexts where the epistemic perspective for *must* is anchored to x , should be acceptable only if *must* is non-maximal. This prediction is confirmed by the acceptability of expressions like (27a-i)–(27a-ii) relative to the context in (27). In contrast, the acceptability of (27a-i)–(27a-ii), given the comparative oddness of (27-b)–(27-d), is hard to explain on maximal accounts according to which *must* ϕ (Strawson) entails any of ϕ , *know* ϕ , or *certain* ϕ .²²

(27) *At half time, Liverpool was beating Arsenal 6-0. It's a knockout round, no draws—there has to be a winner. Two disappointed Arsenal fans, Lisa and James, left the stadium at half-time. The game is now likely over. James wonders out loud whether they should walk to a bar and check the final score. Lisa replies:*

- a. There is no point....
 - (i) I hope Arsenal won, but they must have lost.
 - (ii) Although I hope Arsenal won, they must have lost.
- b. There is no point...
 - (i) #I hope Arsenal won, but they lost.
 - (ii) #Although I hope Arsenal won, they lost.
- c. There is no point...
 - (i) #I hope Arsenal won, but I know they lost.
 - (ii) #Although I hope Arsenal won, I know they lost.
- d. There is no point...
 - (i) ??I hope Arsenal won, but I'm certain they lost.
 - (ii) ??Although I hope Arsenal won, I'm certain they lost.

Summing up, we have seen that, when combined with a suitable semantics for attitudes like *suppose* and *think*, probabilistic accounts of *must* and *might* predict the kinds of contrasts observed by Yalcin (2007) and others. What is crucial for us, however, is that, in contrast to the maximal/minimal account, the non-maximal/minimal accounts correctly predict that embedded epistemic tensions which conjoin *must* ϕ with *possible* $\neg\phi$ (or with any other expression, such as *hope* $\neg\phi$, that entails the bare possibility that $\neg\phi$) should come out as strictly acceptable and have coherent readings. We have explored suggestive evidence, based on patterns like (26) and (27), that this prediction is borne out.

4 Threshold-based vs. conditional non-maximal/minimal accounts

So far, we have examined the predictions of the maximal/minimal and the non-maximal/minimal accounts of *must* and *might* relative to the acceptability patterns

²² I argued in Sect. 2.4 that an assertion by S of *must* ϕ typically entails that $B_S(\phi)$. That result does not conflict with the current explanation of the acceptability of (27-a); for recall that we modeled $B_S(\phi)$ as just requiring that ϕ hold in all of the most plausible worlds. Indeed, in a context like (27), *I hope Arsenal won, but I believe they lost* feels quite acceptable. For further discussion, see Sect. 5.2.

generated by various kinds of epistemic tensions. I have argued that those patterns support the non-maximal/minimal accounts: *certain ϕ* asymmetrically entails *must ϕ* , *must ϕ* doesn't entail *knows ϕ* or *ϕ* , *must ϕ* is compatible with the strict possibility that $\neg\phi$, and *might ϕ* is stronger than *possible ϕ* . Those results, however, do not discriminate between the threshold-based and the conditional account. This section focuses on acceptability patterns that can discriminate between those two non-maximal/minimal accounts. The target patterns involve deductive conclusions, contexts of risk, and downplaying scenarios. Although some of these cases have been used to argue against non-maximal/minimal accounts in general, I will argue that, in general, they present a greater challenge to the threshold-based than to the conditional account.

4.1 Deductive conclusions

von Fintel and Gillies (2010) point out that *must*-claims can be used felicitously in conclusions of deductions, as illustrated in (28-a) and (29-a), and argue that this is problematic for non-maximal accounts. Suppose *must ϕ* didn't entail *ϕ* , or that *ϕ* is certain, then shouldn't these uses be pragmatically odd? For using a *must*-claim as a deductive conclusion would imply or suggest a weaker conclusion than that entailed by the common-ground when updated with the premises of the argument. Indeed, deductive conclusions hedged with expressions of high but non-maximal probability, such as (28-b) and (29-b), feel odd and do not seem like correct paraphrases of the corresponding *must*-claims.

- (28) *The ball is in A or in B. It's not in A.*
 a. So it must be in B.
 b. ??So it's almost certain that it is in B.
- (29) *If x is prime and even, then x is 2. x is prime and x is even.*
 a. So x must be 2.
 b. ??So it is at least 98% likely that x is 2.

This objection is most convincing against accounts which stipulate that *must* is cross-contextually non-maximal, such as the threshold-based account in (9-a). For instance, Lassiter (2016, 2017) defends a version of (9) that allows θ^{must} to vary between values that are high—albeit non-maximal—so long as, at each context, the following condition is satisfied: $\theta^{\text{likely}} < \theta^{\text{must}} < \theta^{\text{certain}} \leq 1$.²³ The challenge for this account is that, when combined with standard views on scalar implicatures, it arguably predicts that an assertion of *must ϕ* will tend to generate an upper-bounded implicature—roughly, that the speaker *S* was not in a position to make the stronger assertion that *ϕ* . If triggered in a deductive conclusion, this implicature would clash with the common ground entailment that the assertion of the bare prejacent was war-

²³ Some quantificational accounts also stipulate that, in its epistemic use, *must ϕ* is cross-contextually weak: e.g., Giannakidou and Mari (2016) hold that *must ϕ* presupposes that *ϕ* does *not* hold in all the worlds of the epistemic modal base, and asserts that *ϕ* holds in all of the 'best' worlds of the epistemic modal base.

ranted. As a result, *must* ϕ conclusions would, in such contexts, be incorrectly predicted to feel odd.

The conditional account, however, allows for felicitous uses of *must*-claims in deductive conclusions. I will only sketch my argument here, but I will refine and defend it in Sect. 5.2, after presenting a detailed account of the interaction between epistemic modals and implicatures. According to the conditional account, *must*-claims concern what follows from the evidence, given certain background assumptions. In everyday contexts, such as deciding from specific information in a website whether a store is open, or whether it will be sunny on the weekend, we use relatively rich sets of normality assumptions, which may include assumptions that we represent as defeasible. But when drawing logical inferences from explicit premises, or whenever the goal is to draw inferences that are at least as secure as the premises, we use few, if any, non-trivial background assumptions. In these contexts, speakers can usually be represented as not only believing but also as being certain about their background assumptions. Since the conditional account entails that, for all contexts, $\theta_c^{\text{must}} = 1$, instead of high but < 1 , uses of *must* ϕ in contexts that call on normality assumptions held with certainty will not entail or implicate any degree of epistemic weakness or lack of certainty.

As currently formulated, then, the conditional account is in a better position than the threshold-based account to deal with patterns like (28)–(29). At the same time, we shouldn't overstate the force of this objection. First, it is easy to revise the threshold-based account so as to block the problematic upper-bounded implicatures in deductive contexts. Simply reformulate it so as to allow for the possibility that, in some contexts c , $\theta_c^{\text{must}} = \theta_c^{\text{certain}} = 1$. This can be done by replacing (11) with (30). One can then add that a class of contexts in which the maximal threshold will be typically selected, are precisely contexts where the interlocutors are interested in drawing deductive inferences.

$$(30) \quad \begin{array}{l} \text{a. } 0 = \theta^{\text{possible}} \leq \theta^{\text{might}} < \theta^{\text{must}} \leq \theta^{\text{certain}} = 1 \\ \text{b. for all } w, g_{\text{poss/cert}}(w) = g_{\text{might/must}}(w) \end{array}$$

Secondly, Giannakidou (1999) and Goodhue (2017) argue, based on cross-linguistic evidence, that *must*-conclusions in deductions are not really epistemic. Advocates of the threshold-based account could try to defend that hypothesis. A third option, presented by Lassiter (2016), is open to threshold-based accounts which stipulate that *must* includes a lexicalized evidential signal. In this case, a hearer might (pragmatically) reason from S 's assertion of *must* ϕ —in, say, a deductive context—that S did not assert bare ϕ because S intended to emphasize its evidential status, rather than because S was uncertain about ϕ given the premises. Whether any of these avenues prove promising, we clearly need additional evidence to discriminate between the threshold-based and the conditional accounts.

4.2 Risk in normal worlds

Yalcin (2016) presents contexts involving risk in which expressions with matrix 'weak necessity modals' in their epistemic reading—e.g., with *should* and *ought*—are odd

whereas minimal variants with operators that explicitly convey high but non-maximal certainty/likelihood are acceptable. A representative context of this kind is presented in (31). Expressions with epistemic modals which explicitly convey high but non-maximal certainty/likelihood in the prejacent, such as (31-a)–(31-c), are acceptable in this context. In contrast, close variants with weak necessity modals are odd, as illustrated in (31-d). Extending Yalcin’s pattern, I add the observation that, in contexts like (31), *must* clearly patterns with the weak necessity modals, as can be seen by the oddness of (31-e).

- (31) *Suppose an urn has 100 marbles, 95 white and 5 black. One marble is drawn at random and the speaker doesn’t yet know what color was drawn.*
- a. It is likely that the marble drawn is white.
 - b. It is almost certain that the marble drawn is white.
 - c. It is 95% certain/likely that the marble drawn is white.
 - d. ??The marble drawn should/ought to be white.
 - e. ??The marble drawn must be white.

This kind of pattern is problematic for the threshold-based account in (9). For on this view, *must* ϕ expresses high but non-maximal degree of likelihood or certainty that ϕ . So why are *must*-claims odd in contexts like (31) that involve a known but small risk in the prejacent? Contexts like (31) admit both quantitative and qualitative expressions of a high but non-maximal degree of certainty or likelihood, as shown by the acceptability of (31-a)–(31-c). Thus, that contexts like (31) don’t also admit *must*-claims like (31-e) undermines accounts of epistemic *must* that model it as semantically expressing something akin to ‘almost certain’ or ‘very likely’.

In contrast, the comparative oddness of *must*-claims like (31-e), in contexts like (31), can be easily explained by the conditional account in (10). According to this account, although *must*-claims involve probability 1, they often have non-maximal epistemic status because they include prior conditionalization on normality assumptions. The key observation, concerning contexts like (31), is that there are no (salient/relevant/natural) background assumptions about the world such that, once the available evidence is conditionalized on those assumptions, it follows that the likelihood of drawing a white marble is 1. That is, the default background assumptions in a context like (31) will usually include information like ‘the draws from the urn are fair’, ‘the likelihood of drawing any marble is 1/100’ and so on, and conditional on those normality assumptions, the likelihood of drawing a white marble is 95/100. This is precisely what the context makes explicit. Accordingly, the oddness of (31-e) can be attributed to a clash between that common ground and what would be required to accommodate the *must*-claim (namely, that assuming the world is as expected—i.e., that the lottery is fair and so on—the likelihood of drawing white is 100/100).

To further probe the conditional account, consider the context in (32). (32) is similar to (31) in that the chance that a white marble was picked is very high but there is still some small risk. The difference is that in (32) the draw is executed by a reliable robot designed to pick just the white marbles. The key observation is that, unlike (31), this context admits explicit expressions of high but non-maximal certainty/likelihood and minimal variants with *should* and *must*, as shown by the acceptability of (32-a)–(32-c).

- (32) *Suppose an urn has half white and half black marbles. State of the art, reliable robots have been designed to pick out marbles of specific colors from urns—‘R-Whites’ pick out white marbles and ‘R-Blacks’ pick out black marbles. A robot has drawn a marble but the speaker doesn’t yet know what color was drawn. The speaker then finds out that the marble was picked by an R-White.*
- a. It is very likely/almost certain that the marble drawn was white.
 - b. The marble drawn should be white.
 - c. The marble drawn must be white.

Why can we say, in contexts like (32), that the marble drawn ‘must’ be white, even if we implicitly/explicitly believe that such state-of-the-art robots, although exquisitely crafted, are not strictly functionally perfect, that is, even if, as in (31), there is a risk (that R-White picked a black marble)? The difference is that in (32)—but not in (31)—there is a suitable set of contextually relevant and salient normality assumptions. Obvious candidates—commonly used in every day reasoning—would include background assumptions like ‘such state-of-the-art artefacts perform their intended function’. Given the evidence *and* conditional on those normality assumptions, the probability that the marble was white is 1. As a result, the conditional account correctly predicts that *must*-claims are licensed in contexts like (32). Furthermore, the conditional account, in contrast to the threshold-based one, makes this prediction without also over-generating acceptability for *must*-claims in contexts like (31), which involve some degree of risk even if we only consider the worlds compatible with both the salient evidence and the background normality assumptions.

It is worth reflecting on why the threshold-based and conditional accounts make different predictions about the degree of acceptability of *must*-claims across contexts like (31) and (32). According to the threshold-based account, in contexts where the evidence entails that $\theta^{\text{must}} < Pr(\phi) < 1$ (i.e., that $Pr(\phi)$ is sufficiently high but non-maximal), assertions of *must* ϕ should, in general, feel appropriate—indeed, as appropriate as assertions of *very/n% likely* ϕ and *almost/n% certain* ϕ (for sufficiently high yet non-maximal choices for n). In contrast, according to the conditional account, we should distinguish between two kinds of contexts in which the evidence entails that $Pr(\phi)$ is high but non-maximal. In contexts of the first kind, illustrated by (31), the small but non-zero risk (captured by $Pr(\neg\phi)$) remains after conditionalizing on salient normality assumptions which capture general expectations about the world. In contexts of the second kind, illustrated by (32), the small risk is eliminated after conditionalizing on a suitable set of normality assumptions. The conditional account predicts that while expressions like *very/n% likely* ϕ and *almost/n% certain* ϕ can be used appropriately in both kinds of contexts, *must* ϕ is only strictly appropriate in the latter kinds of contexts, where the small but non-zero risk is eliminated after conditionalizing on a salient set of normality assumptions.

Interestingly, there is a type of objection, commonly presented as against ‘non-maximal’ accounts of *must* in general, which we can at this point show is effective only against the threshold-based account. The objection appeals to odd uses of *must* ϕ in situations that are structurally like (31)—i.e., where given the evidence *and* normality assumptions the probability that ϕ is still high but < 1 . Here is an example presented by von Fintel and Gillies (2021). Take a situation like (33). Given those

facts, an insurance company trying to minimize its costs may formulate a rule as in (33-a) but not as in the odd variant in (33-b):

- (33) *To establish whether a patient has a particular disease \mathcal{D} there are two tests. T_1 is cheap but not always definitive: it can indicate that the patient has \mathcal{D} but often it merely indicates that the patient is more or less likely to have \mathcal{D} . Test T_2 is definitive, but much more expensive than T_1 .*
- a. T_2 can only be administered if the results of T_1 are that it is not certain that the patient p has \mathcal{D} but that p likely has it.
 - b. $\#T_2$ can only be administered if the results of T_1 are that it is not certain that the patient p has \mathcal{D} but that p must have it.

This contrast is puzzling for the threshold-based view. For according to this view, (33-b) would arguably just convey something like (33-a): namely, that T_2 may be used only when T_1 gives the result that it is not certain but it is likely that the patient has \mathcal{D} , except that the relevant threshold for *must* may be higher than the one for *very likely*. In contrast, according to the conditional account, it is easy to see why (33-b) is odd. In the situation relevant to allowing the use of T_2 , the result of T_1 is inherently risky, i.e., is risky even under the assumption that T_1 is functioning normally (structurally, this is like the probability of losing in a fair lottery draw when you hold, say, one ticket out of a hundred). Yet when the rule is formulated as in (33-b), it says that T_2 can be used when the result of T_1 is that it is not certain unconditionally that the patient has \mathcal{D} , but that it is certain conditional on assumptions such as that T_1 is working properly. That is obviously inconsistent with how T_1 is assumed to function—given the common ground in (33)—in situations when it gives high-likelihood but non-certain result.

Summing up, the conditional account, but not the threshold-based one, correctly predicts an interesting contrast in the acceptability of *must*-claims across contexts where there is a small chance or risk that the prejacent is false. The contrast is that, in general, *must*-claims are odd in contexts that are inherently risky or uncertain, yet they are acceptable in contexts in which a comparable level of risk or uncertainty can be eliminated by conditionalizing on normality assumptions.

4.3 Downplaying scenarios

Downplaying scenarios', illustrated in (34), consist of simple dialogues in which (i) a speaker S makes a modal claim whose prejacent turns out to be false, (ii) S is accused of having made a false assertion, and (iii) S replies by insisting that the original modal claim was strictly speaking correct. The key observation, due to von Stechow and Gillies (2010), is that when the prejacent turns out to be false, downplaying a previous assertion of very high but non-maximal probability or certainty, as in (34-c), is an acceptable conversational move. In contrast, downplaying a previous *must*-claim, when its prejacent turns out to be false, is distinctively odd and arguably unacceptable, as illustrated in (35-c).

- (34) a. Jasmine: [*after checking Google*] It's almost/98% certain that it's raining.
 b. Kate: [*opens curtains*] Not it isn't. You were wrong.
 c. Jasmine: Well, strictly speaking, I was not wrong. I was careful. I only said that it's almost/98% certain that it's raining.
- (35) a. Jasmine: [*after checking Google*] It must be raining.
 b. Kate: [*opens curtains*] Not it isn't. You were wrong.
 c. Jasmine: #Well, strictly speaking, I was not wrong. I was careful. I only said that it must be raining.

The contrast between (34-c) and (35-c) is unexpected given views, such as the threshold-based account in (9), which model *must* ϕ as being roughly semantically equivalent in force to expressions of high but non-maximal certainty or likelihood that ϕ . To see why, continue to assume, for concreteness, a knowledge norm of assertion. From this perspective, intuitions about the justifiability of downplaying claims should reflect the following pattern: the stronger the original modalized claim, the less justified it is to subsequently downplay it if the prejacent turns out to be false. Now, according to the threshold-based account, *must* ϕ entails that the likelihood of ϕ is above some high but non-maximal threshold. It follows that, in otherwise matching contexts, downplaying a previous assertion that *must* ϕ should be roughly as acceptable as downplaying a previous assertion of high but non-maximal certainty in ϕ . Yet this prediction is undermined by the contrast between (34-c) and (35-c).

Proponents of the threshold-based non-maximal account of *must*, however, have denied that the alleged contrast in (34)–(35) captures a general pattern. For example, Lassiter (2016) reports that downplaying the high certainty claim in (34-c) feels (roughly) as unacceptable as downplaying the *must*-claim in (35-c). To resolve this disagreement concerning the patterns observed in these cases, Del Pinal and Waldon (2019) designed a series of experiments to obtain acceptability judgments for various downplaying scenarios. The results unambiguously corroborated von Stechow and Gillies's (2010, 2021) original intuition: across a range of different stimuli, conditions that involved downplaying claims of very high but non-maximal certainty, such as (34-c), were rated as significantly more acceptable than matching conditions that involved downplaying *must*-claims, such as (35-c).

Proponents of the threshold-based account of *must* could respond as follows. Asserting an expression of the form '*n*% certain that ϕ ', where *n*% is a high but non-maximal probability, tends to generate the upper-bounded implicature that '*m*% certain that ϕ ', for any *m, n* such that $m > n$. In contrast, although *must* is also non-maximal, it does not have a stronger scale-mate; as a result, asserting that *must* ϕ does not, in general, generate a parallel upper-bounded implicature. Continue to assume that the weaker the original modalized claim, the easier/more justified it is to subsequently downplay it when its prejacent turns out to be false. It follows that if the '*n*% certain that ϕ ' condition (tends to) generate upper bounded implicatures, whereas the '*must* ϕ ' condition doesn't, we can explain why it is easier to downplay in the former case, even if we hold that truth-conditionally both conditions express high but non-maximal probability in ϕ .

Del Pinal and Waldon (2019) designed an experiment to test this response. The target stimuli are like (34)–(35), but with two important manipulations. First, in one condition

the ‘ $n\%$ certain that ϕ ’ sentences were replaced with ‘at least $n\%$ certain that ϕ ’ (where $n\%$ stands for a high but non-maximal probability). The ‘at least’ modification was introduced to block potential upper bounded implicatures (see Krifka 1999; Mayr 2013). Second, in all the conditions the bare prejacent was mentioned immediately before the downplaying sentence. This was done to increase the likelihood that the bare prejacent would be seen as a salient alternative of the modalized claim (see Katzir 2014). The resulting conditions are illustrated in (36).

- (36) a. ‘ $n\%$ certain’ condition:
 ... Well, strictly speaking, I was not wrong... I didn’t say that it is/was raining. I only said that it’s 98% certain that it’s raining.
 b. ‘at least $n\%$ certain’ condition:
 ... Well, strictly speaking, I was not wrong... I didn’t say that it is/was raining. I only said that it’s at least 98% certain that it’s raining.
 c. ‘must’ condition:
 ... Well, strictly speaking, I was not wrong... I didn’t say that it is/was raining. I only said that it must be/have been raining.

The original pattern of results was replicated under these manipulations, as summarized in (37). Downplaying in the ‘ $n\%$ certain’ and in the ‘at least $n\%$ certain’ conditions was rated as roughly equally acceptable. In addition, downplaying in each of those conditions was rated as significantly more acceptable than downplaying in the ‘must’ condition.

- (37) Results of Del Pinal and Waldon (2019) (‘ $x > y$ ’: := ‘ x was rated as a significantly more justified downplaying move than y ’):
 ‘ $n\%$ certain’ condition \approx ‘at least $n\%$ certain’ condition $>$ ‘must’ condition

This pattern of results presents a serious challenge to the threshold-based account according to which *must* ϕ means, roughly, that the probability of ϕ given the evidence is (very) high but non-maximal. Even when we control for potential upper-bounded implicatures, *must*-claims turn out to be significantly harder to downplay than matching claims of very high but non-maximal certainty.

In contrast, the pattern of results in (37) is predicted by the conditional account. According to this account, *must* ϕ says that ϕ follows with maximal probability given the salient evidence *and* a set of relevant normality assumptions. In everyday dialogues like (34) and (35), background normality assumptions include information like ‘if Google says it is m at l , then it is m at l ’. These are (defeasible) assumptions that interlocutors believe, or take for granted, in certain deliberation and conversational contexts, and which they use to draw inferences from specific bits of information such as that Google says that it is raining in Chicago or Atlanta at a particular time. Due to the doxastic constraint on normality assumptions, it follows, as shown in Sect. 2.4, that an assertion of *must* ϕ by speaker S commits S to believing the bare prejacent—i.e., to $B_S(\phi)$. In contrast, asserting *almost/n% certain* ϕ only commits S to believing that ϕ has a high likelihood, given the evidence, but doesn’t entail any full or unhedged doxastic commitment to ϕ . Given this difference in the strength of their doxastic entailments, and the reasonable principle that the stronger a claim, the

harder it is to subsequently downplay it, it follows that assertions of *must* ϕ should be harder to downplay than assertions of *almost/n% certain* ϕ —which is precisely what we observe in the results in (37).

4.4 Where do we stand?

Let us summarize the argument for the conditional non-maximal/minimal account. In Sect. 3, I argued that the predictions of the maximal/minimal account conflict with the acceptability patterns associated with various kinds of epistemic tensions, which suggest that *must* is not a maximal and *might* is not a minimal ('bare possibility') epistemic operator, as assumed by the two non-maximal/minimal accounts. Next, in Sect. 4, I presented cases that aim to discriminate between the threshold-based and the conditional non-maximal/minimal accounts. I argued that the conditional account is in a better position to make sense of the following three observations. First, unlike explicit, unambiguous claims of high but non-maximal probability or certainty, *must*-claims can be used as deductive conclusions. Second, in cases that involve a small risk or uncertainty given the evidence, *must*-claims do not pattern with—and are more selective in certain specific ways than—claims of high but non-maximal probability or certainty. Third, when the preadjacent turns out to be false, downplaying previous *must*-claims is significantly harder than downplaying claims of high but non-maximal probability or certainty. To complete my argument for the conditional account, we need to discuss one more desiderata—evidential uses of the epistemic auxiliaries.

5 Evidential uses

Basic evidential uses of *must* are illustrated by patterns like (38)–(39), where the key observation, captured in (39-a), is that it is odd to assert *It must be raining outside* if the speaker is directly observing that it is raining at the relevant location:

- (38) *Ann and Jasmine are deciding whether to go out. Ann sees people coming in carrying wet umbrellas.*
- a. Ann: It must be raining outside. We should stay in.
 - b. Ann: It is raining outside. We should stay in.
- (39) *Ann and Jasmine are deciding whether to go out. Ann looks out the window and sees that it is pouring rain.*
- a. Ann: #It must be raining outside. We should stay in.
 - b. Ann: It is raining outside. We should stay in.

Despite their popularity in the literature, it is not easy to determine precisely how—or even whether—evidential uses bear on debates about the strength of *must* and *might*. For when suitably supplemented, both maximal/minimal and non-maximal/minimal accounts cohere reasonably well with basic evidential patterns like those in (38)–(39). Here's a sketch of two popular supplementations, the first based on scalar implicatures and the second on the stipulation of an 'indirectness' presupposition:

- If *must* is non-maximal, by asserting *It must be raining* instead of *It is raining*, Ann implies that she doesn't yet know (or is justified in asserting) that it is raining, or that she isn't really certain that it is raining. Those implicatures are compatible with the sort of non-definitive indirect evidence which is part of the common ground in (38), so (38-a) is felicitous. In contrast, since interlocutors usually hold that if x sees that p , then x is in a position to know and be certain that p , those implicatures clash with the common ground in (39), which explains the oddness of (39-a) (cf. Karttunen 1972; Kratzer 1991; Giannakidou and Mari 2016; Goodhue 2017).
- If *must* is maximally strong, we can't (at least straightforwardly) appeal to a scalar implicatures-based account, but evidential patterns can be explained via an indirectness presupposition (von Fintel and Gillies 2010, 2021). An assertion of *must* ϕ at w says that ϕ holds in all the worlds of the epistemic modal base—i.e., that $\bigcap f(w) \subseteq \phi$ —and presupposes that ϕ is neither entailed by nor inconsistent with any proposition in $f(w)$ that is directly known. This 'indirectness' presupposition is satisfied in (38-a) but not in (39-a), which explains why the latter is odd.²⁴

Although both accounts have some initial plausibility, I will argue in what follows that, as currently formulated, they still have various shortcomings, both empirical and theoretical. My main goal, however, is to defend a combination of the conditional account with a specific version of the 'grammatical' approach to the computation of scalar implicatures. I will show that the resulting account can deal with both basic evidential patterns and several closely related variants, without negatively affecting our previous (good) results *vis-à-vis* the other desiderata for theories of epistemics.

5.1 Evidential uses as default implicatures

Initially, it might seem straightforward to get an adequate conditional plus scalar implicatures-based account of basic evidential patterns like (38)–(39) (cf. Goodhue 2017; von Fintel and Gillies 2010, 2021). Assume for now that we are only dealing with contexts that provide substantive, non-trivial sets of default normality assumptions for *must*.²⁵ The target derivation might then go roughly as follows:

- (40) (P1) *Must* ϕ requires that ϕ have probability 1 in the set of worlds compatible with the evidence *and* given some reasonable assumptions about the world.
- (P2) *Must* ϕ competes with a stronger (or non-weaker) alternative $O(\phi)$ that either does not conditionalize on the set of normality assumptions, or does so on a proper subset of the normality assumptions used by *must*.
- (C) The use of *must* ϕ over $O(\phi)$ suggests that the speaker S was not in a position to assert $O(\phi)$, i.e., that S believes that ϕ follows only given the evidence *and* some reasonable but defeasible assumptions.

²⁴ Mandelkern (2019) develops a novel pragmatic account of evidential uses which can be paired with maximal/minimal accounts (and arguably also with the conditional non-maximal/minimal account). I briefly discuss Mandelkern's account in Sect. 5.2.

²⁵ Contexts in which that assumption is not satisfied are discussed in Sect. 5.2.

This sketch raises two questions, however. (i) What alternatives could play the role of $O(\phi)$? (ii) Why compute scalar implicatures in cases when the resulting enrichment creates a conflict with the common ground which would otherwise not occur? Issue (i) is not trivial because, relative to its syntactic category-matching scale mates, *must* is arguably at the top of its scale, even if it is not maximally strong. Yet suppose there is a principled derivation of alternatives that provides candidates for $O(\phi)$. Issue (ii) is still a problem. In standard neo-Gricean frameworks, implicatures are ultimately computed to ‘increase’ the coherence of assertions/speakers, relative to the common ground. It is thus not clear why interlocutors would systematically compute implicatures which result in enriched readings that are inconsistent with the common ground, in cases when the literal, non-enriched readings would not clash with the common ground. Yet this is precisely what would have to occur in (39-a) relative to the common ground in (39).

To address concerns (i) and (ii), I propose that implicature-based accounts of evidential uses should be implemented in a ‘grammatical’ rather than a neo-Gricean approach to scalar implicatures. Grammatical views hold that scalar implicatures are derived compositionally via a covert exhaustification operator, EXH, which for our purposes can be defined as in (41) (Fox 2007; Chierchia et al. 2012).²⁶ EXH(ϕ) asserts ϕ and the negation of all ‘innocently excludable’ (IE) alternatives to ϕ . As defined in (41-b), an alternative ψ of ϕ is ‘innocently excludable’ just in case (i) we can consistently negate ψ while asserting ϕ , and (ii) accepting both ϕ and $\neg\psi$ doesn’t entail any other alternatives of ϕ (not already entailed by ϕ alone).

- (41) a. $\llbracket \text{EXH}(\phi) \rrbracket(w) = \llbracket \phi \rrbracket(w) \wedge \forall \alpha \in \text{IE}(\phi, \text{Alt}(\phi)) : \neg \llbracket \alpha \rrbracket(w)$
- b. $\text{IE}(\phi, \text{Alt}(\phi)) = \{ \psi \in \text{Alt}(\phi) : \llbracket \phi \rrbracket \not\subseteq \llbracket \psi \rrbracket \wedge \neg \exists \psi' [\psi' \in \text{Alt}(\phi) \wedge (\llbracket \phi \rrbracket \wedge \neg \llbracket \psi \rrbracket) \subseteq \llbracket \psi' \rrbracket] \}$

This basic framework allows for different views on the distribution of EXH and the procedure which determines the set of alternatives, $\text{Alt}(\phi)$. On the implementation I propose, expressions are obligatorily parsed with EXH. One motivation for this move—due originally to Magri (2009, 2014) and defended in Del Pinal (2021)—is precisely to explain patterns in which interlocutors seem to systematically compute implicatures which result in clashes with the common ground that, had the enrichment not been computed, would have resulted in informative, felicitous assertions (i.e., to allow for scalar enrichments that decrease the overall rationality/cooperativeness of speakers). Concerning the determination of $\text{Alt}(\phi)$, the choices range from quite formal (Katzir 2007; Fox and Katzir 2011) to highly context sensitive procedures (Swanson 2010, 2017). Yet most views agree that $\text{Alt}(\phi)$ will usually include expressions obtained by replacing any focused scalar terms in ϕ with their scale mates. Furthermore, there is increasing agreement that contextually salient ad hoc scales and alternatives which are

²⁶ Grammatical accounts have various advantages over standard Gricean accounts of scalar implicatures, some of which I discuss below. One that is particularly important for us is that it allows for the triggering of implicatures in (non-asserted) embedded clauses. Interestingly, evidential readings seem to occur in such positions. For example, it is intuitively rather odd to report Ann’s belief state in a scenario like (39) (i.e., when Ann is directly looking at the pouring rain) as *Ann believes that it must be raining*.

not strictly structural alternatives of ϕ can also enter into $Alt(\phi)$ (Katzir 2014; Magri 2017). This is basically what I will assume here.²⁷

What is the result of exhaustifying *must* ϕ ? It is reasonable to hold that, in general, assertions of modalized sentences make salient or are used in contexts in which other modalized sentences are salient. Accordingly, salient (even if not strictly formal) alternatives to *must* ϕ will often include variations of *certain* ϕ , *clear* ϕ , *obviously* ϕ , and so on—i.e., similar modalized sentences, roughly comparable in terms of structural complexity, with target operators that are either epistemically maximally strong or at least non-weaker than *must*. In addition, these alternatives may also systematically include $K^+(\phi)$, where K^+ is a covert pure (non-evidential) epistemic necessity operator.²⁸ Let ' $\mathcal{E}_{s,l}^+$ ' be a placeholder for any such (strong/ish) epistemic operator l , anchored to speaker S . The alternatives to *must* ϕ that are fed to EXH can then be schematically represented as in (42-b). Assuming all the alternatives are relevant, the output of EXH can then be represented as in (42-d).

- (42) It must be raining (= *must*(ϕ))
- a. EXH[*must*(ϕ)]
 - b. $Alt(must(\phi)) = \{might(\phi), must(\phi), \mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)\}$
 - c. $IE(must(\phi), Alt(must(\phi))) = \{\mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)\}$
 - d. $\llbracket EXH[must(\phi)] \rrbracket = must(\phi) \wedge \neg \mathcal{E}_{s,1}^+(\phi) \wedge \dots \wedge \neg \mathcal{E}_{s,n}^+(\phi)$

Given the interpretation in (42-d), an assertion of (42) would convey that S holds that ϕ (= it is raining) is entailed given the evidence *and* some reasonable yet defeasible assumptions about the world but that S doesn't hold that ϕ follows just from the evidence, or even when the modal base is constrained by a more austere subset of assumptions about the world. This captures the intuitive content of (42) in contexts like (38), where interlocutors observe people coming in with wet umbrellas and clothes but not the rain itself. What about the oddness of assertions of (42) in contexts like (39), where interlocutors directly perceive that it is raining? Recall that, on the reading in (42-d), an assertion of (42) will entail that S does *not* hold that it is, say, obviously/definitely raining. Yet when interlocutors (and in particular the speaker) are directly perceiving the pouring rain, it is natural for them to take it as part of the common ground that it is obviously/definitely raining. This results in a clash between the reading in (42-d) and the common ground. Since EXH is obligatory, this clash cannot

²⁷ I should point out, however, that a quite similar account can be obtained even if we adopt a more constrained approach such that $Alt(\phi)$ only includes strictly structural alternatives of ϕ . For discussion, see footnote 29.

²⁸ Independent evidence for the hypothesis that natural languages include a covert epistemic necessity operator is found in recent work arguing that ignorance implicatures should be derived compositionally (Meyer 2013; Buccola and Haida 2019; Marty and Romoli 2021). In addition, as pointed out to me by an editor of L&P, if one adopts a standard Kratzerian semantics for bare indicative conditionals, one also needs to postulate that natural languages include a covert pure (non-evidential) epistemic necessity modal, which can appear as the main modal of bare conditionals. Finally, it is also important to note that Buccola et al. (2021) have recently argued that covert operators can in general be used to form alternatives of expressions with overt operators.

be resolved by selecting a parse without EXH, which explains the resilient oddness of (42) in contexts like (39).²⁹

It is essential to this account that *must*-claims are obligatorily parsed with EXH. For unless the parses with EXH have the status of a resilient default, precisely in contexts that clash with the common ground EXH could be dropped and *must* ϕ assigned an LF that does not generate any (upper-bounded) implicatures (e.g., an LF without EXH). In this case, *must* ϕ could be pragmatically strengthened so as to entail ϕ or that ϕ is certain or perfectly obvious. Accordingly, without obligatory EXH we would be able to explain weakness intuitions in cases like (38-a), which do not result in oddness, but not in cases like (39-a), which do result in oddness. Again, the hypothesis that EXH is mandatory is not an ad hoc stipulation included here simply to derive the observed patterns with evidential uses of modals. That hypothesis has been independently defended precisely on the grounds that it is needed to explain oddness patterns that arguably involve a clash between the common ground and interpretations enriched with implicatures (cf. Magri 2009, 2011; Marty and Romoli 2021; Del Pinal 2021).

This implicature-based derivation of evidential readings of *must* can easily deal with extensions of basic evidential patterns that remain an open challenge to the package of maximal *must* with an indirectness presupposition. Contrast our original example in (39-a), where Ann can't felicitously assert the *must*-claim when looking at the pouring rain, with variations like the ones in (43-a)–(43-b) (adapted from von Stechow and Gillies 2010; Goodhue 2017), where the *must*-claims substantially improve in acceptability:

²⁹ Crucially, a similar result can be derived while assuming a more constrained procedure (e.g., strictly structural one) for determining $Alt(must(\phi))$. Yet in this case the grammatical theory has to then be supplemented with the (increasingly popular) hypothesis that ignorance implicatures are derived compositionally via the interaction between EXH and a (speaker-centric) epistemic necessity operator K_i^+ (see Meyer 2013; Fox 2016; Buccola and Haida 2019; Marty and Romoli 2021). Suppose EXH is obligatory and K_i^+ optional. (42) can then be parsed as in (ia), (iia) or (iiia). The alternatives in each case—i.e., in (ib), (iib), and (iiib)—are either strict scalar alternatives, or obtained through deletion of focused (overt) constituents (cf. Katzir 2007; Fox and Katzir 2011).

- (i)
 - a. EXH[*must*(ϕ)]
 - b. $Alt(must(\phi)) = \{might(\phi), must(\phi), \phi\}$
 - c. $IE(must(\phi), Alt(must(\phi))) = \{\phi\}$
 - d. $\llbracket EXH[must(\phi)] \rrbracket = must(\phi) \wedge \neg\phi$
- (ii)
 - a. $K_s^+[EXH(must(\phi))]$
 - b. $Alt(must(\phi)) = \{might(\phi), must(\phi), \phi\}$
 - c. $IE(must(\phi), Alt(must(\phi))) = \{\phi\}$
 - d. $\llbracket K_s^+[EXH(must(\phi))] \rrbracket = K_s^+[must(\phi) \wedge \neg\phi]$
- (iii)
 - a. EXH[$K_s^+(must(\phi))]$
 - b. $Alt(K_s^+(must(\phi))) = \{K_s^+(might(\phi)), K_s^+(must(\phi)), K_s^+(\phi)\}$
 - c. $IE(K_s^+(must(\phi)), Alt(K_s^+(must(\phi)))) = \{K_s^+(\phi)\}$
 - d. $\llbracket EXH[K_s^+(must(\phi))] \rrbracket = K_s^+(must(\phi)) \wedge \neg K_s^+(\phi)$

Based on each corresponding derivation, it is easy to check that (ia) and (iia) have interpretations that would in general result in incoherent assertions. In contrast, (iiia) supports the coherent reading that the speaker *S* is certain that ϕ follows from evidence and reasonable (defeasible) assumptions, but is not certain that ϕ follows just from the evidence. This approximates the intuitive reading of *must* ϕ in contexts like (38), and arguably still predicts a clash, hence the resulting oddness, in contexts like (39), i.e., when the interlocutors are likely to hold that it is part of the common ground that *S* has the sort of evidence which licenses being certain that ϕ .

- (43) *Ann and Jasmine are discussing—after a seminar on scepticism which they took very seriously—whether they should go out. They look out the window: it looks as if it's pouring rain.*
- a. Ann: It looks as if it's raining outside. There is no good reason to think that both of our perceptual systems are simultaneously unreliable/malfunctioning. So it must be raining outside. Let's get our coffee in the basement.
 - b. Ann: It looks as if it's raining outside. Although we can't {be certain/totally sure/really know} that our perceptual systems aren't misleading us, there is no reason to take that possibility seriously. So it must be raining outside. Let's get our coffee in the basement.

In (43), Ann has direct perceptual information of the sort that usually warrants asserting that ϕ (= it's raining outside). Yet additional factors in this kind of setting suggest that Ann isn't certain, or completely willing to self-ascribe knowledge, that ϕ . Still, ϕ does follow given the (direct) evidence *and* some reasonable (defeasible) assumptions about the world, such as that human perceptual systems are veridical. Accordingly, the conditional plus grammatical account correctly predicts that *must* ϕ can be felicitously asserted, as in (43-a) and (43-b). In contrast, the indirectness presupposition account faces two open challenges. One is to explain how evidence coming from the same source can change in status from direct to indirect as a function of context. The other is to explain why the *must*-claims improve, relative to original cases such as (39-a), even when the speaker explicitly acknowledges some degree of epistemic doubt in the prejacent.³⁰

5.2 Evidential uses, doxastic strength and strong uses of *must*

According to my conditional plus grammatical account, evidential readings of *must* have a degree of epistemic weakness in the sense that, given the target LFs and contexts, *must* ϕ assertions typically entail that ϕ doesn't hold unconditionally given just the relevant evidence. Yet we have also discussed cases—e.g., downplaying scenarios

³⁰ The conditional plus grammatical account coheres well with other results emphasized in recent work on the evidential patterns of epistemics. First, since *Alt* is sensitive to salient alternatives, this account is flexible relative to which epistemic operators are excluded, and allows for stronger enrichments than the one obtained by adding the negation of strict epistemic necessity (e.g., enrichments can incorporate, depending on the context, negation of certainty, clarity or obviousness, just like 'some' claims can be enriched so as to exclude 'all', 'most' or 'half' claims, depending on the context). Secondly, it predicts that cross-linguistic counterparts of *must* should follow the same evidential patterns. Third, it predicts that other strong(ish/er) epistemic modals should generate similar evidential patterns (if they also conditionalize on defeasible normality assumptions, in a way that renders them compatible with the negation of strict epistemic necessity). Fourth, it explains why *can't*-claims generate evidential patterns similar to those observed for *must*—at least if it turns out that, in general, *can't* ϕ is a spell out of \neg *might* ϕ rather than of \neg *possible* ϕ . For example, *S* can felicitously assert *It can't be raining* if *S* sees people coming in with shorts and dry clothes, but the same assertion would be odd if *S* is looking directly at the sunny and clear sky. The default LF for such *can't*-claims is EXH[\neg *might*(ϕ)]. Since *might* is a 'live' possibility operator, the prejacent is asymmetrically entailed by alternatives such as \neg *possible* ϕ , which will thus be negated by EXH when salient and relevant, giving rise to enriched readings along the lines of \neg *might* $\phi \wedge$ *possible* ϕ . The entailment that, given *S*'s evidence, it's possible that it's raining conflicts with what interlocutors will usually take to be in the common ground when *S* is looking directly at a sunny and clear sky.

and deductive conclusions—in which *must*-claims seem to have maximal doxastic or epistemic strength. I will now argue that there is no problematic tension looming here.

To begin to bring out the sense in which *must*-claims are doxastically strong, consider the examples in (44), inspired by Copley (2004, 2006) and Swanson (2016). Fixing for their epistemic readings, an assertion of *must* ϕ generates oddness when it is conjoined with an assertion of $\neg\phi$ or of $B_s(\neg\phi)$, as illustrated in (44-a)–(44-b). In contrast, *must* ϕ improves when it is conjoined with assertions which entail or presuppose the bare possibility that $\neg\phi$, as illustrated in (44-c)–(44-d).

- (44) *John left an hour ago; there’s no traffic; distance is short. QUD: Is John at the party (where the interlocutors are)?*
- a. John₁ must be here by now. #But he₁ isn’t here yet.
 - b. John₁ must be here by now. #But I believe he₁ isn’t here yet.
 - c. John₁ must be here by now. But there is a small chance that he₁ isn’t.
 - d. John₁ must be here by now. But I sure hope he₁ isn’t.

Why are (44-a)–(44-b) odd, whereas (44-c)–(44-d) are fine or at least significantly improved?³¹

Given the conditional plus grammatical account, and the context and conversational goals in (44), the *must*-claims in (44-a)–(44-d) are parsed by default as in (45-a), and assigned the interpretation in (45-b), where we assume that $\mathcal{E}_{s,1}^+$ is a (nearly or strictly) maximally strong and contextually salient epistemic operator:

- (45) John must be here by now (=*must*(ϕ))
- a. LF: EXH[*must*(ϕ)]
 - b. $\llbracket\text{EXH}[\textit{must}(\phi)]\rrbracket = \textit{must}(\phi) \wedge \neg\mathcal{E}_{s,1}^+(\phi)$ $\models B_s(\phi)$

Recall that, when the epistemic space and normality assumptions are anchored to the speaker *S*, the conditional account guarantees that $\textit{must}(\phi) \models B_s(\phi)$, and our background doxastic logic ensures that *S* can coherently believe ϕ while holding that ϕ isn’t, say, certain or perfectly obvious (see Sect. 2.4). This $B_s(\phi)$ doxastic entailment constrains the kinds of assertions that can be conjoined with contents like (45-b). For example, if conjoined with an assertion which entails or strongly suggests that $B_s(\neg\phi)$, we get an incoherent discourse, which explains why (44-a)–(44-b) are odd. At the same

³¹ In their original examples, Copley (2004, 2006) and Swanson (2016) focused on the observation that while *should* ϕ can be conjoined with $\neg\phi$, as in (ia), other (genuine) epistemics, including live and bare possibility ones, generate oddness in parallel structures, as illustrated in (ib)–(ic):

- (i) *John left an hour ago; there’s no traffic; distance is short.*
- a. So John₁ should/ought to be here by now. But he₁ isn’t here yet.
 - b. So John₁ is probably/must be here by now. #But he₁ isn’t here yet.
 - c. So John₁ is possibly/might be here by now. #But he₁ isn’t here yet.

Contrasts like the one between (ia) and (ib)–(ic) suggest that ‘should/ought’ (can) have a ‘pseudo-epistemic’ reading—which doesn’t use an epistemic modal base—akin to the account of ‘normally’ defended by Yalcin (2016). From this perspective, *should/ought* shouldn’t in general be used to set or try to reveal the baseline behavior of strong-ish (but non-maximal) epistemic operators in specific constructions/contexts (as von Stechow and Gillies (2010) sometimes do).

time, (45-b) is strictly compatible with the bare possibility that $\neg\phi$, which explains the improved acceptability of (44-c)–(44-d).³²

I have argued that the conditional plus grammatical account of evidential uses doesn't reduce, in unattested ways, the doxastic strength of *must*-claims. The next task is to show that this account is also compatible with the kinds of cases that seem to require maximal epistemic strength, e.g., felicitous uses of *must*-claims in deductive conclusions.

According to the conditional account, the discourse context provides a set of relevant normality assumptions, picked out by g , which the speaker (and interlocutors) believe, at least for the purposes at hand. In everyday contexts—deciding whether to go outside, whether someone's at the party, and so on—those background assumptions usually include substantive propositions that, although believed, are explicitly represented as defeasible, non-trivial claims about the world. Strictly speaking, it is relative to those kinds of ordinary contexts and corresponding set of normality assumptions that we derive the standard evidential interpretation of *must*-claims from their default parses of the form EXH[*must* ϕ]. For given the definition of EXH in (41), a $\mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)$ alternative is excludable only if its negation is compatible with *must* ϕ , a condition that may be satisfied when *must* ϕ is given a non-maximal interpretation (by being restricted with normality assumptions at least some of which are not held with maximum certainty). These conditions are schematically captured in (46-a)–(46-c):

- (46) *must* ϕ (typical everyday contexts)
- a. LF: EXH[*must* ϕ]
 - b. Given c, w, g includes defeasible, non-trivial assumptions.
 - c. $\llbracket \text{EXH}[\textit{must } \phi] \rrbracket^{c,w,e,g} = \textit{must}(\phi) \wedge \neg \mathcal{E}_{s,1}^+(\phi) \wedge, \dots, \wedge \neg \mathcal{E}_{s,n}^+(\phi) \models B_s(\phi)$

Yet consider a context c' that provides a value for g that only includes trivial assumptions, e.g., some basic tautologies and inference rules. This may occur whenever it is part of the common ground that interlocutors are only interested in the deductive consequences of their premises or information. Such a set of normality assumptions will usually also satisfy the doxastic requirement that they be believed; but it may well include no assumptions that are also represented as defeasible.³³ In a context like c' , then, the speaker S can be represented as not only believing but also as being certain about those background assumptions. What follows from this? According to our grammatical account, the default LFs are still as in (46-a), repeated in (47-a). Let ' K_s^+ ' stand for a maximally strong, unrestricted epistemic necessity operator anchored to

³² Some readers have pointed out to me that (44-d) feels more natural than (44-c). This judgment isn't surprising from the perspective of the conditional plus grammatical account. Strictly speaking, the bare possibility assertion in (44-c) is borderline redundant (I say 'borderline' because it may still clarify to other interlocutors that S is using a set of normality assumptions that includes defeasible propositions). Redundancy can generate oddness, but the conditions under which it does so are intricate and the judgments usually less strong than in cases of incoherence. Still, from this perspective, the *hope*-claim in (44-d) should feel improved because, although it entails/presupposes that, relative to S 's evidence, it's strictly possible that John isn't at the party, it also adds the novel information that S would prefer it if John isn't at the party (yet), which is obviously not conveyed by the initial *must*-claim.

³³ To be sure, interlocutors can represent even trivial assumptions as defeasible in special cases, such as in discussions of metaphysics and philosophical logic.

S . In a context like c' where g returns only propositions of which S is certain, *must* ϕ entails not just $B_s(\phi)$ but also $K_s^+(\phi)$. As a result, most candidates for $\mathcal{E}_{s,l}^+(\phi)$ will not count as excludable alternatives of *must* ϕ , hence EXH will be vacuous, returning only its prejacent, and we are left with a strong reading for (47-a), as captured in (47-c). In these specific conditions, then, we allow for felicitous uses of *must* ϕ in deductive conclusions.³⁴

- (47) *must* ϕ (deductive contexts)
- a. LF: EXH[*must* ϕ]
 - b. Given c' , w , g includes just trivial assumptions.
 - c. $\llbracket \text{EXH}[\textit{must } \phi] \rrbracket^{c',w,e,g} = \textit{must}(\phi) \models B_s(\phi), K_s^+(\phi)$

According to the conditional plus grammatical account, then, the context sensitivity of normality assumptions is such that, although in many everyday contexts *must* is assigned a non-maximal epistemic reading, there are specific conditions in which it's assigned a maximally strong reading. This proposal raises three concerns which I want to briefly address.

First, if g can pick a trivial set of normality assumptions for *must*, why don't interlocutors simply go for that option in cases, like (39-a), when selecting a more substantive set results in non-maximal readings which trigger obligatory implicatures that generate oddness? This strategy is not generally available for the following reason. The conditional plus grammatical account rests on the assumption that discourse contexts must provide a value for g —i.e., a set of relevant normality assumptions—in a way that is, to some degree, independent of the goal of simply trying to make the speakers' utterances felicitous and correct. The goals of interlocutors, conventions and standards appropriate to specific domains, and the question under discussion all contribute to determine specific sets of background normality assumptions. Crucially, similar factors and constraints guide domain restrictions of quantifiers in general. Consider this example:

- (48) *A and B are roommates and just got back from grocery shopping. They bought 4 bottles of coconut water.*
- a. A: Where's the coconut water we got today?
 B: All the coconut water is in the fridge.

³⁴ To be clear, I'm not suggesting that the *only* way to allow for felicitous uses of *must*-claims in deductive conclusions is via selection of 'trivial' normality assumptions (thus blocking the exclusion of any $\mathcal{E}_{i,l}^+(\phi)$ alternative of *must* ϕ). Given the conditional plus grammatical package, other possibilities naturally emerge. In some cases, EXH can associate, in LFs of the form EXH[*must* ϕ], with (constituents of) ϕ , rather than with *must*: e.g., an assertion of *it must be raining*_F can express (i) that the evidence given background assumptions entails that it is raining and (ii) that it is not the case that they entail that it is, say, snowing. In cases like this, there's no obligatory enrichment to instances of $\neg\mathcal{E}_{i,l}^+(\phi)$. In other cases, the discourse may make it clear that the only relevant alternatives are, say, *might* vs. *must*-claims. And since to be considered for exclusion by *exh*, alternatives of the prejacent should also be relevant, in these cases EXH[*must* ϕ] will not implicate weakness. When considering specific variations of evidential patterns and their interaction with deductive uses, it is important to keep in mind these additional mechanisms for generating enriched readings.

Suppose that, in w_1 , two bottles of coconut water are in the fridge and two in the floor by the garage door. In general, we would hold that B's response in (48-a) is false in w_1 , even though there are various domain restrictions of 'all' that would make it true (e.g., 'all the coconut water [*in the kitchen*] is in the fridge'). Suppose that, in w_2 , all four bottles are in the fridge. In general, we would hold that B's response in (48-a) is true in w_2 , even though there are various coconut water bottles within, say, a three mile radius of A and B's home that are not in their fridge. Accordingly, the salient domain restriction, in a context like (48), is roughly 'the coconut water that A and B just bought that is somewhere in their house/car'. Interlocutors can't just freely modify that salient domain to e.g. ensure the correctness of the resulting utterance. In this respect, the normality assumptions picked out by g work just like other types of domain restrictions.³⁵

The second concern is this, take a context like c' , from (47) above, which captures schematically uses of *must* in deductive inferences. Why would a speaker S ever choose to use *must* ϕ , in c' , instead of the apparently less ambiguous ϕ ? In other words, why would S risk being misinterpreted in c' by using a sentence that could be interpreted as non-maximal and even substantially hedged, such as *must* ϕ , instead of simply asserting the bare ϕ (which given a knowledge or similar norm of assertion would generally convey a strong epistemic commitment)? The reason is that *must* ϕ —even when used strongly by conditionalizing on a slim or trivial set of normality assumptions—conveys additional information not typically conveyed by an assertion of ϕ alone: namely, that the reason S holds ϕ is because it follows from this or that specific set of relevant/salient evidence. That is, an assertion of *must* ϕ , more reliably than an assertion of ϕ , highlights information about the specific argument or grounds that S has for holding ϕ .³⁶

The third concern stems from the observation, due to Mandelkern (2019), that there may be additional felicity constraints on *must*-claims even in conclusions of deductions. To illustrate, Mandelkern points to the contrast between odd uses of *must* in conclusions of deductive inferences that are too obvious, such as (49-a), vs. improved uses in conclusions of deductions that are slightly more complex or involved, as in (50-a).

- (49) A: *How many marbles do you have?* B: *I have two bags of marbles. There are (exactly) two marbles in one bag, and (exactly) three marbles in the other...*
- a. B: ??So I must have (exactly) five marbles.
 - b. B: So I have (exactly) five marbles.

³⁵ To be sure, this doesn't exclude the possibility that interlocutors sometimes do tinker with their assignments of normality assumptions precisely to rescue a *must*-claim that would otherwise be odd or too obviously true/false. This might happen when they are unsure about elements of the common ground, incl. the broad goals/standards/topics of the conversation. Imagine a tourist wandering through a hotel lobby which, unbeknownst to them, is holding a philosophy conference, and trying to make sense of utterances like 'we all have a visual representation as if it is pouring rain outside; there is no reason to think we are hallucinating in perfect synchrony; so it must be raining'.

³⁶ For discussion of this 'specific support for the preadjacent' component of *must*-claims, see Stone (1994), Mandelkern (2019), Murray (2020) and Waldon (2021).

- (50) A: *How many marbles do you have?* B: *I have twenty-six bags of marbles, and each bag has at least seven marbles...*
- a. B: So I must have at least one hundred and eighty two marbles.
 - b. B: So I have at least one hundred and eighty two marbles.

Given evidence E , precisely which inferences count as ‘mutually obvious’ depends on the context. Furthermore, the operative notion of obviousness seems to be sensitive to kinds or domains of deductive inferences in a way that is not yet fully understood. For example, why would simple arithmetic operations, such as those in (49-a), count as too (mutually) obvious to support a *must* ϕ -conclusion, whereas inferences based on simple applications of the disjunctive syllogism or modus ponens—standard examples of felicitous uses of *must* ϕ in deductive conclusions, as in (28) and (29)—do not count as obvious to the same degree (indeed, developmentally and cross-culturally, there is evidence that the latter logical inferences tend to be easier in the sense that they are more widely available and independent of formal education)? Still, I think that the basic pattern illustrated in (49)–(50) holds within specific domains, in the sense that the degree of felicity of a *must* ϕ deductive conclusion tends to improve as the complexity of the supporting inference increases.

Using the unique resources of the conditional plus grammatical account, there is a natural way of approaching these patterns which incorporates a key insight from Mandelkern’s own account. Recall why, according to the conditional plus grammatical account, *must* ϕ typically conveys, in a case like (49-a), a maximally strong epistemic claim: the background assumptions are ‘trivial’ principles of arithmetic and inference rules, which are generally held with certainty, and as a result, *must* ϕ conveys that ϕ follows with full certainty from the explicit premises/evidence. So although the *must* ϕ claim is exhaustified, no relevant alternatives—schematically represented as $\mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)$ —are excludable because for none of them can their negation be consistently conjoined with a maximally strong reading of *must* ϕ . However, one could argue that, if we look more carefully into the semantic structure of the alternatives in $\mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)$, the previous prediction should be revised in a subtle but important way. This is because there are alternatives, potentially salient in the relevant cases—think of instances of *obviously* ϕ , *plainly* ϕ , and so on—which are arguably semantically conjunctive in that they convey not just that ϕ is certain, or has maximal probability, given evidence E , but also that ϕ follows from E in a simple or transparent way (cf. Barker 2009). Crucially, maximal uses of *must* ϕ can be consistently conjoined with the negation of such semantically conjunctive epistemic operators.³⁷

³⁷ To see this, assume at least one of the alternatives in $\mathcal{E}_{s,1}^+(\phi), \dots, \mathcal{E}_{s,n}^+(\phi)$ is interpreted as semantically entailing both an epistemic necessity claim and an evidential ‘not too obvious or clear or direct’ condition, which we can schematically represent as $K^+(\phi) \wedge EV(\phi)$. Negating that we get $\neg K(\phi) \vee \neg EV(\phi)$, which can be consistently conjoined with a maximally strong interpretation of *must* ϕ , and the result would be equivalent to $K^+(\phi) \wedge \neg EV(\phi)$. Now, recent work on evidentials suggests that some of the relevant operators—involved in contextually salient alternatives for *must*—may well have a non-trivial at issue vs. non-at issue/presupposed semantic structure, rather than a flat conjunctive semantic structure (see Murray 2020). This might complicate the previous result when the relevant alternatives are negated (since e.g., the evidential part, $EV(\phi)$, may project out of negation if modeled as presupposed). However, even assuming that an operator like, say, *obviously* presupposes rather than asserts either $K^+(\phi)$ or its $EV(\phi)$ entailments, we can still maintain the target result by appealing to a local accommodation operator, which may be

The resulting enriched reading would be, roughly, that ϕ is entailed by the evidence and background assumptions (in this case held with certainty), but not in a way that is completely obvious or transparent.³⁸

At this point, I hope to have shown that the conditional plus grammatical account provides a promising and flexible approach to evidential uses of *must* (and other epistemics), including infelicitous uses, puzzling variations of the basic cases, and oddness patterns which point to a systematic restriction within maximally strong uses.

6 Conclusion

We have examined three prima facie reasonable accounts concerning the strength of epistemic *must* and *might*: the maximal/minimal account in (8), the threshold-based non-maximal/minimal account in (9), and the conditional non-maximal/minimal account in (10). While each account can explain some of the target desiderata, I have argued that only the conditional one can make sense, in a uniform way, of the intricate behavior of *must* and *might* in (embedded) epistemic tensions, deductive contexts, contexts of risk, downplaying dialogues, and various kinds of evidential uses. To be sure, some details need to be worked out before the conditional account can be considered part of a general theory of the epistemic auxiliaries, including its compositional implementation, integration with non-epistemic readings, and connections to related terms and constructions.

Yet even at this preliminary stage, this result issues in a corrective lesson for proponents of probabilistic approaches. Suppose one accepts Kratzer's claim that *must* should not be modeled as simple necessity nor *might* as simple possibility. Given a probabilistic-measure semantics, it is tempting to take advantage of its expressive power and implement that insight by directly tweaking the thresholds for *must* and *might* (e.g., Swanson 2006; Lassiter 2016, 2017). Yet our investigation suggests that we should reject that move and instead implement Kratzer's insight by appealing to an operation that conditionalizes by default on normality assumptions. From this perspective, we should think of *must* and *might* not so much as vehicles for expressing

licensed in fairly standard ways by the pressure to avoid inconsistencies or empty/vacuous applications of EXH.

³⁸ Another option for dealing with patterns like (49)–(50), which is compatible with the conditional plus grammatical account, is to endorse Mandelkern's proposal directly. This proposal is based on the interaction between the semantics of *must* (esp., the component which says that the prejacent follows from the relevant/salient evidence) and some independently motivated pragmatic constraints (esp., a version of the principle that assertions shouldn't be redundant given the information in the common ground). From those premises, Mandelkern derives a felicity constraint which says, roughly, that *must* ϕ assertions are infelicitous if the way in which ϕ follows from the evidence is too obvious to the interlocutors. That explains why (49-a) is odd, while (50-a) is comparatively better. From the perspective of the conditional plus grammatical account, the premises of Mandelkern's account are satisfied at least in contexts that result in strong uses of *must* ϕ . Accordingly, such strong uses would be subject to the fully general pragmatic principles that, according to Mandelkern, further restrict their distribution. I won't try to empirically separate Mandelkern's original account with the grammatical EXH-based implementation I proposed above, but a key difference might be whether we also observe an anti-obviousness constraint in embedded, maximally strong uses of *must*. For such cases are directly expected on the grammatical account, since EXH may appear in embedded positions, but would require some arguably non-trivial modification of the fully pragmatic account so as to get a plausible notion of redundancy relative to local contexts.

what is ‘almost certain’ and ‘practically possible’, but rather as vehicles for expressing what is certain and possible given the relevant evidence *and* a set of contextually appropriate background assumptions about the world.

Theoretically, the conditional account has substantial appeal. Reasoning purely on the basis of what is known, or what we take ourselves to know, has an important place in our discursive and deliberative practices. Yet in most everyday contexts, we reason not just from evidence but also from various non-trivial background assumptions about the way the world normally is or goes, default assumptions which we (tend to) believe even when we don’t represent ourselves as strictly knowing them. From this perspective, we *expect* to find some conventionalized ways of expressing this mode of common sense reasoning from evidence, independently of whether the language system interfaces with, or has access to, a kind of natural probabilistic logic. According to the conditional account, this is precisely the function of *must*, *might* and their cross-linguistic counterparts.

As I said at the outset, the main goal of this paper is to discriminate amongst various accounts of the semantic strength of the epistemic auxiliaries, and not directly to motivate the move to probabilistic frameworks. Yet one might suspect that my case for the conditional account bears on the latter issue. At first glance, only the threshold-based account seems to require a probabilistic implementation. By tinkering with the thresholds as in (11) to capture the relative force of epistemics, it uses the unique resources of a measure semantics. In addition, the view that *must* means something like ‘very likely’ has the consequence that it is not obvious how to satisfactorily translate this account into a standard ordering semantics.³⁹ In contrast, the conditional account can be translated into an ordering semantics without affecting its descriptive adequacy relative to the patterns examined here. So my argument for the conditional account could be taken to suggest that, at least for modeling the epistemic auxiliaries, there is no need for a measure semantics, even less for a genuinely probabilistic one. However, there are other reasons to adopt a probabilistic measure semantics: e.g., the potential for gradability of at least some epistemic auxiliaries (Santorio and Romoli 2017; Lassiter 2017), interactions between nested auxiliaries under other epistemics (Moss 2015; Cariani 2016), subtle differences in the evidential behavior and doxastic implications of epistemics with similar force (Swanson 2016), and theoretical uniformity/simplicity should we conclude that other natural language operators and expressions have access to a probabilistic measure semantics.⁴⁰ Whether these are ultimately good reasons to adopt, for our models of the auxiliaries, not just a measure semantics but a genuine probabilistic one is still an open question.

³⁹ Whether this is ultimately a reason to adopt (i) a measure semantics and (ii) a probabilistic one depends on open debates about the logic needed to model epistemics like *likely* and *probably*. Yalcin (2010) and Lassiter (2015, 2017) develop measure semantic accounts that respect finite additivity and capture various desirable inference patterns not captured by standard ordering accounts. But Holliday and Icard (2013) show that one can capture the target patterns with a weaker measure semantics with qualitative additivity or an ordering semantics with certain lifting functions.

⁴⁰ For example, Santorio and Romoli (2017) implement a probabilistic measure semantics using standard degree semantics, and present an attractive and uniform account of free choice inferences for epistemic adjectives (relying on the scope relations between *exh* and degree operators). If the auxiliaries are modeled in analogous ways, one could extend their account to free choice inferences with epistemic auxiliaries.

Acknowledgements For many helpful discussions, I would like to thank Gabe Dupre, Brian H. Kim, Eleonore Neufeld, Calum McNamara, Matt Mandelkern, Uli Sauerland and Brandon Waldon. Special thanks to Elise Woodard for a crucial suggestion that led to a substantial refinement of my account of evidential patterns, and to Paolo Santorio, Eric Swanson and two anonymous reviewers for L&P for many incisive comments and suggestions on various drafts of this paper.

References

- Anand, P., & Hacquard, V. (2013). Epistemics and attitudes. *Semantics and Pragmatics*, 6(8), 1–59. <https://doi.org/10.3765/sp.6.8>.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind & Language*, 24(3), 253–273. <https://doi.org/10.1111/j.1468-0017.2009.01362.x>.
- Buccola, B., & Haida, A. (2019). Obligatoriness irrelevance and the computation of ignorance inferences. *Journal of Semantics*, 36(4), 583–616. <https://doi.org/10.1093/jos/ffz013>.
- Buccola, B., Kriz, M., & Chemla, E. (2021). Conceptual alternatives: Competition in language and beyond. *Linguistics and Philosophy*. <https://doi.org/10.1007/s10988-021-09327-w>.
- Cariani, F. (2016). Deontic modals and probabilities: One theory to rule them all? In N. Charlow & M. Chrisman (Eds.), *Deontic Modality* (pp. 11–46). Oxford, UK: Oxford University Press.
- Carr, J. (2015). Subjective Ought. *Ergo*, 2, 678–710. <https://doi.org/10.3998/ergo.12405314.0002.027>.
- Chierchia, G., Fox, D., & Spector, B. (2012). Scalar implicature as a grammatical phenomenon. In C. Maienborn, K. von Stechow, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. III, pp. 2297–2331). Berlin: Mouton de Gruyter.
- Copley, B. (2004). So-called epistemic ‘should’. *Snippets*, 9, 7–88.
- Copley, B. (2006). What should ‘should’ mean? Paper presented at *Language Under Certainty Workshop*, Kyoto University. <https://halshs.archives-ouvertes.fr/halshs-00093569>.
- Del Pinal, G. (2021). Oddness, modularity, and exhaustification. *Natural Language Semantics*, 29, 115–158. <https://doi.org/10.1007/s11050-020-09172-w>.
- Del Pinal, G., & Waldon, B. (2019). Modals under epistemic tension. *Natural Language Semantics*, 27(2), 135–188. <https://doi.org/10.1007/s11050-019-09151-w>.
- DeRose, K. (1991). Epistemic possibilities. *The Philosophical Review*, 4(100), 581–605.
- Dorr, C., & Hawthorne, J. (2013). Embedding epistemic modals. *Mind*, 122(488), 867–913. <https://doi.org/10.1093/mind/fzt091>.
- Dowell, J. L. (2011). A flexible contextualist account of epistemic modals. *Philosopher's Imprint*, 11(14). <http://hdl.handle.net/2027/spo.3521354.0011.014>.
- Egan, A., Hawthorne, J., & Weatherston, B. (2004). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 131–170). Oxford, UK: Oxford University Press.
- Fox, D. (2007). Free choice and the theory of scalar implicatures. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 71–120). New York: Palgrave Macmillan.
- Fox, D. (2016). On why ignorance might be part of literal meaning: Comments on Marie-Christine Meyer. Handout for MIT Workshop on Exhaustivity.
- Fox, D., & Katzir, R. (2011). On the characterization of alternatives. *Natural Language Semantics*, 19, 87–107. <https://doi.org/10.1007/s11050-010-9065-3>.
- Giannakidou, A. (1999). Affective dependencies. *Linguistics and Philosophy*, 22(4), 367–421. <https://doi.org/10.1023/A:1005492130684>.
- Giannakidou, A., & Mari, A. (2016). Epistemic future and epistemic must: Nonveridicality, evidence and partial knowledge. In J. Blaszczak, A. Giannakidou, D. Klimek-Jankowska, & K. Migdalski (Eds.), *Mood, aspect and modality: New answers to old questions* (pp. 75–117). Chicago, IL: University of Chicago Press.
- Goodhue, D. (2017). *Must* ϕ is felicitous only if ϕ is not known. *Semantics and Pragmatics*, 10(14:EA). <https://doi.org/10.3765/sp.10.14>.
- Hacquard, V. (2010). On the event relativity of modal auxiliaries. *Natural Language Semantics*, 18(1), 79–114. <https://doi.org/10.1007/s11050-010-9056-4>.
- Harman, G. (1986). *Change in view*. Cambridge, MA: The MIT Press.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404. <https://doi.org/10.1007/s11098-015-0553-7>.

- Holliday, W. H., & Icard, T. F. (2013). Measure semantics and qualitative semantics for epistemic modals. *Semantics and Linguistic Theory*, 23, 514–534.
- Karttunen, L. (1972). Possible and must. In J. P. Kimball (Ed.), *Syntax and Semantics* (pp. 1–20). New York: Academic Press.
- Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy*, 30(6), 669–690. <https://doi.org/10.1007/s10988-008-9029-y>.
- Katzir, R. (2014). On the roles of markedness and contradiction in the use of alternatives. In S. P. Reda (Ed.), *Pragmatics, semantics and the case of scalar implicatures* (pp. 40–71). London, UK: Palgrave Macmillan.
- Klecha, P. (2014). *Bridging the divide: Scalarity and modality*. PhD thesis, University of Chicago.
- Kratzer, A. (1981). The notional category of modality. In H. Eikmeyer & H. Reiser (Eds.), *Words, worlds, and contexts: New approaches in word semantics* (pp. 38–74). Berlin: De Gruyter.
- Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An international handbook of contemporary research* (pp. 639–650). Berlin: De Gruyter.
- Kratzer, A. (2012). *Modals and conditionals*. Oxford: Oxford University Press.
- Krifka, M. (1999). At least some determiners aren't determiners. In K. Turner (Ed.), *The semantics/pragmatics interface from different points of view* (pp. 257–291). Amsterdam: Elsevier Science Publishers.
- Lasersohn, P. (1999). Pragmatic halos. *Language*, 75(3), 522–551. <https://doi.org/10.2307/417059>.
- Lassiter, D. (2015). Epistemic comparison, models of uncertainty, and the disjunction puzzle. *Journal of Semantics*, 32(4), 649–684. <https://doi.org/10.1093/jos/ffu008>.
- Lassiter, D. (2016). *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2), 117–163. <https://doi.org/10.1007/s11050-016-9121-8>.
- Lassiter, D. (2017). *Graded modality: Qualitative and quantitative perspectives*. Oxford, UK: Oxford University Press.
- Lyons, J. (1977). *Semantics*. Cambridge, UK: Cambridge University Press.
- MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In A. Egan & B. Weatherston (Eds.), *Epistemic modality* (pp. 144–178). Oxford, UK: Oxford University Press.
- Magri, G. (2009). A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics*, 17(3), 245–297. <https://doi.org/10.1007/s11050-009-9042-x>.
- Magri, G. (2011). Another argument for embedded scalar implicatures based on oddness in downward entailing environments. *Semantics and Pragmatics*, 4(6), 1–51. <https://doi.org/10.3765/sp.4.6>.
- Magri, G. (2014). Two puzzles raised by oddness in conjunction. *Journal of Semantics*, 33(1), 1–17. <https://doi.org/10.1093/jos/ffu011>.
- Magri, G. (2017). Blindness, short-sightedness, and Hirschbergs contextually ordered alternatives, a reply to Schlenker. In S. Pistoia Reda & F. Domaneschi (Eds.), *Linguistic and psycholinguistic approaches on implicatures and presuppositions* (pp. 9–54). New York: Palgrave Macmillan.
- Mandelkern, M. (2016). A solution to Karttunen's problem. *Proceedings of Sinn und Bedeutung*, 21(2), 827–844. <https://doi.org/10.18148/sub/2018.v21i2.170>.
- Mandelkern, M. (2019). What 'must' adds. *Linguistics and Philosophy*, 42(3), 225–266. <https://doi.org/10.1007/s10988-018-9246-y>.
- Marty, P., & Romoli, J. (2021). Presupposed free choice and the theory of scalar implicatures. *Linguistics and Philosophy*. <https://doi.org/10.1007/s10988-020-09316-5>.
- Mayr, C. (2013). Implicatures of modified numerals. In I. Caponigro & C. Cecchetto (Eds.), *From grammar to meaning: The spontaneous logic of language* (pp. 139–171). Cambridge: Cambridge University Press.
- Meyer, M. C. (2013). *Ignorance and grammar*. PhD thesis, MIT.
- Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5), 1–81. <http://doi.org/10.3765/sp.8.5>.
- Moss, S. (2018). *Probabilistic knowledge*. Oxford, UK: Oxford University Press.
- Murray, S. E. (2020). Evidentiality, modality, and speech acts. *Annual Review of Linguistics*, 7, 213–233.
- Ninan, D. (2018). Relational semantics and domain semantics for epistemic modals. *Journal of Philosophical Logic*, 47(1), 1–16. <https://doi.org/10.1007/s10992-016-9414-x>.
- Nuyts, J. (2001). Subjectivity as an evidential dimension in epistemic modal expressions. *Journal of Pragmatics*, 33(3), 383–400. [https://doi.org/10.1016/S0378-2166\(00\)00009-6](https://doi.org/10.1016/S0378-2166(00)00009-6).
- Portner, P. (2009). *Modality*. Oxford, UK: Oxford University Press.

- Roberts, C. (2015). The character of epistemic modality: Evidentiality, indexicality, and what's at issue. Ms., Ohio State University.
- Rothschild, D. (2020). What it takes to believe. *Philosophical Studies*, 177, 1345–1362. <https://doi.org/10.1007/s11098-019-01256-6>.
- Rudin, D. (2020). Deriving a variable-strength *might*. *Proceedings of Sinn und Bedeutung*, 20, 587–603. <https://ojs.uib.uni-konstanz.de/sub/index.php/sub/article/view/283>.
- Santorio, P., & Romoli, J. (2017). Probability and implicatures: A unified account of the scalar effects of disjunction under modals. *Semantics and Pragmatics*. <https://doi.org/10.3765/sp.10.13>.
- Stalnaker, R. (2014). *Context*. Oxford, UK: Oxford University Press.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4), 487–525. <https://doi.org/10.1007/s10988-008-9023-4>.
- Stone, M. (1984). The reference argument for epistemic must. In H. Bunt, R. Muskens, & G. Rentier (Eds.), *International Workshop on Computational Semantics* (pp. 181–190). Tilburg: Tilburg University.
- Swanson, E. (2006). *Interactions with context*. PhD thesis, MIT.
- Swanson, E. (2010). Structurally-defined alternatives and lexicalizations of XOR. *Linguistics and Philosophy*, 33(1), 221–242. <https://doi.org/10.1007/s10988-010-9074-1>.
- Swanson, E. (2011). How not to theorize about the language of subjective uncertainty. In A. Egan & B. Weatherston (Eds.), *Epistemic modality* (pp. 249–269). Oxford, UK: Oxford University Press.
- Swanson, E. (2016). The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic*, 45(2), 121–146. <https://doi.org/10.1007/s10992-015-9367-5>.
- Swanson, E. (2017). Omissive implicatures. *Philosophical Topics*, 45(3), 117–137. <https://doi.org/10.5840/philtopics201745216>.
- van Benthem, J., & Smets, S. (2015). Dynamic logics of belief change. In H. van Ditmarsch, J. Y. Halpern, W. van der Hoek, & B. Kooi (Eds.), *Handbook of epistemic logic* (pp. 313–385). Milton Keynes, UK: College Publications.
- von Fintel, K., & Gillies, A. S. (2010). *Must... stay... strong!* *Natural Language Semantics*, 18(4), 351–383. <https://doi.org/10.1007/s11050-010-9058-2>.
- von Fintel, K., & Gillies, A. S. (2021). Still going strong. *Natural Language Semantics*, 29, 91–113. <https://doi.org/10.1007/s11050-020-09171-x>.
- Waldon, B. (2021). Epistemic *must* and *might*: Evidence that argumentation is semantically encoded. Paper presented at the 56th Annual Meeting of the Chicago Linguistic Society (CLS), virtual only.
- Willer, M. (2013). Dynamics of epistemic modality. *Philosophical Review*, 122(1), 45–92. <https://doi.org/10.1215/00318108-1728714>.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464), 983–1026. <https://doi.org/10.1093/mind/fzm983>.
- Yalcin, S. (2010). Probabilistic operators. *Philosophy Compass*, 5(11), 916–937. <https://doi.org/10.1111/j.1747-9991.2010.00360.x>.
- Yalcin, S. (2016). Modalities of normality. In N. Charlow & M. Chrisman (Eds.), *Deontic modality* (pp. 230–255). Oxford, UK: Oxford University Press.

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