



Strategic Use of (Un)certainty Expressions

Alexandra Lorson*, Chris Cummins and Hannah Rohde

Linguistics and English Language, University of Edinburgh, Edinburgh, United Kingdom

Speakers have a number of options when introducing propositions which they take to be uncertain: for instance, they can use verbs such as ‘know’, ‘believe’ or ‘think’. The production of uncertainty expressions is highly context dependent. One promising approach to capturing the semantic meaning of these expressions takes them to be available only when the speaker’s confidence in the proposition exceeds some threshold. However, it is unclear whether this approach deals satisfactorily with the full range of usages of uncertainty expressions. For instance, speakers may also use them to achieve social goals such as toning down the force of their assertion. In this case they pursue another communicative goal than just being cooperative: they also aim to be polite. The current study investigates the speakers’ motivations in choosing between uncertainty expressions such as ‘believe’ or the factive ‘know’ in two controlled contexts. More specifically, we show that speakers’ choice of expression is influenced by (i) how likely they estimate an event to be and (ii) strategic considerations relating to the communicative context in which they are working. Thus, speakers adjust their language as a manipulative process. We situate these results in the context of threshold semantics.

OPEN ACCESS

Edited by:

Louis A. De Saussure,
Université de Neuchâtel, Switzerland

Reviewed by:

Stavros Assimakopoulos,
University of Malta, Malta
James Olumide Olufowote,
University of Oklahoma, Norman, OK,
United States

*Correspondence:

Alexandra Lorson
A.Lorson@sms.ed.ac.uk

Specialty section:

This article was submitted to
Language Sciences,
a section of the journal
Frontiers in Communication

Received: 29 November 2020

Accepted: 01 February 2021

Published: 18 March 2021

Citation:

Lorson A, Cummins C and Rohde H
(2021) Strategic Use of (Un)certainty
Expressions.
Front. Commun. 6:635156.
doi: 10.3389/fcomm.2021.635156

Keywords: (un)certainty expressions, modality markers, factive presuppositions, degrees of belief, communicative strategies

1 INTRODUCTION

Speakers have access to a vast repertoire of tools, such as uncertainty expressions (e.g., ‘believe’, ‘think’) and factive verbs (e.g., ‘know’, ‘notice’), to convey their degrees of belief about a specific state of the world, or whether a particular event has taken place. Uncertainty in communication and interaction has been theorised from different perspectives (e.g., Littlejohn et al., 2017), several of which address the idea that the communicator’s goal in interaction is often to reduce or manage their (cognitive) uncertainty (e.g., Berger and Calabrese, 1975; Berger, 1995; Brashers, 2006). Work in linguistic semantics and pragmatics has paid particular attention to the use of expressions that convey information about the (un)certainty of propositional information, which constitute an important tool for reducing a hearer’s uncertainty as to the current state of affairs in the world. Among these expressions are verbs which take sentential complements and which convey different degrees of speaker confidence in the factuality of those complements.

Factive verbs such as ‘know’ are argued to presuppose the truth of their complements, under which assumption we might expect them to be used only by speakers who are certain about the factuality of those complements. By contrast, verbs such as ‘believe’ convey no such presupposition. For example, a speaker uttering (1a) might only have plausible reason for thinking that the glasses are on the kitchen table. In contrast, in (1b) the speaker seems to convey certain knowledge about the

location of the glasses. Focusing on the contrast between (1a) and the bare assertion in (1c), the uncertainty expression in (1a) can be understood to convey the speaker's confidence (or the lack thereof) in the truth of the proposition that the glasses are on the kitchen table.

- (1a). I believe that you left your glasses on the kitchen table.
 (1b). I know that you left your glasses on the kitchen table.
 (1c). You left your glasses on the kitchen table.

Recent work attempting to formalise the semantics of uncertainty expressions (e.g., Yalcin, 2010; Lassiter, 2017) has appealed to the idea of thresholds of probability: For any expression there exists some threshold in the range $[0, 1)$, and an utterance containing that expression is true if the probability of the event it describes exceeds this threshold. Given that cooperative speakers are expected to provide as much relevant information as they can (following Grice, 1975), we would pragmatically expect a speaker to choose the utterance with the highest threshold they can: that is, they should utter (1b) rather than (1a) if they hold that the probability of the event exceeds the threshold that would make (1b) true. Hence, the hearer of (1a) might infer that the speaker is not certain enough to assert (1b).

However, this inference relies on the assumption that there is no reason for a speaker to have used a weaker alternative: that is, to have said 'believe' when they were in a position to say 'know'. In practice, we might intuit that a speaker will sometimes opt for the weaker option despite their knowledge state. One possible motivation for this would be politeness. For example, suppose that a valuer has inspected a painting which the owner thinks is worth a fortune, and the valuer is certain that it is not. Out of politeness, they might still utter (2a) rather than (2b).

- (2a). I believe that your painting isn't worth much.
 (2b). I know that your painting isn't worth much.

In this paper, we examine language users' expectations about the choice between 'know' and uncertainty expressions such as 'believe'. In particular, we investigate the effect of the strength of evidence available to the speaker, across two different scenarios in which the expectations of appropriate speaker behavior differ. As such we examine ways in which participants estimate speakers' strategic use of (un)certainty expressions, or to phrase it differently, ways in which they posit that speakers may be using language as a manipulative tool. Our broad aim is to explore the usefulness of a threshold semantics approach in capturing language usage in this domain. In experiment 1, we look at a variety of verbs conveying different degrees of belief: 'know', 'notice', 'to be sure', 'think', 'believe', 'see', 'guess'. Henceforth, we will use the term *(un)certainty expressions* when referring both to uncertainty expressions and to factive verbs. In experiment 2 we will zoom in onto the (un)certainty expressions 'believe' and 'know'.

2 PREVIOUS WORK ON UNCERTAINTY EXPRESSIONS

2.1 Threshold Semantics

Uncertainty expressions indicate the degree to which speakers are committed to the content of their utterances and reveal the speakers' knowledge about the truth of the presented proposition. Interlocutors rely on contextual factors to produce and interpret uncertainty expressions, since there is no straightforward translation between uncertainty expressions and event probabilities (Clark, 1990). For example, hearers have been found to have difficulties in understanding verbal probability descriptors such as 'common' when being informed about the risk of medical side effects Knapp et al. (2004).

One way of formalising uncertainty expressions, semantically, would be to assume a threshold semantics (Yalcin, 2010; Lassiter, 2017): For any expression there exists some threshold $[0, 1)$, and an utterance containing that expression is true if the probability of the event it describes exceeds this threshold. For example, if the threshold for 'believe' is 0.6, then the utterance 'I believe that Scotland will win the match tomorrow' is available for a speaker who believes that the probability of Scotland winning the match tomorrow exceeds the threshold of 0.6. Thus, a threshold semantics account acknowledges the influence of context on the interpretation of uncertainty expressions with regards to varying event probabilities.

However, research suggests that the usage of (un)certainty expressions is more dynamic than predicted by threshold accounts. For example, there is a considerable amount of inter-speaker variability (e.g., Wallsten et al., 1986) that might not be explainable with varying event probabilities. In a recent study, Schuster and Degen (2020) suggest that hearers use the speaker's identity when interpreting utterances: what expressions does a speaker preferably use and what thresholds do they associate with specific expressions? Schuster and Degen (2020)'s study tests how hearers adapt to speaker-specific thresholds when interpreting uncertainty expressions. In one part of their study, participants were asked to listen to either a cautious or confident speaker producing a bare assertion or utterances containing 'probably' or 'might'. The confident speaker would use 'might' and 'probably' to describe lower event probabilities than a cautious speaker. Participants were then asked to make guesses about the production of the speaker they were introduced to in the first part choosing between 'might', 'probably' and a *something else* option. The results show that participants in the confident speaker condition gave high ratings for 'probably' for a larger range of event probabilities than for 'probably' in the cautious speaker condition. For 'might' the opposite was observed: participants gave high ratings for 'might' for a larger range of event probabilities in the cautious speaker condition than in the confident speaker condition. In a third experiment, participants inferred higher event probabilities for 'might' and 'probably' when produced by the cautious speaker than the confident speaker. Thus, hearers seem to be able to adjust to speaker-specific thresholds, e.g., use higher thresholds for

expressions such as ‘probably’ when interpreting utterances communicated by a cautious speaker.

Schuster and Degen (2020)’s findings suggest that hearers adapt to the speakers’ identity. Crucially, this line of research focused on the utility of an utterance as being determined primarily by its informativeness and the speakers’ utterance preference. Thus, speakers aim to reduce uncertainty for the hearer by choosing the most informative utterance while also having personal preferences as to which utterance to choose. As was pointed out by Schuster and Degen (2020), while these considerations successfully capture inter-speaker variability, it is not clear how such an approach would capture the effect of additional communicative goals speakers may have when using uncertainty expressions. For example, a speaker may use ‘might’ to make a statement more moderate, or to be polite—goals which may vary by context even for the same speaker and which conflict with the goal of presenting the strongest possible information, in order to reduce uncertainty for the hearer. In the following section, we discuss these goals in more detail.

2.2 Communicative Strategies

2.2.1 Politeness

It has long been noted that uncertainty expressions that convey modal meaning also give rise to particular discourse effects (e.g., Fraser, 1975; Holmes, 1982). Uncertainty expressions may function as *downtoners* (Holmes, 1982) or *hedges* weakening the assertive strength of an utterance, and yielding discourse effects such as vagueness or politeness. The concept of *hedges* was popularised by Lakoff (1973) to capture linguistic expressions, such as ‘sort of’, that can signal different degrees of category membership of a particular expression, see (3).

(3) This paper is sort of long.

This type of hedging was later called *propositional hedging* (Fraser, 1975), since the hedge affects the truth value of the propositional content. Hedging was further investigated by Brown and Levinson (1978) who extended hedging to expressions that modify the speaker’s commitment to a proposition, such as ‘think’ (sometimes discussed as *speech act hedges* Fraser, 2010). This type of hedging is the focus of this paper. For example in (4), a teacher might want to gently introduce the subject about a pupil having to repeat the 4th grade to the parents (4a) than to directly assert it (4b).

(4a). I believe that your son will have to repeat the 4th grade.

(4b). Your son will have to repeat the 4th grade.

Within politeness theory (Brown and Levinson, 1987), being polite is often analysed through the lens of facework: that is, the interlocutors’ aim to maintain their positive or negative face. Whereas positive face reflects the interlocutors’ maintenance of a positive self-image, negative face reflects the interlocutors’ freedom to act on their own terms. Face-threatening actions can damage the face of either the speaker or hearer. For example, in a scenario where two teachers are discussing a student’s mark,

see (5), Teacher2 is threatening their colleagues’ positive face by criticizing them. Similarly, this situation could be potentially face threatening for Teacher2 because they could come across as uncompassionate. A strategy to lessen the severity of the threat would be the use of an uncertainty expression such as ‘believe’ (5a)—i.e., hedging—instead of asserting the criticism directly (5b).

(5). Teacher1: This essay was submitted on time.

(5a). Teacher2: I believe you are mistaken; the student submitted the essay a day late.

(5b). Teacher2: You are mistaken; the student submitted the essay a day late.

Empirical research suggests that speakers indeed consider motives such as politeness when communicating their degrees of belief. Juanchich and Sirota (2015) investigated the way in which speakers communicate more or less face-threatening news to their friend. More specifically, participants were confronted with a scenario that had less severe consequences (friend’s car breaking down), and another that had more severe consequences (friend making a bad investment). Both events were characterised as being 50% likely to occur. Juanchich and Sirota (2015)’s findings suggest that tactful speakers who are concerned about the hearer hedge their utterances, understating their confidence by using expressions such as ‘a very small probability’, ‘a small probability’, ‘slightly probable’.

Holtgraves and Perdeu (2016) extended this line of research by looking at scenarios with varying event probabilities (20, 50, and 80%) from the production and comprehension perspective. The findings of a production task indicate that speakers hedge their utterances by using expressions conveying lower degrees of belief when hearers are more severely affected by the event the speakers are describing (e.g., ‘It’s somewhat unlikely/likely that the car needs a new transmission.’). This was the case even when both events—severe and less severe—were equally likely.

From the hearers’ perspective, Holtgraves and Perdeu (2016) found that participants assigned high probabilities (value ranging from 0 to 100) to a given event when hearing expressions conveying high certainties, such as ‘definitely’ (e.g., ‘The car definitely needs a new battery/transmission.’). At the same time, participants assigned lower probabilities to severe events (car needs a new transmission) than for less severe events (car needs a new battery). This is a striking result in the context of prior research by Bonnefon and Villejoubert (2006) showing that participants tend to judge severe negative outcomes to be more likely than the speaker conveys. Taken together, these studies suggest that speakers downplay the probability of severe adverse events but hearers are aware of this and compensate accordingly in their interpretation. In general, speakers who use hedging strategies are perceived as less authoritative and confident than those who do not (Hosman, 1989; Crismore and Kopple, 1997). At the same time, speakers are also perceived as warmer, as we might predict if hearers take into account politeness theoretic constraints.

2.2.2 Power Dynamics

The severity of a face threatening action depends partly on the power relation between the interlocutors (Brown and Levinson, 1987). Low-power speakers may be more concerned about facework than speakers who have the same or a higher social status than the hearer. For example, if a teacher were talking to their student, as in (6), rather than a colleague, as in (5), they might be less inclined to resort to hedging (6b) and more inclined to contradict the speaker directly (6a).

- (6). Student: I submitted my essay on time.
 (6a). Teacher: You are mistaken; you submitted the essay a day late.
 (6b). Teacher: I believe you are mistaken; you submitted the essay a day late.

The previously discussed studies show that speakers use hedging when they are on a par with the hearer (communicating with friends, Juanchich and Sirota (2015)) or in a lower social position (communicating with parents, Holtgraves and Perdeu, 2016). It is less clear to what extent high-power speakers engage in facework. From the hearers' perspective, high-power hearers (role of parent) may dismiss the possibility that the low-power speaker (child) could use hedging (Holtgraves and Perdeu, 2016). On the other hand, as shown by Bonnefon and Villejoubert (2006) in the context of medical communication, low-power hearers (patients) appear to be aware that a high-power speaker (doctor) may hedge their statements.

However, the effects of power relations might arise here for two separate reasons: because the bare assertion is taken to be face-threatening to the hearer in a way that is inappropriate in the presence of a particular power dynamic, or because the risk to the speaker of making a false statement is higher in such a case. That is, it's not obvious whether the speaker who hedges does so because they could confidently make a stronger statement but that would be inappropriate, or because they require higher confidence to make a stronger statement in the presence of an unfavourable power dynamic. In a threshold semantics analysis, this latter case could be treated as a case in which the threshold for making the stronger statement has increased. For example, immediately after the vote count for the 2020 US presidential election, Republican senators who refused to say that Trump had lost may have done so because they feared reprisals for saying so even though they thought it was certainly true, or because they didn't want to take the risk of asserting his defeat falsely in a circumstance where they thought he might still have a very slim chance of winning.

Previous studies have singled out specific power dynamics. One way of interpreting these findings is as evidence of speakers' selection of communicative strategies. Criticising someone who is in a high-power position requires a different strategy than criticising someone who is in a low-power position. A successful strategy for a low-power speaker could be to hedge/downplay their certainty about their interlocutor's mistake. However, there are scenarios where a speaker might choose a quite different strategy: conveying high certainties. For example, in a scenario where the hearer is thought to have lied about what they did last night and the speaker has at least a suspicion about

what the hearer was up to, the speaker might want to act as if they are already certain about what the hearer did in order to elicit a strong reaction and ultimately learn the truth. Ways of conveying high certainty would include using the bare assertion, see (7b), or to use 'know' (7a).

- (7a). I know that you went to the party without me.
 (7b). You went to the party without me.

Note that the use of the factive presupposition trigger 'know', although indirect, might confer an advantage over using the bare assertion: by using 'know' the speaker not only conveys that they are highly certain that the hearer went to the party without them but also acts as if this content is a fact and something that everyone including the hearer can agree on. By definition, presuppositions are presumably already shared knowledge. If presupposed content happens to be new, one way of repair is for the interlocutors to accommodate the presupposition: that is, to act as if it was in the common ground. Presupposed content triggered by 'know'—the interlocutor being at the party without the speaker—might be less likely to be challenged by the hearer, even if it is in fact controversial. Lorson et al. (2019) contrasted the usage of assertions vs. presuppositions and their findings suggested that (formally) presupposed content was less likely to be challenged by the hearer than asserted content.

In summary, speakers use uncertainty expressions to communicate their degrees of belief. The willingness to use particular expressions seems to vary between speakers, and interlocutors are flexible in adapting to each other's production preferences. But, at the same time, speakers may also use a certain expression when the event they are describing does not exceed the threshold of that expression. The thresholds of uncertainty expressions seem to vary between interlocutors and interlocutors seem to be flexible when it comes to adjusting to each other's way of using uncertainty expressions. On top of that, speakers may use uncertainty expressions as hedging devices or downtoners following communicative goals such as being polite. Thus, it seems that there are a multitude of factors contributing in different ways to the speakers' production of uncertainty expressions. In this paper, we explore whether we can appeal to a threshold-based semantic analysis to systematise the effects of these competing considerations.

3 THE PRESENT STUDY

The current study investigates the speakers' motivations in choosing between uncertainty expressions such as 'believe' or factive verbs such as 'know'. More specifically, we explore whether participants' choices of expressions are influenced by (i) how likely they estimate an event to be and (ii) strategic considerations relating to the communicative context in which they are working. We will extend previous research by introducing a different way of assessing participants' degrees of belief, and by introducing a within-subjects manipulation that examines the effect of context on a speaker's strategic utterance choice.

The majority of studies that we discussed have used quantitative prompts to manipulate the probability that an event takes place. Using probabilities is problematic for three reasons: (i) In daily life speakers don't usually know the exact probability with which an event takes place; (ii) speakers rely on evidence/arguments/experience/intuition rather than reasoning about event probabilities when communicating degrees of belief; (iii) interlocutors usually perform quite poorly when it comes to understanding probabilities (Kahneman, 2011). In our study, we ask participants to choose between utterances to produce, while showing them pictures and documents that can be implicitly evaluated for the event certainty they denote. We assess the participants' degree of belief after the production task by asking participants to rate how likely they thought that a specific event took place given a piece of evidence. By taking this approach, we aim to elicit production choices and event certainty judgements that are more similar to those occurring in daily communication.

Prior work has shown that participants can adjust to speaker-specific utterance thresholds. Here we ask whether speakers' use of (un)certainly expressions also varies depending on the scenario they are in and who they are talking to. We target two scenarios that vary in their power dynamics to compare production choices made by a speaker in a high-power position to those made by a speaker on a par with their interlocutor. The advantage of a within-subjects manipulation (in contrast to prior work that has tested only one power dynamic or another) is that we can examine the extent of context-specific adaptation while holding the speaker constant. In our study, we test whether participants adapt to a change of context to change the way they convey their (un)certainly.

We conducted two experiments—the first experiment examines seven (un)certainly expressions and the second experiment focuses on 'know' vs. 'believe'. Both involved a production task where participants were asked to choose between utterances to convey messages about different events, followed by an evaluation task where participants were asked to adjust a slider to indicate their evaluation of the certainty of those events. Note that, while we characterise this primarily as a production task, participants' choices in this task can also be understood to involve elements of comprehension: the participant is asked to indicate what a character in a scene is likely to say, by evaluating and choosing among different candidate utterances. This decision-making process reflects language users' awareness that what they say is guided by how it may be understood by the hearer, meaning that speakers and hearers engage in what can be termed 'mutual vigilance' (Sperber et al., 2010). The only difference between these experiments in terms of task is the number of utterances participants had to choose from. The full set of data for both experiments is available here: <https://osf.io/e5av9/>.

3.1 Experiment 1: Seven (Un)certainly Expressions

In the first experiment we investigated the production of a wide range of (un)certainly expressions ('know', 'notice', 'to be sure', 'think', 'believe', 'see', 'guess') in two controlled, contrasting scenarios. The goal of this experiment was to (i) test the

experimental design and the materials and (ii) narrow down which (un)certainly expressions we should contrast in the second experiment.

The experiment consisted of two tasks. For the production task, participants were asked to play the role of a detective in an investigation of an art heist where they briefed a colleague and interrogated a suspect, relying on pieces of evidence about the suspect's whereabouts. In this way we elicited expressions that correspond to degrees of belief without providing explicitly quantitative prompts. After the production task, participants then evaluated their confidence in each piece of evidence retrospectively. We call this evaluation our evidentiality measure. Thus, participants produce utterances in two scenarios corresponding to a range of confidence levels in the propositions uttered.

We expect participants to be pragmatically cooperative in the sense of producing the most informative statement that they truthfully can. If the threshold-based account of the semantics of these verbs is correct, we would expect speakers to select the verb with the highest threshold that does not exceed their degree of belief in the complement proposition. This means that, assuming 'know' has a higher threshold than 'believe', speakers are predicted to choose 'believe' over 'know' as long as the probability of the event they are describing does not exceed the threshold for 'know'. Following Schuster and Degen (2020)'s results, we expect speakers to differ in their thresholds for producing specific expressions.

Considering work in the domain of politeness theory, we generally predict that verbs expressing higher degrees of belief will be more widely used in the interrogation than in the briefing. We examine whether this difference can be captured by assuming that the scenario exerts a general effect on the thresholds for producing particular utterances. This is based on the assumption that speakers might have a systematic tendency to downplay their certainty when speaking to their colleague in a cooperative scenario, compared to when they are interacting with a suspect in an uncooperative scenario. We could also reformulate this in terms of power relations: In the interrogation the speaker is in a high-power position compared to the briefing, which potentially obviates their need to consider politeness constraints.

3.1.1 Participants

We tested 35 participants recruited over the crowd-sourcing platform Prolific, specifying participants with an approval rate above 90. Participants were paid an average of £7.53/h (the average duration of the experiment was 35 min). The age of the participants ranged from 18 to 52 years, with a mean of 24 years. 23 participants stated their preferred pronoun as she/her, 11 chose he/him and 1 chose they/their.

3.1.2 Design and Materials

For the production task, each participant was exposed to both scenarios, Briefing and Interrogation. The order of scenario was counterbalanced across participants. Each critical item was presented to each participant once, either in the Briefing or Interrogation scenario, and paired either with evidence that we

estimated to be weak in evidentiality or evidence that we considered to be strong in evidentiality. However, for the analysis, we did not rely on this categorisation of evidential strength, but instead on the participants' evaluation of that in a post-test (see below). In order to ensure that participants saw a particular item only once but at the same time were exposed to both scenarios, we introduced two suspects. In this way, participants interrogated suspect1 and briefed about suspect2 or the other way round. This yielded 40 critical items: 10 briefing items per suspect and 10 interrogation items per suspect, each accompanied by either weak or strong evidence. Participants consequently saw 20 of these items—10 briefing items and 10 interrogation items, each scenario being about a different suspect. In each scenario block we used 20 filler items of which 10 were control items which functioned as attention checks. Within each scenario block the order of the items was randomised.

The critical items consisted of a picture containing a question/answer pair and a picture of a piece of evidence. The picture with the question/answer pair set the scene for the scenario manipulation: Either participants saw a picture of a briefing room facing a colleague, or they saw an interrogation room where they would be confronted with a suspect; see **Supplementary Material** for full details. The question in the heading of the picture was meant to be an already asked question by the participant. The answer to the question was provided by either the colleague or suspect, depending on the scenario, in form of a speech bubble. The participants were asked to react to the colleague's/suspect's answer by filling the gap in a sentence by choosing between 'know', 'notice', 'to be sure', 'think', 'believe', 'see', or 'guess', highlighted in the briefing (8) and (9) items. Alternatively, they were able to choose the option 'other' and formulate their own utterance. The order in which the expressions were displayed was randomised.

(8) Briefing item:

Participant: Did Emily Brown have any financial problems?
 Colleague: Financially the suspect was doing alright.
 Participant: I ___ that the suspect was in need of money.
 [know|believe|notice|amsure|think|see|guess]

(9) Interrogation item:

Participant: Did you have any financial problems?
 Suspect: Financially I was doing alright.
 Participant: I ___ that you were in need of money.
 [know|believe|notice|amsure|think|see|guess]

The manipulated pieces of evidence ranged from pictures to statements. As was mentioned above, we roughly manipulated the evidence to be weak or strong but will rely on the participants' evaluation of the evidence in our analysis. For example, for the items (8) and (9) we provided a bank statement as strong evidence and a statement of a friend who mentions potential, financial difficulties as weak evidence; see **Supplementary Material** for full details.

The filler items were turns between the suspect/colleague and the participant that had nothing to do with the case. For the 10 control items, the information was provided in the picture and

the participant had to choose the correct answer (here 11am), see (10). The option 'other' was also available.

(10) Control item: Picture: Clock in rooms says it is 11am.
 Participants: Oh look at the time. Is it already? [11am|noon|2pm]

After the production task, participants were asked to evaluate the pieces of evidence they had seen in both scenarios: *Given the piece of evidence below, how certain are you that p?*, where *p* is the complement proposition from earlier in the experiment. This would have been either the bank statement (strong evidence) or the statement of a friend (weak evidence), depending on which piece of evidence they had seen in the preceding scenario. Participants saw 20 pieces of evidence in total. Each piece of evidence dealt with a separate proposition which is why we assume that evidentiality ratings will not be influenced by anything else taking place in the experiment before we ask participants to rate the evidence. To communicate their certainty, participants adjusted a slider from 0 (not at all certain) to 100 (very certain).

3.1.3 Procedure

Before the experiment, participants were asked to give informed consent to take part in a fictional investigation of an art heist in the role of a detective. We also informed them about the structure of the experiment: (1) production task, engaging in two discussions, (2) rating evidence, (3) demographic questionnaire. Then the task was introduced in the form of a story about an art heist in Edinburgh involving two suspects. Since the lead detective on the case went missing, we asked the participants to help out solving the case. Both scenarios were introduced by stating that one of the suspects had been arrested. The participants were then asked either to prepare for the interrogation of the suspect with a colleague (Briefing scenario) or to interrogate the suspect right away (Interrogation scenario). Participants were instructed to converse with the colleague/suspect about different topics including questions about the case and they were told that for parts of the interaction they would need to look at the evidence that had been collected. In the Briefing scenario we then asked them to find the best way to help their colleague, and in the Interrogation scenario to find the best way to interrogate the suspect. After having completed the production task we asked participants to rate the quality of the pieces of evidence they had seen. The evaluation task was followed by the voluntary, demographic questionnaire. The experiment lasted approximately 35 min.

3.1.4 Analysis

We analysed our data fitting a Bayesian categorical regression model with maximal random effects structure using the R (R Core Team, 2020) package brms (Bürkner, 2018) which provides an interface to fit Bayesian mixed models using Stan (Stan Development Team, 2017). The Bayesian framework was chosen because the models with maximal random effects structure did not converge in the frequentist framework. Since

TABLE 1 | Raw counts for each utterance of the first experiment ordered from most frequent to least frequent.

Response	Counts
believe	185
know	171
think	110
see	80
notice	58
am sure	54
guess	14

Bayesian models can fall back to prior information they converge more easily. We chose not to analyse our data with a reduced effects structure within the frequentist framework to prevent inflated false positive effects (Barr et al., 2013).

The experimental factor scenario (Briefing/Interrogation) and the continuous variable evidentiality ([0,100]), were included to predict the probability of choosing 'believe'/'notice'/'am sure'/'think'/'see' over 'know'. This makes 'know' the reference category of the model. The factor scenario was sum-coded: -1 as Interrogation and 1 as Briefing. Evidentiality was standardised, such that the variable was centred at zero with a standard deviation of 1. The model included varying intercepts and slopes for participants and items, assuming that the effect of scenario and evidentiality on the participants' utterance choices varies between participants and items.

We used weakly regularising priors, which allowed a reasonably wide range of parameter values and at the same time penalised very extreme values. The priors for the by-expression intercepts were normal distributions with mean 0 and standard deviation 10. This means that we are 68% certain that the by-expression intercepts would fall within -10 and 10 on the log-odds scale which translates approximately to a range of 0 and 1 on the probability scale. For both fixed effects, we used normal priors with a mean of 0 and a standard deviation of 1. Random effects were modelled as a correlation matrix and a vector of standard deviations. The standard deviations were assigned half-normal priors with a mean of 0, and a standard deviation of 1. For the correlation matrix, a LKJ(2) prior was used such that smaller correlations are favored over extreme values such as ± 1 (Sorensen et al., 2016; Stan Development Team, 2017).

Samples were drawn from the posterior distributions of the model parameters using the NUTS sampler (Hoffman and Gelman, 2013). Four sampling chains were run, each collecting 4,000 iterations whereby the first 1,000 iterations were disregarded as part of the warm-up phase leading to 12,000 iterations available for analysis.

Unlike the frequentist analysis, the Bayesian analysis will not produce point estimates but instead posterior distributions over parameters quantifying the probability of each possible parameter value given the data. We will report the posterior mean $\hat{\beta}$ and the 95% credible interval (95%-CrI). The 95%-CrI is the range around the posterior mean within which the true value of the parameter lies with a probability of 0.95. We could roughly interpret the evidence as reliable if zero lies outside the parameters' 95% credible interval (Kruschke et al., 2012).

The response 'other' was excluded from the analysis. We included 'other' mainly to give participants more freedom in their utterance choice and to create a more natural experience. The response 'other' made up only 8% of the data.

3.1.5 Results

The accuracy of participants was at 95% for the control items which suggest that they paid attention during the experiment. Overall, participants used a wide range of expressions, see **Table 1** for details. 'Guess' was attested much less frequently than the other options, so we omit it from the detailed statistical analysis.

In the evaluation task, participants assessed the evidence by using the whole range of the slider: the evidentiality ratings ranged from 0 to 100, with a mean of 72. Furthermore, evidentiality ratings varied between utterances. Participants used 'know' when they saw fairly convincing evidence (mean evidentiality rating was 81), whereas they used 'think' for weaker evidence (mean evidentiality rating was 62); see **Table 2** for details.

Contrasting the by-expression evidentiality ratings for both scenarios reveals that, overall, the median evidentiality ratings across expressions were lower for events mentioned in the interrogation than in the briefing; see **Figure 1**. The only exception is 'am sure'. If we were to assume that the median evidentiality ratings give us insight to the thresholds of the individual expressions, the following order of expression threshold for each scenario emerges:

Briefing :

know > notice > see > believe > am sure > think

Interrogation :

know > see > am sure > notice > believe > think

Compared to the briefing, 'see' and 'am sure' are much higher ranked in the interrogation, whereas 'notice' and 'believe' move down in the ranking. The differences in order between the two scenarios would be surprising if one were to strictly assume that speakers choose the expression with the highest possible threshold compatible with their degrees of belief.

For the expressions 'notice', 'am sure', 'think' and 'believe', an increase in evidentiality meant a decrease in log-odds. This effect was most pronounced for 'think' ($\hat{\beta} = -1.13$, CrI: [-1.55, -0.73]), and weakest for 'notice' ($\hat{\beta} = -0.38$, CrI: [-0.88, 0.18]), see **Table 3** for all expressions. This suggests that averaging over scenarios, speakers are less likely to choose 'notice', 'am sure', 'think' or 'believe' than 'know' when they are highly certain about the content they want to communicate. The effect was absent for the

TABLE 2 | Mean evidentiality ratings by utterance for the experiment 1.

Response	Mean	SD	Median	Min	Max
know	81	19	83.0	16	100
notice	72.3	23.3	71.5	15	100
am sure	73	22.3	75.5	20	100
see	79.3	22.2	81.0	0	100
think	62.2	28.2	63.5	0	100
believe	67.4	22.5	70	0	100
guess	54.2	18.7	60	15	74

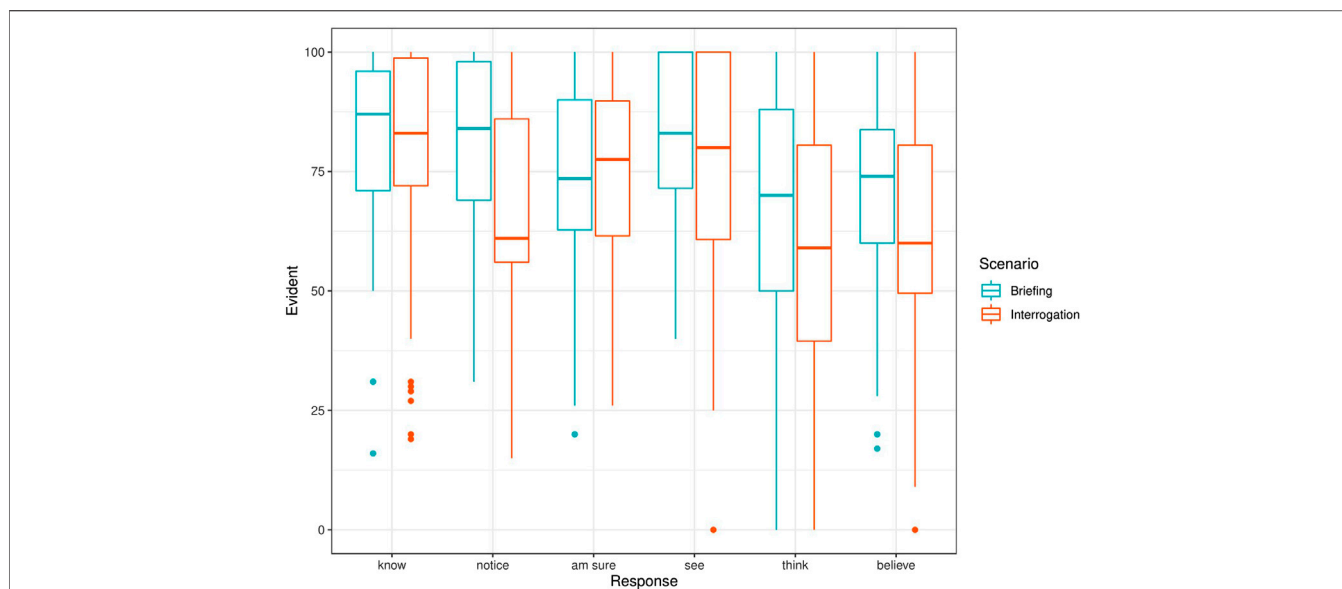


FIGURE 1 | By-expression evidentiality ratings for each scenario (briefing in blue, interrogation in red). The plot shows the median of the evidentiality ratings (line) and the upper quartile and lower quartile (box). Whiskers extend to the scores outside the quartiles. Dots represent outliers.

TABLE 3 | Population-level estimates of the categorical regression model in log-odds with the standard errors and 95% credible intervals.

Utterance	Coefficient	Posterior mean	Est. Error	l-95% CrI	u-95% CrI	\hat{R}
notice	Intercept	-1.39	0.31	-2.06	-0.84	1.00
am sure	Intercept	-2.03	0.46	-3.05	-1.26	1.00
see	Intercept	-0.95	0.24	-1.47	-0.52	1.00
think	Intercept	-0.70	0.26	-1.25	-0.21	1.00
believe	Intercept	0.02	0.23	-0.43	0.47	1.00
notice	Evidentiality	-0.38	0.27	-0.88	0.18	1.00
am sure	Evidentiality	-0.62	0.26	-1.12	-0.10	1.00
see	Evidentiality	0.00	0.22	-0.41	0.44	1.00
think	Evidentiality	-1.13	0.20	-1.55	-0.73	1.00
believe	Evidentiality	-0.77	0.21	-1.19	-0.36	1.00
notice	Scenario1	0.18	0.19	-0.19	0.56	1.00
am sure	Scenario1	0.84	0.26	0.35	1.39	1.00
see	Scenario1	0.12	0.17	-0.21	0.44	1.00
think	Scenario1	0.69	0.17	0.36	1.02	1.00
believe	Scenario1	0.47	0.15	0.18	0.77	1.00

In the table the by-expression intercepts are listed first, then the estimates for the evidentiality effect followed by the estimates for the scenario effect. The effect scenario is the change in log-odds for the briefing (-1 interrogation, 1 briefing). \hat{R} is a convergence diagnostic which compares the between- and within-chain estimates. Values larger than 1 suggest that the chains have not mixed well.

expression ‘see’ ($\hat{\beta}=0$, CrI:[-0.41, 0.44]) which suggests that, averaging over scenarios, speakers are free to choose ‘know’ or ‘see’ when they are highly certain.

The effect of scenario was overall less pronounced than the effect of evidentiality. For the expressions ‘am sure’, ‘think’ and ‘believe’, the probability to be chosen increased in the briefing and decreased in the interrogation. This effect was most pronounced for ‘am sure’ ($\hat{\beta}=0.84$, CrI:[0.35, 1.39]), followed by ‘think’ ($\hat{\beta}=0.69$, CrI:[0.36, 1.02]) and ‘believe’ ($\hat{\beta}=0.47$, CrI:[0.18,

0.77]). This suggests that, given average evidentiality, speakers choose ‘am sure’, ‘think’ or ‘believe’ in preference to ‘know’ in the briefing, but less so in the interrogation. The effect was absent for ‘notice’ ($\hat{\beta}=0.18$, CrI:[-0.19, 0.56]) and ‘see’ ($\hat{\beta}=0.12$, CrI:[-0.21, 0.44]). This suggests that given average evidentiality, the extent of speakers’ preference for ‘know’ over ‘see’ and ‘notice’ is independent of scenario.

Overall, the findings suggests that in the briefing, stronger evidence was needed in order for participants to choose expressions associated with a high degree of certainty such as ‘know’. In the interrogation, however, less convincing evidence might have sufficed. Furthermore, evidentiality was found to play a bigger role than scenario in expression choice.

3.1.6 Interim Summary

The results of the first experiment suggest that speakers indeed base their utterance choices on their degrees of belief and moreover adjust their choices depending on the scenario they are in. Overall, participants chose (un)certainly expressions associated with lower degrees of belief more often in the briefing than in the interrogation. For example, ‘think’, and ‘believe’ received lower mean evidentiality ratings than ‘know’ (62 and 67 respectively, see Table 2), and were more likely to be used in the briefing than in the interrogation. There were no differences between scenarios for ‘know’ and ‘see’ and the mean evidentiality ratings for ‘see’ (79.3) were almost as high as the ones for ‘know’ (81). The results of ‘am sure’ and ‘notice’ are less straightforward to interpret: While ‘am sure’ was more often used in the briefing than in the interrogation, we find no effect of scenario for ‘notice’, despite both expressions having received almost the same evidentiality ratings (73 and 72 respectively).

On the face of it, the results of the first experiment suggest a complex pattern of usage preferences without clear indications of the kind of stratification of choices that would be predicted by threshold semantics account. However, it is possible that this reflects the crowded space of possible options, and differences of opinion between participants as to the relative strength of, for instance, ‘believe’ and ‘am sure’. Hence, in the following experiment, we constrain the space of choices further in a bid to obtain clearer results. Specifically, we consider ‘know’ in comparison with ‘believe’ across a wide range of confidence levels (‘know’ for higher, and ‘believe’ for lower confidence levels). We use ‘believe’ rather than ‘think’ because ‘believe’ was chosen by participants more frequently.

3.2 Experiment 2: ‘Know’ Versus ‘Believe’

In the second experiment, we aimed to investigate the usage of (un)certainly expressions further by focusing on ‘know’ and ‘believe’ based on the results of the first experiment. The experimental design, materials and procedure are the same as for the first experiment with the difference that participants could only choose between ‘know’, ‘believe’ and ‘other’. Similar to the first experiment our hypothesis is that speakers decide between ‘know’ and the epistemically weaker option ‘believe’, depending on (i) their degree of confidence, (ii) and the communicative setting they are in.

3.2.1 Participants

We tested 85 participants recruited over the crowd-sourcing platform Prolific, specifying participants with an approval rate above 90 and restricted to people that had not previously participated in experiment 1. Participants were paid with an average of £7.53/h (the average duration of the experiment was 30 min). After data collection two participants were excluded because their accuracy for the control items was below chance level. The age of the remaining 83 participants ranged from 18 to 66 years, with a mean of 32 years 50 participants stated their preferred pronoun as she/her, 32 chose he/him and 1 chose they/their.

3.2.2 Materials

The materials were the same as for the first experiment, see **Section 3.1.2**. The only difference was that the utterance choice was limited to ‘believe’ and ‘know’, see (11) as an example for the Briefing scenario.

- (11) Briefing item: Participant: Did Emily Brown have any financial problems? Colleague: Financially the suspect was doing alright. Participant: I that the suspect was in need of money. [know|believe]

3.2.3 Procedure

The procedure matched that from the first experiment, see **Section 3.1.3**.

3.2.4 Analysis

We analysed our data fitting a Bayesian binary logistic regression model with maximal random effects structure using the R (R Core Team, 2020) package brms (Bürkner, 2018). The experimental

TABLE 4 | Mean evidentiality ratings by utterance for the experiment.

Response	Mean	SD	Median	Min	Max
know	83.45	27.12	91	0	100
believe	64.67	27.12	70	0	100

factor scenario (briefing/interrogation) and the continuous variable evidentiality ([0,100]), were included to predict the probability to choose ‘know’. The factor scenario was again sum-coded (−1 as interrogation and 1 as briefing) and evidentiality was standardised. The model included varying intercepts and slopes for participants and items.

We used the same priors as for the data of the first experiment: The intercept was normal distribution with mean 0 and standard deviation 10. For the fixed effects, we used normal priors with a mean of 0 and a standard deviation of 1. The standard deviations were assigned half-normal priors with a mean of 0, and a standard deviation of 1 and for the correlation matrix, we used a LKJ(2) prior.

The sampling process was the same as for the experiment 1 analysis, see **Section 3.1.4** for a detailed description. The response ‘other’ was again excluded from the analysis. As anticipated, providing participants with a reduced set of utterances, ‘other’ was chosen more frequently than in the first experiment and made up 19% of the data. After examining the data, ‘other’ was often chosen when evidentiality ratings were rather low. We will leave the analysis of these responses for future research but will not focus on them in this study.

3.2.5 Results

Two participants were excluded for below-chance performance on the control items. The remaining participants responded to the control items with an accuracy of 97%.

Similar to the first experiment, in the evaluation task, participants assessed the evidence by using the whole range of the slider: the evidentiality measure ranged from 0 to 100, with a mean of 73. Overall, on average ‘know’ (mean 83) was chosen for higher evidentiality measures than ‘believe’ (mean 65), see **Table 4** for more information. Looking at the by-expression evidentiality ratings contrasting both scenarios, it seems that the evidentiality ratings across expressions were lower for events mentioned in the interrogation than in the briefing; see **Figure 2**. The difference seems to be more pronounced for ‘believe’. This suggests that in the briefing, stronger evidence was needed in order for participants to choose ‘know’. In the interrogation, however, less convincing evidence might have sufficed.

These observations are supported by the outcome of our analysis. The estimate of the main effect of evidentiality was ($\hat{\beta}=1.34$, CrI:[1.02, 1.70]), suggesting that the probability of saying ‘know’ increases when the speakers’ degrees of belief increase, see **Table 5** for more details. The estimate of the scenario effect was ($\hat{\beta}=-0.29$, CrI:[−0.50, −0.10]), suggesting that speakers were more likely to say ‘know’ in the Interrogation scenario than in the Briefing scenario. Thus, we could replicate the effects we found in the first experiment.

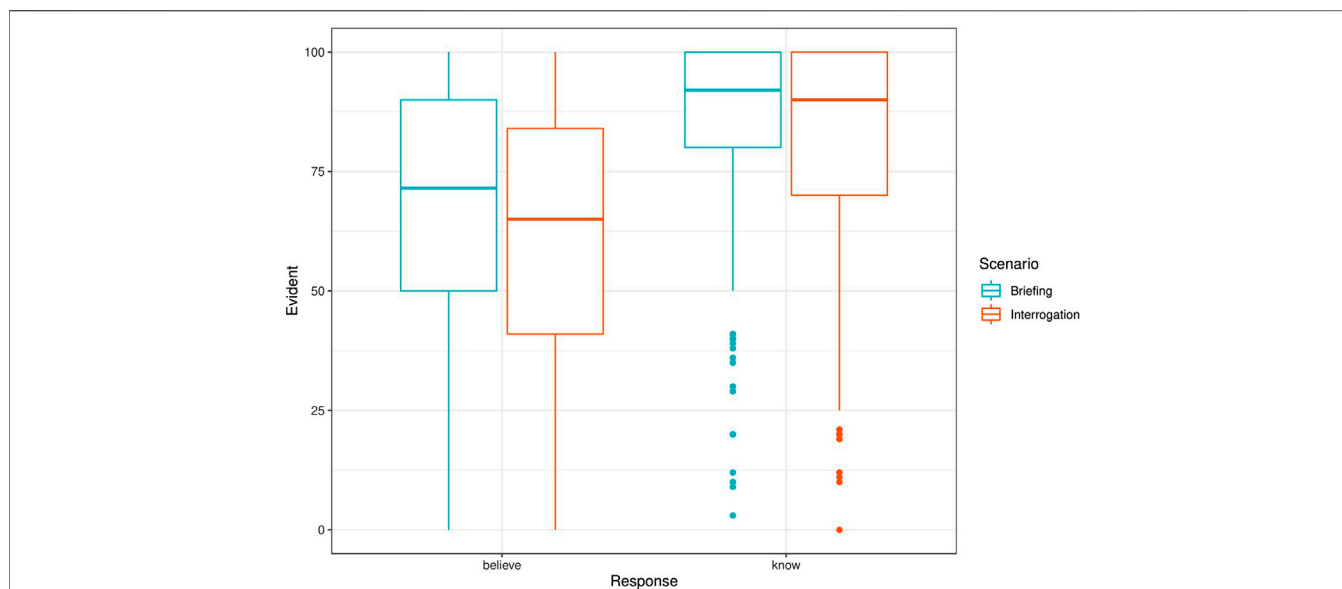


FIGURE 2 | By-expression evidentiality ratings for each scenario (briefing in blue, interrogation in red). The plot shows the median of the evidentiality ratings (line) and the upper quartile and lower quartile (box). Whiskers extend to the scores outside the quartiles. Dots represent outliers.

TABLE 5 | Population-level estimates of the categorical regression model in log-odds with the standard errors and 95% credible intervals.

Coefficient	Posterior mean	Est. Error	l-95% CrI	u-95% CrI	\hat{R}
Intercept	-0.48	0.19	-0.87	-0.10	1.00
Evidentiality	1.34	0.17	1.02	1.70	1.00
Scenario	-0.29	0.10	-0.50	-0.10	1.00

The effect scenario is the change in log-odds for the briefing (-1 interrogation, 1 briefing). \hat{R} hat is a convergence diagnostic which compares the between- and within-chain estimates. Values larger than 1 suggest that the chains have not mixed well.

We visualised the models predictions given our data in **Figure 3**. In order to facilitate understanding we back-transformed the data from log-odds to probabilities. The plot shows the predicted probabilities of using ‘know’ for each scenario. The x-axis represents the standardised evidentiality ratings, whereby 0 means average evidentiality (corresponding to 73 on the original scale). The lines represent the means of the fixed effects and the faded area the 95% credible intervals for the effects. The plot illustrates that with increasing evidentiality, the probability to use ‘know’ is predicted to increase as well. Furthermore, ‘know’ is predicted to be more likely to be used in the Interrogation scenario than in the Briefing scenario. For example, with an average evidentiality rating (0 in the plot), the probability of using ‘know’ is predicted to be approximately 0.31 in the briefing and approximately 0.45 in the interrogation.

However, inspecting credible intervals of the posterior means, we can see that the estimated distribution for evidentiality is further away from 0 than the estimated distribution of scenario. This suggests that evidentiality plays a bigger role for speakers deciding between ‘know’ and ‘believe’ than the scenario they are in.

3.2.6 Discussion

The findings of the second experiment replicated the findings from the first experiment: speakers seem to base their utterance choices on their degrees of believe and moreover to adjust their choices depending on the scenario they are in. For experiment 2 this means concretely that participants chose ‘believe’ for lower evidentiality ratings than ‘know’ and that they were more likely to choose ‘believe’ in the briefing than in the interrogation. One difference between the two experiments lies in the evidentiality ratings for ‘believe’ which were lower for the second experiment than for the first. It is reasonable to assume that participants were willing to use ‘believe’ for lower degrees of confidence in the second experiment due to the lack of explicitly suggested alternatives.

4 General Discussion

There are two different ways to interpret our findings. First, they could suggest that in a cooperative scenario speakers want to truthfully communicate their certainty to be as informative and cooperative as possible. In such a scenario a speaker chooses a particular expression based on their degrees of belief such that, if they deem the event probability to exceed the threshold of an expression, the speaker chooses that expression as long as there is no other, more informative expression. In contrast, in a scenario where speakers are faced with an uncooperative interlocutor who might not tell the truth, speakers might instead act strategically, and use expressions that are associated with higher certainties such as ‘know’ or ‘see’ in order to come across as authoritative and confident (Hosman, 1989; Crismore and Kopple, 1997).

Alternatively, our findings could suggest that in a cooperative scenario speakers are obliged to engage to some extent in polite

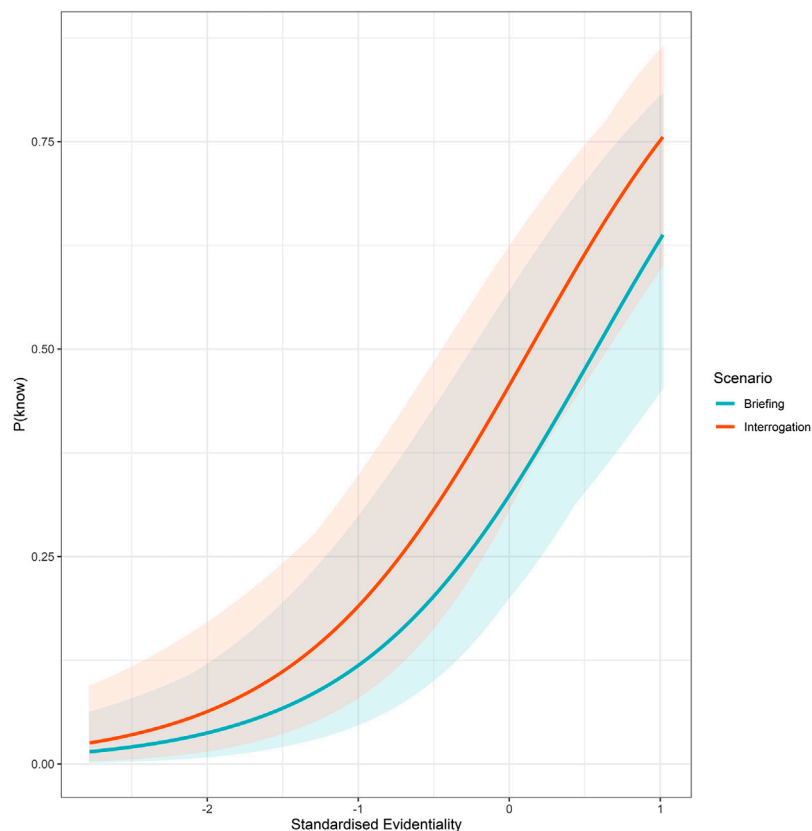


FIGURE 3 | Predictions of our model given our data. Log-odds were back-transformed to probabilities (y-axis). The x-axis is the standardised evidentiality measure: 0 stands for an evidentiality of 73. An increase of one standard deviation on the standardised scale means an increase of 26 on the original scale. The lines represent the means of the fixed effects and the faded area depicts the 95% credible interval of the fixed effects.

conversation. This would mean that, besides pursuing the aim to be cooperative and most informative, speakers also engage in facework. Hence, a briefing with a colleague would not only entail acting cooperatively and conveying information to each other but also trying not to step on each other's toes. Recall that, in hedging accounts and politeness theory (e.g., Fraser, 1975; Brown and Levinson, 1987), this means that cooperative speakers tend to downtone their utterances when communicating their certainty in order to be polite. By contrast, in an uncooperative scenario speakers might feel less obliged to engage in facework especially when they are in a high-power position. Thus, instead of hedging their statements, speakers communicate their degrees of belief directly.

While we incline toward the first explanation, we acknowledge that the second explanation is hard to rule out experimentally, given the putatively ubiquitous nature of facework. Although the Briefing scenario is designed to be cooperative, in the sense of prioritising the exchange of accurate information, politeness considerations are doubtless still at play to some extent. These considerations may mediate between the speaker's degree of belief and choice of (un)certainly expression in a way that has yet to be fully theorised. This is due to the fact that there is no traditional baseline condition. We purposely chose a detective story as a cover story because it enabled us to create a coherent experiment where being part of a cooperative and an uncooperative scenario

is plausible for the participant. Even in a more neutral context, such as talking to a friend, conventions such as being polite will be involved to some extent so it is unclear whether this would constitute a suitable baseline condition.

In practice, our results suggest that speakers use (un)certainly expressions much more dynamically than expected by a strict threshold semantics account, if that account is coupled with standard pragmatic assumptions about cooperativity. For example, we did not find a fixed ordering of (un)certainly expressions. Recall that, under these assumptions, we would expect each expression to be restricted to evidentiality levels between its threshold and the threshold for the next stronger expression. For instance, given the choices of 'think' and 'know', we would expect to see 'think' attested for evidentiality above the threshold for 'think' and below the threshold for 'know'. However, we did not find a fixed ordering of the kind this model would predict. Considering the median evidentiality rankings for each (un)certainly expression from experiment 1, repeated below, we found differences in their ordering between scenarios.

Briefing :

know > notice > see > believe > am sure > think

Interrogation :

know > see > am sure > notice > believe > think

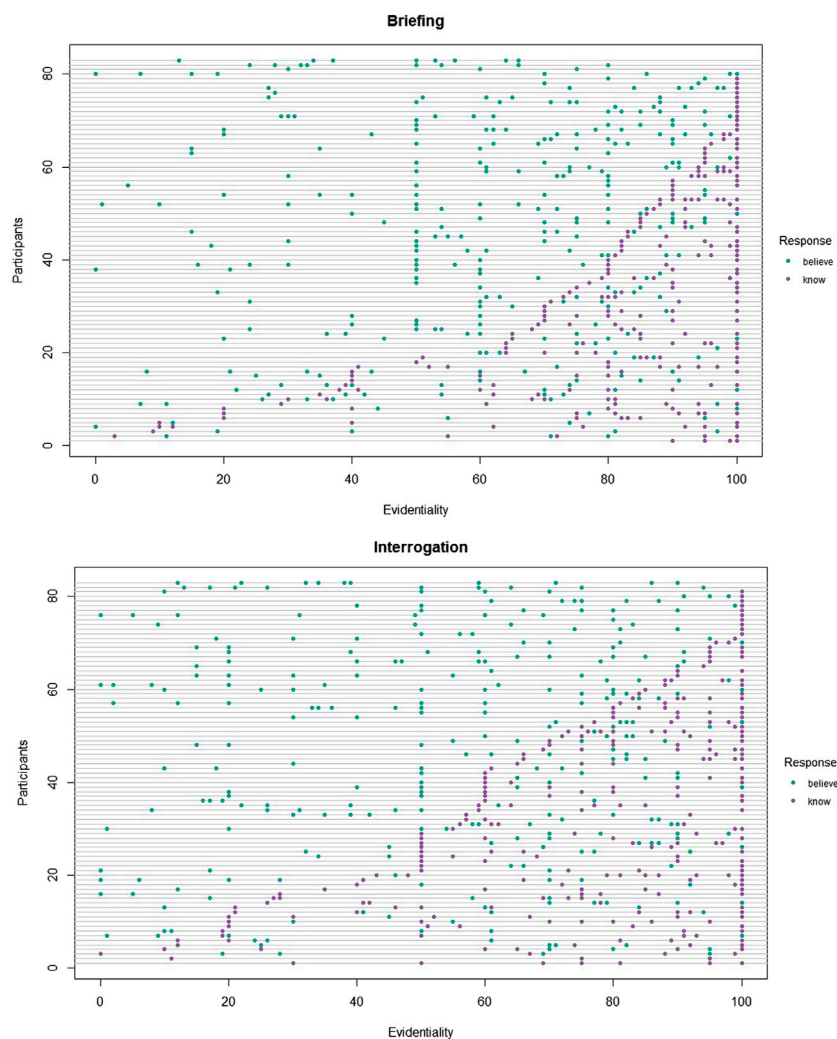


FIGURE 4 | The plots show the by-participant usage of ‘know’ (purple dots) and ‘believe’ (green dots) in the two different scenarios. The y-axis represents the individual participants, whereby each gray line belongs to one participant. Participants are ordered from top to bottom according to the lowest degree of belief for which they used ‘know’. The x-axis represent the evidentiality measure. This way we can see the degrees of belief of the participants when choosing a particular expression. A green dot that appears immediately to the right of a purple dot indicates a case where a speaker used ‘believe’ for stronger evidence than for which they elsewhere used ‘know’.

As the statistical analysis above suggests, many participants appeared to be willing to use ‘know’ for lower evidentiality ratings in the Interrogation scenario than in the Briefing scenario. This could be accommodated within a threshold-based account by assuming that an individual’s threshold for the use of ‘know’ can vary between scenarios, just as thresholds are argued to vary between speakers in general. However, a more challenging result for the threshold-based account is that participants frequently used ‘believe’ for levels of evidentiality which exceeded those for which they elsewhere used ‘know’. This is shown in **Figure 4** where we plotted the by-subject utterance choices across evidentiality ratings.

These findings could be reconciled with the threshold-based account in a couple of ways. One possibility is to assume that an

individual’s threshold for using ‘know’ varies throughout the experiment (or that the participants’ reported evidentiality ratings do not correspond with the ratings on which they based their productions at the time of utterance). Another possibility is that speakers are simply not pragmatic in the way we have assumed, and freely produce less informative utterances than they are entitled to (i.e., asserting ‘believe’ in just some specific circumstances when they could assert ‘know’), or produce more informative utterances than are warranted (i.e., asserting ‘know’ when they do not have sufficient confidence, by their own criteria, to license this). However, it should be noted that the claims of threshold semantics are difficult to falsify if we are not committed to speakers being pragmatic in this kind of way: one could simply posit lower

thresholds. That is to say, without this kind of assumption about a speaker's pragmatic behavior, the explanatory usefulness of a threshold semantic account is vitiated.

In terms of the broader implications for theories of communication, our results suggest that speakers choose an utterance not only based on their perception of the world (here their degrees of belief) but also on the effect that their communicative action may have on the hearer. Sperber and Wilson (1986/1995) claim that speakers' intention is twofold: speakers (i) want to be understood (informative intention), and (ii) aim to convince their hearers to think or act in accordance with the speakers' beliefs (communicative intention). In particular, our results suggest that speakers are willing to make stronger claims than appear to be semantically warranted, in a context in which doing so could be communicatively effective in the latter sense.

This in turn raises the question of how a hearer should act in such a case—that is, where they suspect that a speaker is overstating the probability that a proposition is true, according to the usual assumptions about the meanings of (un)certainty expressions. Sperber et al. (2010) suggested that interlocutors automatically engage an 'epistemic vigilance' mechanism whose purpose is to detect misinformation, by assessing the quality and plausibility of the communicated content as well as the reliability of the speaker. In the Interrogation scenario, it would be rational for a hearer to engage in an even more conscious evaluation of credibility—that might go beyond the automatic application of the epistemic vigilance mechanism—and opt out of the usual cooperative assumption that the speaker's claims are to be accepted at face value.

Research on epistemic vigilance (e.g., Mazzarella et al., 2018) has typically focused on cases where a proposition is categorically accepted or rejected, so it is potentially interesting to consider what epistemic vigilance means in the context of gradable degrees of belief. Confronted with a claim that the speaker 'knows' *p*, should the vigilant hearer conclude that the speaker in fact merely 'believes' *p*, or should they reject it outright?

Moreover, keeping in mind that hearers are vigilant, speakers have to be cautious and check to what extent their utterance coheres with the beliefs of the hearer in order to make it more probable for their communicative intention to succeed (Sperber et al., 2010). In our experiment, participants might have reasoned similarly following an overall goal of gaining compliance (e.g., Dillard, 1990). In the Briefing scenario, for example, they might have deemed it strategically advantageous to occasionally communicate their degrees of belief truthfully without under- or overstating their confidence. Here, the aim of truthfully communicating their confidence would be to prove their reliability which could facilitate their overall goal: briefing their colleague successfully. Similarly, in the Interrogation scenario, the participant might wish to be perceived as reliable, in order to avoid the suspect engaging in epistemic vigilance and thus discounting subsequent overstated claims. Thus, throughout the experiment, there could be higher-order strategies in play, which might explain participants' varied use of (un)certainty expressions within one scenario—the speaker may wish to give

the hearer a particular impression of how they use uncertainty expressions, much like how a good poker player will fold on some bad hands in order to make a subsequent bluff more effective. This could explain participants' varied usage of (un)certainty expressions within one scenario. However, these claims have to be tested further.

Moreover, it remains to be investigated whether speakers can employ these strategies successfully, especially in the interrogation setting where the speakers' goal is to convince the hearer and the hearer has a great incentive to hide the truth. Future research may also shed light on what parts of the communicated content hearers might refuse. Overall, our study illustrates participants' expectations of what someone else would say in a given situation, and as such, indirectly gives insights into the comprehension of (un)certainty expressions. In the end, speakers and hearers are both interlocutors who inevitably reason about each other and constantly exchange their roles within their dialogue.

In the broader context of uncertainty, we could see these interactions as attempts to reduce the interlocutors' uncertainty about the truth-values of various propositions under discussion. From this point of view, a police officer's interrogation of a suspect could be seen as an attempt to eliminate their subjective cognitive uncertainty as to whether the suspect is guilty. The use of an expression that conveys a false degree of confidence in the factuality of a proposition—e.g., 'I know you were there'—could be one way of seeking confirmation of a hitherto uncertain claim, and thus building a consensus as to what is true. The briefing of a colleague, by contrast, does not aim at the complete elimination of uncertainty, but toward the building of a consensus between the interlocutors as to how (un)certain various propositions subjectively should be, given the available evidence. This difference potentially underlies the difference in communicative strategies evident in our experiments.

5 CONCLUSION

In this paper, we tested the production of (un)certainty expressions in two contrasting scenarios. The first experiment contrasted a wide range of (un)certainty expressions, and the second experiment focused on the production of 'know' vs. 'believe'. We found that speakers choose between (un)certainty expressions based on their degrees of belief, and furthermore, adjust their choices depending on the scenario they are in. These findings supports hedging and politeness accounts which assume that speakers may use (un)certainty expressions strategically. By contrast, our findings are surprising under a strict threshold semantics accounts paired with pragmatic assumptions about cooperativity, since we found that speakers use (un)certainty expressions much more freely. Besides using 'know' for lower evidentiality ratings in the Interrogation scenario than in the Briefing scenario, participants frequently used 'believe' for levels of evidentiality which exceeded those for 'know'.

DATA AVAILABILITY STATEMENT

The full set of data for both experiments is available here: <https://osf.io/e5av9/>.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the PPLS Research Ethics Committee (School of Philosophy, Psychology and Language Sciences). The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

AL, CC, and HR contributed to conception and design of the study. AL collected the data and performed the statistical analysis. AL wrote the first draft of the manuscript. CC and HR wrote sections of the manuscript. All authors contributed

to manuscript revision, read, and approved the submitted version.

FUNDING

This work was supported by the Scottish Graduate School for Arts and Humanities in conjunction with Scottish Funding Council.

ACKNOWLEDGMENTS

We are grateful to the reviewers for their helpful and constructive comments.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fcomm.2021.635156/full#supplementary-material>.

REFERENCES

- Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: keep it maximal. *J. Mem. Lang.* 68, 255–278. doi:10.1016/j.jml.2012.11.001
- Bürkner, P.-C. (2018). Advanced Bayesian multilevel modeling with the r package BRMS. *R. J.* 10, 395–411. doi:10.32614/rj-2018-017
- Berger, C. R., and Calabrese, R. J. (1975). Some explorations in initial interaction and beyond: toward a developmental theory of interpersonal communication. *Hum. Commun. Res.* 1, 99–112.
- Berger, C. R. (1995). *Communication and social influence processes. Inscrutable goals, uncertain plans, and the production of communicative action*. Michigan: Michigan State University Press, 1–28.
- Bonnefon, J.-F., and Villejoubert, G. (2006). Tactful or doubtful? Expectations of politeness explain the severity bias in the interpretation of probability phrases. *Psychol. Sci.* 17, 747–751. doi:10.1111/j.1467-9280.2006.01776.x
- Brashers, D. E. (2006). Communication and uncertainty management. *J. Commun.* 51, 477–497.
- Brown, P., and Levinson, S. (1987). *Politeness. Some universals in language usage*. Cambridge: Cambridge University Press.
- Brown, P., and Levinson, S. (1978). “Universals in language usage: politeness phenomena,” in *Questions and politeness. Strategies in social interactions*. Editor E. Goody (Cambridge: Cambridge University Press), 56–311.
- Clark, H. H. (1990). Quantifying probabilistic expressions: Comment. *Stat. Sci.* 5, 12–16.
- Crismore, A., and Kopple, W. J. V. (1997). *Hedging and discourse. Hedges and readers: effects on attitudes and learning*. Berlin, Boston: De Gruyter, 83–114.
- Dillard, J. (1990). *A goal-driven model of interpersonal influence. Seeking compliance: the production of interpersonal influence messages*. Berlin: Gorsuch Scarisbrick, 41–56.
- Fraser, B. (1975). “Hedged performatives,” *Syntax and semantics*. Editors P. Cole and J. L. Morgan (New York: Academic Press), 187–210.
- Fraser, B. (2010). “Pragmatic competence: the case of hedging,” in *New approaches to hedging*. Editors G. Kaltenböck, W. Mihatsch, and S. Schneider (London: Emerald), 15–34.
- Grice, H. P. (1975). “Logic and conversation,” in *Syntax and semantics. Speech acts*. Editors P. Cole and J. L. Morgan (New York: Academic Press), 41–58.
- Hoffman, M. D., and Gelman, A. (2013). The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Machine Learn. Res.* 15, 1593–1623. doi:10.1186/s12711-019-0515-1
- Holmes, J. (1982). Expressing doubt and certainty in English. *RELC J.* 14, 9–28.
- Holtgraves, T., and Perdue, A. (2016). Politeness and the communication of uncertainty. *Cognition* 154, 1–10. doi:10.1016/j.cognition.2016.05.005
- Hosman, L. A. (1989). The evaluative consequences of hedges, hesitations, and intensifiers powerful and powerless speech styles. *Hum. Commun. Res.* 15, 383–406.
- Juanchich, M., and Sirota, M. (2015). A direct and comprehensive test of two postulates of politeness theory applied to uncertainty communication. *Judgment Decis. Making* 10, 232–240. doi:10.4135/9781412959384.n286
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Knapp, P., Raynor, D. K., and Berry, D. C. (2004). Comparison of two methods of presenting risk information to patients about the side effects of medicines. *BMJ Qual. Saf.* 13, 176–180. doi:10.1136/qshc.2003.009076
- Kruschke, J. K., Aguinis, H., and Joo, H. (2012). The time has come: Bayesian methods for data analysis in the organizational sciences. *Org. Res. Methods* 15, 722–752. doi:10.1177/1094428112457829
- Lakoff, G. (1973). Hedges: a study in meaning criteria and the logic of fuzzy concepts. *J. Philos. Logic* 2, 458–508.
- Lassiter, D. (2017). *Graded modality: qualitative and quantitative perspectives*. Oxford: Oxford University Press.
- Littlejohn, S. W., Foss, K. A., and Oetzel, J. G. (2017). *Theories of human communication*. 11 edn. Long Grove, Illinois: Waveland Press.
- Lorson, A., Cummins, C., and Rohde, H. (2019). “When objecting to presupposed content comes easily,” Proceedings of the 23rd workshop on the semantics and pragmatics of dialogue. London: SemDial, 54–60.
- Mazzarella, D., Trouche, E., Mercier, H., and Noveck, I. (2018). Believing what you’re told: politeness and scalar inferences. *Front. Psychol.* 9, 90. doi:10.3389/fpsyg.2018.00908
- R Core Team (2020). *R: a language and environment for statistical computing*. Vienna, Austria: SemDial.
- Schuster, S., and Degen, J. (2020). I know what you’re probably going to say: listener adaptation to variable use of uncertainty expressions. *Cognition* 203, 175–200. doi:10.31234/osf.io/8w6xc
- Sorensen, T., Hohenstein, S., and Vasisht, S. (2016). Bayesian linear mixed models using stan: a tutorial for psychologists, linguists, and cognitive scientists. *Quant. Methods Psychol.* 12, 175–200. doi:10.20982/tqmp.12.3.p175

- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., et al. (2010). Epistemic vigilance. *Mind Lang.* 25, 359–393.
- Sperber, D., and Wilson, D. (1986/1995). *Relevance: communication and cognition*. Hoboken: Blackwell.
- Stan Development Team (2017). *Stan modeling language users guide and reference manual. Version 2.17.0*.
- Wallsten, T. S., Budescu, D. V., Rapoport, A., Zwick, R., and Forsyth, B. (1986). Measuring the vague meanings of probability terms. *J. Exp. Psychol.* 4, 348–365.
- Yalcin, S. (2010). Probability operators. *Philos. Compass* 11, 916–937. doi:10.1111/j.1747-9991.2010.00360.x

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Lorson, Cummins and Rohde. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.