

1 We might be wrong, but we think that hedging doesn't protect your reputation

2

3

4

5

Peter J. Collins¹

6

Ulrike Hahn^{1,2}

7

8

¹Munich Center for Mathematical Philosophy

9

²Birkbeck, University of London

10

11

Author Note

12

Peter Collins is now at the Department of Psychology, Goldsmiths, University

13

of London

14

15

Corresponding Author:

16

Dr Peter Collins

17

Department of Psychology

18

Goldsmiths, University of London

19

New Cross

20

London

21

SE14 6NW

22

p.j.collins.65@gmail.com

23

24

Copies of the surveys, data files and analysis scripts can be found on the Open

25

Science Framework at osf.io/r9cna.

26

1

Abstract

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

We gain much of our knowledge from other people. Since people are fallible - they lie, mislead, and are mistaken - it seems essential to monitor their claims and their reliability as sources of information. An intuitive way to do this is to draw on our expectations about claims and sources: to perform expectation-based updating (Hahn, Merdes, & von Sydow, 2018). But this updating can have damaging consequences, leading us into a kind of confirmation bias. An alternative is to keep track of outcomes and record whether a claim proves true or false: to perform outcome-based updating (Hahn et al., 2018). This form of updating does not have the negative repercussions on belief accuracy. But both forms of updating might undermine the trust and cooperation assumed to be necessary for successful communication. We explore a potential boundary condition on these types of updating. We investigate whether speakers can protect their reputation when they make claims with low prior probability, with and without knowledge of the final outcome. We explore suggestions from McCready (2015) that speakers can protect themselves by hedging with evidential language: in particular with weaker propositional attitudes ('I suspect that...') and so-called double hedges ('I might be wrong, but I think...'). We find that both forms of updating are robust to hedging with this evidential language, and find no clear evidence for a protective effect. We discuss extra ingredients that may allow successful hedging.

21

Keywords: cooperation; hedging; pragmatics; testimony; belief change

22

23

24

1

2 Each day we receive information from other people, from the mundane
3 ("There is heavy traffic on the motorway") to the highly consequential ("This pension
4 scheme will give you security in retirement"). We can use this information to form or
5 change our beliefs. When we update our beliefs in this way, we are engaging with
6 people's testimony: with their saying, telling, or asserting; their committing to the
7 truth of some proposition intending that a recipient (or recipients) will accept the
8 proposition as true (see, e.g., Adler, 2015; Coady, 1992). Testimony seems both
9 ubiquitous and essential¹. Most of us have not experienced that the world is spherical,
10 witnessed the events in the news, collected evidence of historical events, or studied
11 the safety of medical treatments. But much of the time, we will accept other people's
12 testimony on such matters. Testimony allows a division of labour: we all want to
13 know how to treat our ailments, but we cannot all be doctors. But it also raises a
14 normative question: when should we trust, and when distrust, our information
15 sources?

16 Information sources can clearly be unreliable, leaving us at risk of being
17 misinformed, misled, or deceived. Sources can, for instance, be well-intentioned but
18 inaccurate: people forget, and their memories are reconstructive (for discussion, see
19 Hahn, Oaksford, & Harris, 2012; on reconstructive memory, see, e.g., Loftus, 1975).
20 Sources can also be dishonest: people routinely engage in small acts of deception,
21 cheating when they can do so without being caught or without challenging their self-

¹ This section draws substantially on (Collins, 2017).

1 conception (see, e.g., Chance, Norton, Gino, & Ariely, 2011; Mazar, Amir, & Ariely,
2 2008).

3 Systematic unreliability seems to compel us to be "epistemically vigilant": to
4 monitor the source and content of messages, and to distinguish between
5 comprehending and accepting a message (Sperber et al., 2010). Although this
6 vigilance may be fragile (on this debate, see Gilbert, Krull, & Malone, 1990; Gilbert,
7 Tafarodi, & Malone, 1993; Hasson, Simmons, & Todorov, 2005; Mandelbaum, 2014),
8 we clearly can be epistemically vigilant. For instance, developmental studies show
9 that even 14-month-olds can distinguish between reliable and unreliable informants
10 (Poulin-Dubois, Brooker, & Polonia, 2011; Poulin-Dubois & Chow, 2009); that 3-
11 year-olds are sensitive to verbal uncertainty (Sabbagh & Baldwin, 2001), expertise
12 (Robinson, Champion, & Mitchell, 1999), and accuracy (Ganea, Koenig, & Millett,
13 2011); and that older children (between 6 and 8 years' old) are sensitive to deception
14 (Mascaro & Sperber, 2009), self-interest (Mills & Keil, 2005), and partiality (Mills &
15 Keil, 2008). Such vigilance continues into adulthood, as shown in experimental
16 studies of persuasion (for recent reviews, see Briñol & Petty, 2009; Petty & Briñol,
17 2008).

18 One form of vigilance is to monitor a source's claims and the actual outcomes:
19 whether the claims prove true or false. The recipient can use the outcomes to revise
20 their judgments about a source's reliability. When a source's claim (hypothesis, belief)
21 proves to be true, the source's reliability increases; when a source's claim proves to be
22 false, the source's reliability decreases. Following Hahn, Merdes, and von Sydow
23 (2018) we call this form of vigilance "outcome-based updating". Outcome-based
24 updating is highly effective. So long as the recipient monitors claims and eventual

1 outcomes, the recipient will converge on a true belief about the reliability of the
2 source (Hahn, Merdes, & von Sydow, 2018). But outcome-based updating does not
3 apply when outcomes are unknown or when reasoning about a singular event, such as
4 an individual piece of witness testimony in court (Hahn et al., 2018).

5 Epistemic vigilance is also embodied in a form of updating seen in putatively
6 normative models of testimony from social epistemology: the models of Bovens and
7 Hartmann (2003) and Olsson and Angere (reported in Olsson, 2011; Olsson &
8 Vallinder, 2013). These models capture belief change from a source's assertion in
9 terms of subjective probability (degrees of belief). The models both assume that
10 source reliability and information content interact. In these models, the recipient of
11 some claim (proposition, assertion, statement) assigns subjective probabilities to two
12 key variables: the reliability of the source and the information content of the claim.
13 Source reliability is understood as the (subjective) probability that a particular source
14 will say something true. The recipient has (or can form) an initial judgment of
15 reliability, prior to the specific claim being made; the recipient can, further, revise that
16 judgment in light of the claim. With respect to the information content, the recipient
17 has (or can form) some subjective prior probability. We can gloss this prior as the
18 degree of belief in the claim before receiving any evidence, whether from observation
19 or someone's assertion; or as how expected it is that the claim will prove true, again
20 before receiving any such evidence. We will use the terms "high expectedness" and
21 "low expectedness"² to indicate the degree to which the claim is expected to prove
22 true. The recipient updates the prior in light of the claim and the source's reliability.

² In this paper we deal exclusively with binary hypotheses, for which "low prior" (low expectedness) and "high prior" (high expectedness) refer to either side of the midpoint of the

1 On these models, reliable sources should tend to increase belief³; high prior
2 (high expectedness) information should tend to increase perceived reliability; and
3 low-prior (low expectedness) information should tend to decrease perceived
4 reliability. Following Hahn et al. (2018) we call this form of vigilance "expectation-
5 based updating".

6 Expectation-based updating can apply with single events, such as witness
7 testimony in a courtroom, where a single underlying historic event is at issue. This
8 updating only requires that a recipient has - or can generate - (subjective) prior
9 probabilities for the source's reliability and for the claim in hand. Since these
10 probabilities are subjective, they do not need to refer to relative frequencies of events.
11 However, expectation-based updating can have problematic effects. Through time,
12 sources may display a kind of confirmation bias, amplifying (subjectively) plausible
13 evidence and down-weighting (subjectively) implausible evidence (Hahn et al., 2018).

14 Given that people seem to perform expectation-based updating (see, e.g.,
15 Collins, Hahn, von Gerber, & Olsson, 2018), perhaps we should worry about
16 testimony: perhaps our natural, intuitively rational tendencies lead us astray. But
17 expectation-based updating has been evidenced so far with claims that are bold
18 assertions: assertions that admit no possibility of error (Collins et al., 2018). Claims
19 do not, of course, have to be so strong. Their sources can draw on language for

scale (.5). The midpoint reflects agnosticism: the hypothesis is as likely to be true as false. For multivalued hypotheses - for instance, that a person has a specific telephone number - the relevant midpoint separating low/high prior (expectedness) is $1/N$, where N is the number of rival hypotheses considered.

³ The models differ in their approach to unreliable sources: the Bovens and Hartmann model assumes that hearers, in effect, randomize on whether to believe a maximally unreliable source; the Olsson and Angere model assumes that hearers will decrease their belief in information from a maximally unreliable source.

1 making weaker, less certain claims. Perhaps, then, such language establishes
2 boundary conditions on expectation-based updating. We will consider one boundary
3 condition in the experiments reported below.

4 Expectation-based updating and outcome-based updating share a potential
5 disadvantage for speakers, in that they may undermine the basis of successful
6 communication. It is widely assumed that communication is, in some sense,
7 cooperative (Grice, 1975; McCready, 2015). This assumption can be fleshed out in
8 various ways. As Sperber et al. (2010) put it, speakers and hearers cooperate by
9 investing effort: speakers invest effort to communicate something, and hearers invest
10 effort to attend to and interpret it. This cooperative effort is sustained by the
11 expectation of benefit for each party. The benefit for the speaker is to produce some
12 desired effect; the benefit for the hearer is to acquire true, relevant information
13 (Sperber et al., 2010).

14 Fulfilling the hearer's expectations of benefit is no mean feat. In particular,
15 truth and relevance can be conflicting goals. Imagine a speaker who prioritizes
16 truthfulness. That speaker might decide only to communicate information they are
17 certain about. When that speaker has information that is relevant but uncertain, they
18 would withhold that information. If the withholding comes to light, would the speaker
19 not suffer reputational damage? There is no guarantee that speakers possess
20 information that is both relevant and certain. But we expect speakers to be
21 informative, and an uninformative speaker may damage their reputation (McCready,
22 2015). Conversely, if a speaker decides to communicate relevant but uncertain
23 information, they may also damage their reputation if that information turns out to be

1 false (McCready, 2015) or merely has a low prior probability (Collins, Hahn, von
2 Gerber, & Olsson, 2015; Collins et al., 2018).

3 Cooperation is integral to recent formal models of communication: models
4 that use the tools of game theory (e.g. Goodman & Frank, 2016; McCready, 2015).
5 One such model is the Rational Speech Act Model (RSA), which has had considerable
6 success in explaining phenomena such as scalar implicatures, hyperbole, and
7 vagueness (for discussion, see, e.g., Goodman & Frank, 2016). This model is
8 premised on cooperation. Simplifying somewhat, hearer and speaker reason about
9 each other⁴, and hearers assume that the speaker chooses an utterance proportionately
10 to its expected utility, which is determined by the "social benefit of providing
11 epistemic help to the listener" (Goodman & Frank, 2016, p. 820).

12 But what happens to cooperation when a hearer downgrades their belief in a
13 speaker's reputation? McCready (2015) points out a fundamental dilemma raised by
14 game-theoretic accounts. Cooperation requires speakers to communicate relevant
15 information, but when that information turns out to be false, it threatens the mere
16 possibility of subsequent communication. When a speaker's reputation drops enough,
17 that speaker may struggle to establish enough cooperation for communication to
18 proceed (McCready, 2015). A speaker with a bad reputation will presumably seem
19 unlikely to offer truthful, relevant information⁵ - information that motivates hearers to
20 attend and interpret (Sperber et al., 2010). So speakers must balance two, often
21 conflicting goals, cooperation and reputation management. This means that speakers

⁴ In fact, in this model the speaker reasons about an imagined literal hearer to avoid an infinite regress.

⁵ We will focus in this paper on truth, rather than relevance, and on hedging of truth. See McCready (2015) for illuminating discussion of hedging beyond truth.

1 must possess a strategy for trying to "ensure that honest mistakes don't destroy the
2 possibility of future trust and cooperation"(McCready, 2015, p.3) - and for
3 "signal[ing] that they do not wish to take responsibility for the signals they use"
4 (McCready, 2015, p.39). One such strategy, McCready notes, is the use of *hedges*:
5 evidential language ("I suspect X is the case...") that may serve as "grammatical
6 mechanisms for protecting reputations" (McCready, 2015, p. 3).

7 **Hedging**

8 The term "hedging" has a number of difference senses. Hedging is said to
9 occur when speakers qualify their assertions to lessen the impact in some way: for
10 instance, to soften bad news, to make an utterance more polite, or to avoid a firm
11 commitment to the truth or to an action (see, e.g., Holtgraves, 2002). We restrict our
12 attention, here, to hedging which qualifies assertions to allow for exceptions or
13 possible falsehood (McCready, 2015). This hedging can be likened to disclaimers in
14 advertising (McCready, 2015). Take, for instance, the following examples (the hedges
15 are italicized):

16 (a) John is *sort of* stupid

17 (b) *I suspect that* it is cold outside

18 (c) *I might be wrong, but* Palin is not going to be elected.

19 (d) *This might not be true, but* she doesn't really care about you.

20 (McCready, 2015, p. 39)

21 McCready (2015) argues that hedging with evidential language allows
22 speakers to maintain cooperation without having to be perfectly reliable sources.
23 Hedging does so by allowing sources to protect their reputations. Reputations and
24 reliability, in McCready's theory, correspond well to the notion of reliability seen in

1 the Bayesian models of testimony above. McCready takes hearers to remember their
2 interactions with particular sources; to record whether sources' claims proved true,
3 false, or indeterminate; and to judge the proportion of claims that proved true. This
4 proportion is then the reliability of the source (McCready, 2015, p. 9). Reliability, in
5 this sense, can easily be understood as the prior probability, or base rate, that a
6 particular source will say something true. When hearers interact with a new source,
7 they have no history of interactions to draw on; but they can presumably use
8 background information, such as stereotypes⁶, to form a rapid judgment of reliability.

9 A wide range of evidential language could, in principle, function as hedges,
10 but we focus on two types of evidential hedge distinguished by McCready (2015)
11 which can be seen in examples (a) to (d) above. On McCready's account, while all the
12 hedges act as disclaimers, hedges (a) and (b) work differently from (c) and (d).
13 Hedges (a) and (b), McCready argues, modify the assertions directly: they soften the
14 embedded proposition to the point at which it is assertable. Hedges (c) and (d), in
15 contrast, modify the assertions more indirectly: a bold assertion is made, but a 'shield
16 hedge', or disclaimer, is added (McCready, 2015). Hedges (c) and (d) do not, however,
17 strike all hearers as acceptable (McCready, 2015). For such hearers - the present
18 authors included - a second hedge is needed. Hence:

19 (e) *I might be wrong, but I think* that Palin is not going to be elected.

20 (f) *This might not be true, but I think* that she doesn't really care about you.

21 McCready terms these 'double hedges'. While McCready (2015) focuses on shielding,

⁶ We are not implying that stereotypes provide accurate information (though see, e.g., Jussim, Crawford, & Rubinstein, 2015, for such a claim), but they nevertheless provide information that hearers might presume is accurate and use accordingly.

1 we consider both types of evidential hedge, since both offer a way to defuse the
2 reputational hit and its potential damage to cooperation.

3 What, we might ask, is the semantic content of these hedges? We will assume
4 that these hedges express considerably less than full certainty, which we will
5 understand as a probability considerably below 1. It is a widespread assumption in
6 psychology that expressions of uncertainty can be modelled with probabilities, as
7 shown in the large literature on verbal probability expressions ("It is virtually
8 impossible/possible/likely/virtually certain")(see, e.g., Budescu & Wallsten, 1985;
9 Karelitz & Budescu, 2004; Wallsten & Budescu, 1995). Similarly, probabilities
10 feature in a recent account of epistemic modality (Lassiter, 2010, 2017). We do not
11 have conceptual grounds for fixing an exact point value for this probability. Indeed,
12 we take research on verbal probabilities to show how difficult it is to associate such
13 words with a fixed point value or range. But a loose point estimate of (somewhere
14 around) a "60% chance" seems reasonable⁷.

15 This account of hedging as reputation management is intuitively appealing and
16 compatible with existing research. Theories of argumentation, for instance, hold that it
17 is a key argumentative skill to calibrate claims to the strength of evidence or certainty
18 using appropriate evidential language (Kuhn, 1991; Toulmin, Rieke, & Janik, 1979).
19 It seems to follow that a speaker using such language should be given credit for doing
20 so and should not be treated as though they are making an absolute claim. The
21 account of hedging is also compatible with evidence that hedging can improve

⁷ Intuitively, example (a) - "sort of stupid" - appears to work differently, by adjusting the meaning of "stupid", say, to "stupid in some relevant respect". We thank an anonymous reviewer for this point.

1 perceptions of a source's credibility (Jensen, 2008; though see Longman, Turner,
2 King, & McCaffery, 2012), and that using some kinds of probabilistic language can
3 protect a speaker's reputation (Teigen, 1988). To the best of our knowledge, however,
4 no study has directly tested whether evidential hedging protects reputations.

5

6 **Overview of the Experiments**

7 In this paper we report six novel experiments that treat the following question:
8 does hedging allow speakers to make low expectedness (low prior) claims without
9 damaging their reputation? We try to answer this question for cases of pure
10 expectation and cases when the outcome is known. Throughout we operationalize a
11 speaker's reputation as their perceived reliability as measured on a ratings scale. The
12 first four experiments are variations on the same fundamental design; we report these
13 as a mini meta-analysis. These experiments explored hedging using a propositional
14 attitude ("I suspect") or a double hedge ("I might be wrong, but I think..."). We then
15 report two further experiments that extend the results of the meta-analysis.

16 As a whole, the experiments test between two contrasting sets of predictions,
17 termed, here, the *Vigilance Position* and the *Hedging Position*. The *Vigilance Position*
18 is inspired by expectation-based updating; the *Hedging Position* is inspired by
19 McCready's (2015) account of hedging:

20 *Vigilance Position*: Reputation updating is robust to hedging.

21 *Weak Form*: Low expectedness claims cause perceived reliability to
22 decrease even in the presence of hedges.

23 *Strong Form*: Hedging offers no protection; there are no differences
24 between hedged and unhedged claims.

1 *Hedging Position:* Hedging protects against reputation updating

2 *Weak Form:* Hedged low expectedness claims give rise to better
3 ratings of perceived reliability than unhedged low expectedness claims.

4 *Strong Form:* Hedging offers complete protection; hedged low
5 expectedness claims cause no decrease in perceived reliability.

6 The weak forms of the *Vigilance* and *Hedging Positions* are compatible (in as much as
7 vigilance may decrease perceived reliability, but hedging cushions that blow to
8 speaker reputation); otherwise, the positions conflict. *Strong Vigilance* implies the
9 negation of *Weak Hedging*: it is equivalent to a null hypothesis for hedging. *Strong*
10 *Hedging* implies the negation of *Weak Vigilance*: it is equivalent to a null hypothesis
11 for vigilance. In this paper, then, we focus on *Weak Vigilance* and *Weak Hedging*.

12 Later, in Experiments 5 and 6, we also test for a difference between hedging
13 with propositional attitudes and double hedging. On a plausible reading of McCready
14 (2015), hedging with propositional attitudes may offer some protection against a
15 reputational hit by weakening the assertion. But double hedging should be more
16 effective because of the disclaimer "I might be wrong, but..." Hence, double hedging
17 should lead to higher ratings of perceived reliability than propositional attitudes.

18 To preview the results, across the experiments there was no clear evidence for a
19 reliable effect of hedging. Perceived reliability was lower in the presence of low
20 expectedness (low prior) claims, whether hedged or unhedged; and there were no
21 reliable differences between hedged and unhedged claims or between hedging with
22 propositional attitudes and double hedging.

23

24 **Mini Meta-Analysis: Experiments 1 to 4**

1 We begin our test of the *Vigilance* and *Hedging Positions* with a mini meta-
2 analysis of four experiments⁸. These experiments followed the same basic design, and
3 were simple belief-change tasks based on Collins, Hahn, von Gerber, and Olsson
4 (2015, 2018), who set out to test predictions from two Bayesian models of testimony:
5 the models of Bovens and Hartmann (2003) and Olsson and Angere (reported in
6 Olsson, 2011; Olsson & Vallinder, 2013). Most relevant for present purposes is that,
7 as the models predicted, low expectedness (low prior) information decreased the
8 perceived reliability of its source (in Experiment 2 in Collins et al., 2015, which is
9 Experiment 1b in Collins et al., 2018). In other words, low expectedness information
10 damaged its source's reputation.

11 To introduce the present experiments, we consider, first, the basic method used
12 by Collins et al. (2015, 2018): a pre-test/post-design design. Participants read an
13 initial description of a source⁹, whose reliability they then rated on a scale from 0 (not
14 at all reliable) to 10 (completely reliable). For instance, participants saw the following
15 item:

16 Michael is a clinical nurse specialist.

17 *How reliable do you think he is, from 0 (not at all reliable) to 10 (completely*
18 *reliable)?*

⁸ For brevity's sake we compress some details. We note, also, that Experiments 1 to 3 also included a condition in which claims were embedded under the phrase "I am certain that" so that we could explore the effect of propositional attitudes more generally. For fuller details and data on these additional conditions, see the unpublished PhD dissertation of Collins (2017).

⁹ In other experiments, Collins et al. (2015, 2018) measured belief change instead of source reliability. Initially participants read a claim, rated their belief in it, and then saw the claim again, this time uttered by a source.

1 They then read the source making a claim, after which they rated the source's
2 reliability again on the same scale. Some participants saw a high expectedness claim,
3 that is, a claim with intuitively high prior probability:

4 Now imagine that Michael told you the following: 'One of the best remedies
5 against a severe cough is lots to drink, hot or cold.'

6 Other participants saw a low expectedness claim, that is, a claim with intuitively low
7 prior probability:

8 Now imagine that Michael told you the following: 'One of the best remedies
9 against a severe cough is valium.'

10 In each case, participants answered the following question¹⁰:

11 *Now how reliable do you think he is, from 0 (not at all reliable) to 10*
12 *(completely reliable)?*

13 Participants' prior ratings were subtracted from their posterior ratings to create change
14 scores. Positive change scores meant an increase in perceived reliability; negative
15 change scores, a decrease in perceived reliability. In an initial data set, high
16 expectedness claims lead to a statistically reliable increase in perceived reliability
17 (reliably positive change scores); low expectedness claims lead to a statistically
18 reliable decrease in perceived reliability (reliably negative change scores). A
19 replication yielded the same pattern for high expectedness claims, but the pattern was
20 different for low expectedness claims. There was a trend in the anticipated direction,
21 but this proved statistically inconclusive.

¹⁰ We note that this question is somewhat vague: reliable about what topic and on which occasion? Experiments 1 to 4 use this vague question, but Experiments 5 and 6 insert a topic to address this issue.

1 With some adaptations - and some replication - the paradigm above offers a
2 potential method for testing between the *Vigilance* and *Hedging Positions*. The same
3 basic adaptations feature in Experiments 1 to 4. The experiments used only the low
4 expectedness claims. Since, in Collins et al. (2015, 2018), participants gave these
5 claims low initial ratings of convincingness¹¹ - in other words, low prior probabilities
6 - these claims are an appropriate test case. The experiments modified these claims
7 with (a) hedge(s), and used only reliable sources to avoid any floor effects. The
8 experiments still followed a pre-test/post-test format, with participants rating sources'
9 reliability before and after seeing the source make the low expectedness claim. The
10 key manipulation was hedging: whether the claim was unhedged - a bold assertion -
11 or hedged. This manipulation was between-participants.

12 Experiments 1 to 4 had a number of aims, including exploring the effects of
13 propositional attitudes more generally. For present purposes we focus on hedging with
14 the weak propositional attitude 'suspect' (see Methods for justification of that choice),
15 as in Experiments 1 to 3, and the double hedge 'I might be wrong, but I think...'), as in
16 Experiment 4.

17 Now that the design is clear, we can address some assumptions we made to
18 generate predictions for Experiments 1 to 4, which were necessary because there is an
19 imperfect fit between expectation-based updating and McCready's (2015) hedging.
20 Firstly, McCready takes reliability to be induced from a history of interactions. In this

¹¹ The relevant data are from Experiments 1a and 2a in Collins et al. (2018). Pooling the data for both experiments (as is legitimate in a Bayesian analysis) for the claims we include in the present paper, a Bayesian one-sample *t*-test in the BEST package (Meredith & Kruschke, 2013) showed a mean estimate rating of 3.62, 95 % HDI[3.23, 4.01]. Hence the claims were rated comfortably below the mid-point of the scale (0 to 10).

1 and all other experiments, we instead used background knowledge or context to
2 suggest sources' reliability, based on previous research (e.g. Collins et al., 2018). We
3 assumed that this is a reasonable simplification. Secondly, McCready focuses on the
4 case where "hedges shield the speaker from blame [reputational damage] if it turns
5 out that her assertion fails to represent the facts correctly"(McCready, 2015, p. 3). In
6 Experiments 1 to 4 (and, later, in Experiment 5) we, instead, focussed on
7 expectations: participants never learnt the truth of the matter. We took expectations to
8 set the bar lower. However different outcomes and expectations might be from a
9 rational point of view (see the distinction about accuracy in the General Introduction),
10 from the recipient's point of view there seems to be a continuum between things they
11 know to be false and things they expect to be false. If a recipient finds a claim
12 implausible, they presumably assign it a low probability; if a recipient learns that
13 claim is false, they presumably assign it zero probability¹². It is unclear why hedging
14 should work in the more extreme, zero-probability case, and not the low-probability
15 case. The less extreme case seems, to us, an easier target. We will return to this point,
16 though, in Experiment 6 and the General Discussion.

17 These assumptions having been made, the predictions were as follows. This was
18 a single-factor between-participants design, with the variable Hedging with levels
19 'Unhedged' and 'Hedged-Suspect' (Experiments 1 to 3) or 'Hedged-Double' (i.e. 'I
20 might be wrong, but I think...'; Experiment 4). The pre-test/post-test structure allowed
21 various comparisons: one-sample comparisons (versus zero) for each condition, and
22 comparisons between hedged and unhedged conditions. We predicted the following.

¹² Or near zero: a recipient might want to acknowledge some small chance of error.

1 *Basic Reputational Hit*: in the unhedged condition, perceived reliability should
2 reliably decrease (one-sample comparison).

3 *Weak Vigilance*: perceived reliability should reliably decrease in the unhedged,
4 hedged-suspect and hedged-double conditions (one-sample comparison).

5 *Weak Hedging*: in Experiments 1 to 4 change scores should differ reliably
6 between the unhedged and hedged-suspect and the unhedged and hedged-
7 double conditions, with the unhedged condition reliably lower than both of the
8 hedged conditions, indicating a bigger decrease in perceived reliability
9 (between-groups comparison)

10

11 **Methods**

12 *Participants*. Table 1 summarizes the number of participants for each
13 experiment. The numbers reported are for the experiment as a whole, including the
14 conditions not reported in this paper. Sample sizes were selected to give high power
15 (.8) for a medium effect size in frequentist analyses (the original analyses were
16 frequentist). Numbers are reported for the original samples and for the samples
17 retained after exclusions. For all experiments, participants were excluded if they
18 stated that their first language was not English; for Experiments 2 to 4, they were also
19 excluded if they failed a memory check or indicated that they did not participate
20 seriously (see Procedure). Gender and mean age are reported for the participants
21 whose data were retained.

22

23

24

1

2 *Table 1.* Participants for Experiments 1 to 4

Experiment	N(Sample)	N(Retained)	N(Female)	Mean Age
1	160	159	84	37.35
2	170	165	96	39.93
3	161	151	75	37.69
4	128	124	71 (and 1 non- binary)	38.36

3

4

5 **Materials.** There were five items, which took the following form:6 *Initial source information:* Michael is a clinical nurse specialist.7 *Claim:* Now imagine that Michael told you the following: ‘One of the best
8 remedies for a severe cough is valium.’9 The item above corresponds to the unhedged condition. In the hedged condition,
10 participants read the claim embedded under "I suspect that..."11 We adopted the items from Collins et al. (2018), selecting reliable sources
12 making low expectedness claims. A clinical nurse specialist claimed that valium was
13 one of the best remedies for a severe cough. A baker claimed that the varying
14 temperatures of a particular oven made it perfect for fluffy and crispy bread. A
15 journalist with a good track record predicted that a horse would beat a competitor
16 despite losing to that horse in the majority of races. A retired meteorologist claimed
17 that the maximum summer temperature in Stockholm in June 2013 was 15 degrees.

1 And a respected DJ claimed that a club in Detroit had the reputation of one of coolest
2 in the world.

3 The hedged conditions embedded the claims underneath the propositional
4 attitude 'suspect' (hedged-suspect; Experiments 1 to 3) or the double hedge 'I might be
5 wrong, but I think...' (hedged-double; Experiment 4). Both types of hedge feature in
6 McCready's analysis. In the present study, Experiments 3 and 4 also included
7 manipulation checks in which participants clearly found hedged claims weaker than
8 their unhedged counterparts: in Experiment 3 "suspect" was rated on average 4.92
9 ($SD = 2.04$) on a scale from 0 (very weak) to 10 (very strong), whereas the unhedged
10 claims were rated on average 7.70 ($SD = 2.43$); in Experiment 4, "suspect" was rated
11 on average 5.27 ($SD = 1.91$) and "I might be wrong, but..." was rated on average 4.96
12 ($SD = 1.93$), whereas the unhedged claims were rated on average 8.02 ($SD = 2.39$).
13 The hedged (hedged-suspect and hedged-double) claims were rated as conveying
14 significantly weaker claims than unhedged claims ($ps < .001$; for details, see Collins,
15 2017).

16 **Procedure.** The experiments reported here and below were carried out
17 following the guidelines, and with the approval, of the ethics committee at the
18 Department of Psychological Sciences, Birkbeck, University of London. The
19 experiments were posted on Amazon Mechanical Turk via the intermediary MTurk
20 Data (www.mturkdata.com). We set high qualifications for the task: participants had to
21 be resident in the US and have an overall approval rating of 99% for 1,000 previously
22 completed tasks. These criteria maximized the number of first-language English
23 speakers. Each participant was only able to take part in a single experiment. On the
24 consent page, participants were told that the task would assess how people judge

1 information they receive from other people. No information was given about the
2 specific manipulation. After giving informed consent, participants were assigned, in a
3 round-robin manner, to a condition. They were shown the following instructions:

4 Thank you for taking part in this study. You'll be shown some descriptions of
5 people, and will be asked to indicate how reliable these people seem, by
6 selecting the appropriate number on a scale from 0 (not at all reliable) to 10
7 (completely reliable).

8 Participants then saw 5 items, comprising initial source information and then a claim.
9 Participants gave a prior rating after the initial source information, and a posterior
10 rating after the claim. The order of presentation was counterbalanced. Finally,
11 participants were given debriefing information, and were paid \$0.75, a fee chosen to
12 exceed the US minimum wage for reasonable completion times.

13 In Experiments 2 to 4, there were the following changes to the procedure. To
14 ensure that participants processed the materials deeply enough, participants were
15 informed that they may have to perform a memory test. Participants in the hedged
16 conditions did, indeed, perform a memory test at the end of the experiment and in a
17 fixed order. The test comprised a recall test and a recognition test, in that order, details
18 of which can be found at osf.io/r9cna. The recall test gave an indication of how
19 deeply participants were processing the materials: for instance, in Experiment 2,
20 approximately 76% passed the test, rising to 84% if near synonyms are included. But
21 the recognition test was the minimum threshold for inclusion of the data. This
22 threshold was waived if participants passed the recall test but failed the recognition
23 test, since these participants may have misinterpreted the recognition question as a
24 cue that their recall was incorrect. Then, participants answered a sincerity question,

1 based on Aust, Diedenhofen, Ullrich, and Musch (2013); for details see osf.io/r9cna.
2 Experiments 2 and 3 also placed the propositional attitude 'suspect' in capitals as a
3 simple attentional device; this device was dropped in Experiment 4. And Experiments
4 3 and 4 included a manipulation check between the recognition test and the sincerity
5 question; for the wordings, see osf.io/r9cna.

6

7 **Results**

8 Copies of the data files and analysis scripts are available on the Open Science
9 Framework at osf.io/r9cna. For all data, we calculated change scores by subtracting
10 the prior rating from the posterior rating. We then averaged change scores across
11 items to produce a single score per participant¹³. For the present meta-analysis, we
12 analysed the data in two ways: we first calculated meta-analytic means and
13 confidence intervals in R (R Core Team, 2016) using the 'meta' package (Schwarzer,
14 2007). Since we are only interested in summarizing the present studies, we used a
15 fixed-effects model. We then entered the relevant *t*-statistics from the experiments into
16 a meta-analytic Bayesian *t*-test in the 'Bayes Factor' package (Morey & Rouder,
17 2015). Bayes Factors allow us to directly assess the support for the null and
18 alternative hypotheses. We assess each of the predictions in turn.


19 ***Basic Reputational Hit.*** Fig. 1 is a forest plot of the mean change scores and
20 confidence intervals for the unhedged condition across the four experiments (studies).
21 The figure shows reliably negative change scores: participants reliably decreased their

¹³ Given this averaging, we must be cautious about generalizing beyond the existing items, but see Experiment 5.

1 belief in the reliability of the source when s/he made a low expectedness claim.

2 Hence, the data replicate the basic reputational hit.

3



4

5 *Figure 1.* Forest plot of mean change scores and confidence intervals for the unhedged
6 condition of Experiments 1 to 4.

7 A meta-analytic Bayesian t-test confirmed this result. Table 2 reports the one-

8 sample *t*-statistics, number of participants, and degrees of freedom for the unhedged

9 condition in each study.

10 *Table 2.* One-sample *t*-statistics, degrees of freedom, and number of participants for the
11 unhedged condition for Experiments 1 to 4.

Study	<i>t</i> -statistic (df)	<i>N</i>
1	-6.23 (53)	54
2	-6.53 (58)	59
3	-5.83 (51)	52
4	-4.43 (57)	58

12


13 The Bayesian analysis found strong support for a reputational hit from low

14 expectedness claims. To gauge sensitivity to priors, the analysis used the three

15 different pre-set Cauchy priors in the Bayes Factor package for the standardized effect

1 size r : medium, wide, and ultra-wide. These each yielded strong support for a
 2 reputational hit, $BFs_{Alternative} > 100$.

3 **Weak Vigilance.** Fig. 2 is a forest plot of the mean change scores and
 4 confidence intervals for the hedged condition across the four experiments (hedged-
 5 suspect in Experiment (Study) 1 to 3; hedged-double in Experiment (Study) 4). The
 6 figure shows reliably negative change scores: participants decreased their belief in the
 7 reliability of the source when s/he made a low expectedness claim even when that
 8 claim was hedged.



9

10 *Figure 2.* Forest plot of mean change scores and confidence intervals for the hedged condition
 11 for Experiments 1 to 4.

12

13 A meta-analytic Bayesian t -test confirmed this result. Table 3 reports the t -
 14 statistics, number of participants, and degrees of freedom for the hedged condition.

15

16

17 *Table 3.* One-sample t -statistics, degrees of freedom, and number of participants for the
 18 hedged condition for Experiments 1 to 4.


Study	t -statistic (df)	N
1	-8.99 (53)	54

2	-9.66 (51)	52
3	-6.64 (46)	47
4	-7.89 (66)	67

1

2 The Bayesian analysis found strong support for a reputational hit with hedged
 3 claims, hence for *Weak Vigilance*. As above, the analysis used the three pre-set
 4 Cauchy priors, which each yielded strong support for *Weak Vigilance*, $BF_{S_{Alternative}} >$
 5 100.

6 **Weak Hedging.** Fig. 3 is a forest plot summarizing the difference in change
 7 scores between the unhedged and hedged conditions across the four experiments
 8 (studies). For each study, the confidence intervals include 0; the meta-analytic
 9 confidence interval also includes 0, and the meta-analytic mean is slightly negative,
 10 with non-significantly lower ratings in the hedged conditions. The figure offers no
 11 support for the predicted effect of hedging.



12

13 *Figure 3.* Forest plot of differences in mean change score between hedged and unhedged
 14 conditions (and confidence intervals) for Experiments 1 to 4.

15 A meta-analytic Bayesian t-test confirmed this result. Table 4 reports the
 16 independent-samples *t*-statistics, number of participants in hedged and unhedged
 17 groups, and the degrees of freedom across the four experiments.

18

1

2 *Table 4.* Independent-samples *t*-statistics, degrees of freedom, and participant numbers for
 3 difference in mean change scores for hedged and unhedged conditions for Experiments 1 to 4.

Study	<i>t</i> -statistic (df)	<i>N</i> (Hedged)	<i>N</i> (Unhedged)
1	1.26 (104.27)	54	54
2	-.33 (96.41)	52	59
3	-.14 (95.88)	47	52
4	1.02 (106.73)	67	58

4

5 The Bayesian analysis found moderate support for the null hypothesis, hence
 6 moderate evidence against *Weak Hedging*. As above, the Bayesian analysis used the
 7 three pre-set Cauchy priors for the effect size. In each case there was moderate
 8 support for the null hypothesis, but there was some sensitivity to those priors: with the
 9 default medium prior, $BF_{Null} = 6.23$; with the wide prior, $BF_{Null} = 8.68$; and with the
 10 ultra-wide prior, $BF_{Null} = 12.18$.

11

12 Discussion

13 The mini meta-analysis is informative on a number of points. Firstly, it
 14 replicated the basic reputational hit from low expectedness claims which was found in
 15 the initial study by Collins et al. (2015, 2018) but which was inconclusive in a
 16 replication (Collins et al., 2018). This finding is consistent with Bayesian models of
 17 testimony (Bovens & Hartmann, 2003; Olsson, 2011; Olsson & Vallinder, 2013) and
 18 with the idea of epistemic vigilance (Sperber et al., 2010). The finding generalizes,
 19 moreover, to low expectedness claims which are embedded under hedges. The data

1 support at least *Weak Vigilance*: low expectedness claims cause a reputational hit even
2 when hedged. The data do not support *Weak Hedging*: hedging did not reliably protect
3 reputations. Indeed, the data suggest moderate evidence for the relevant null
4 hypothesis. In other words, epistemic vigilance seems robust against hedging. The
5 data do not sit well with the idea that a key function of hedges is the protection of
6 reputation, in the sense of perceived reliability (McCready, 2015).

7 The results above arise from experiments with slightly different procedures.
8 Experiments 2 to 4 sought to ensure that the hedges were not simply ignored.
9 Experiments 2 and 3 placed the hedge, 'suspect', in capitals; Experiments 2 to 4 all
10 included a memory test, which encouraged participants to process the text more
11 deeply; and Experiments 3 and 4 also included manipulation checks. These procedural
12 differences make no obvious difference to the results, and suggest that participants in
13 all experiments processed the items in some depth.

14 Nevertheless, Experiments 1 to 4 have a number of limitations. The
15 experiments used a design and statistical analysis that limit generalizability. While we
16 might expect our results to generalize to new participants, we have no grounds to
17 expect them to generalize to new materials (Clark, 1973). The experiments also used a
18 pre-test/post-test format, a method that may give rise to demand characteristics
19 (Levine & Parkinson, 1994). In particular, participants may have deduced that we
20 intended them to use the content of the claims to revise their judgments of source
21 reliability. Since the hedging manipulation was between-participants, participants had
22 no evidence to guess the manipulation itself. But the results are worth replicating with
23 an alternative, more natural design. Finally, the experiments did not include any
24 definition of 'reliability' for participants; we assumed, instead, that participants would

1 use an intuitive notion. This intuitive notion may not, however, map neatly onto the
2 notion of reliability that is theoretically relevant.

3

4 **Experiment 5**

5 Experiment 5 set out to address the limitations of the preceding studies. It used
6 a design that yields better evidence of generalizability, addresses potential demand
7 characteristics of the pre-test/post-test format, and tightens up the dependent measure.

8 The basic design was as follows. The experiment included 20 items in total,
9 including adapted versions of the materials from Experiments 1 to 4. To reduce task
10 demands and vary demand characteristics, Experiment 5 abandoned the pre-test/post-
11 test format, and adopted a fully between-participants design which comprised the
12 following 4 conditions.

13 *Null Condition:* participants read only information about 20 sources.

14 *Unhedged Condition:* participants read these 20 sources making low
15 expectedness claims (one claim per source).

16 *Propositional-Attitude Condition:* participants read these 20 sources making
17 the same low expectedness claims but hedged with a propositional attitude ('I
18 suspect that...').

19 *Double-Hedged Condition:* participants read these 20 sources making the
20 same low expectedness claims but hedged with 'I might be wrong, but I
21 think...'

22 Participants rated source reliability, this time using the (potentially) more
23 sensitive measure of fixing a slider on a scale from 0 to 100 (on the definition, see
24 below).

1 In this design, the null condition provides a baseline for the reliability of a
2 source. It anchors the reputational hit: a reputational hit occurs in any condition where
3 ratings are reliably lower than the null condition. The basic effect of low expectedness
4 claims should be seen in the difference between null and unhedged conditions. The
5 effects of hedging can be seen in the difference between unhedged and hedged
6 (Propositional-Attitude, Double-Hedged) conditions. If hedging protects the source's
7 reliability, then ratings should be reliably higher in hedged conditions than in the
8 unhedged condition. This design also permits direct comparison between hedging
9 with a propositional attitude and double hedging.

10 The new design has certain advantages over the preceding studies but prompts
11 further changes. First, consider some advantages. Since the design is fully between-
12 participants, it reduces (or at least varies) the potential demand characteristics. In the
13 between-participants design, participants give a single rating after reading the items,
14 and only a control group (the Null Condition) rates the source when s/he has not made
15 a claim. The design also includes more items, allowing us to be more confident about
16 generalizing our findings.

17 The fully between-participants design might, however, alter how people
18 understand the task. Accordingly we introduced new features to protect construct
19 validity. These features largely concern the dependent variable, perceived reliability.
20 Since different conditions present different amounts of information, participants in
21 these conditions might interpret reliability differently. In particular, participants in the
22 null condition do not see a claim at all, which risks them giving more general ratings
23 than participants in the other conditions. It is one thing to rate a nurse's reliability as a
24 source in general and another to rate the nurse's reliability as a source about medical

1 matters. To aid interpretation, we made the following four changes across all
2 conditions. Firstly, we defined reliability as a person's "credibility, trustworthiness, or
3 expertise - as whether they tend to say the truth". Including this definition also
4 mitigates the concern that participants might understand reliability differently from
5 the theoretically important sense. Secondly, we added information to each item to
6 provide a context for the rating. So, for example, the new version of the now-familiar
7 Michael item was as follows, with the relevant information italicized:

8 *You are discussing treatments for coughs and colds, which are mostly*
9 *treated with simple over-the-counter medicines. Michael is a nurse*
10 *practitioner who specializes in minor illnesses.*

11 Thirdly, we fixed the scope of the reliability question for each item by identifying a
12 topic. For the Michael item (topic italicized):

13 How reliable is Michael *about remedies for minor illnesses?*

14 Since there is no objective measure of topic scale - Is the evolution of the domestic cat
15 a bigger/smaller topic than the history of the Napoleonic wars? - we relied on intuition
16 to keep the scale comparable. In each case, the topic was broad enough to avoid
17 collapsing to a simple judgment of belief in a claim in the conditions where a claim
18 was presented (unhedged, propositional attitudes, and double hedged). Fourthly, we
19 gave participants three practice items to sensitize them to the full range of the scale.
20 These items described a high reliability (highly trustworthy and expert) source, a
21 middle-reliability source (with a 50:50 hit rate), and a low reliability (highly inexperienced)
22 source.

23 A final innovation concerns expectedness. In Experiments 1 to 4, the items
24 varied in how expectedness was determined. In some items, expectedness was

1 determined by (presumed) background beliefs: for instance, we presumed that
2 participants would have a sense of likely and unlikely summer temperatures. In other
3 items, expectedness was determined by contextual information: for instance, we told
4 participants about the record of a pair of horses in races against each other. In
5 Experiment 5, when items invoked background knowledge, they also explicitly stated
6 that knowledge: hence, for example, an item about summer temperatures in North
7 Africa explicitly stated that summers there are very hot.

8 For this design we make the following predictions:

9 *Reputational Hit*: the Unhedged Condition should have reliability ratings
10 reliably lower than the Null condition.

11 *Weak Vigilance*: reliability ratings should be reliably lower than the Null
12 Condition in all other conditions.

13 *Weak Hedging*: the hedged conditions (Propositional Attitude and Double-
14 Hedged Conditions) should have reliability ratings reliably higher than the
15 Unhedged Condition.

16 *Shielding*: the Double-Hedged condition should be rated reliably higher
17 than the Propositional Attitude condition.

18

19 **Methods**

20 *Participants*. 206 participants completed a web survey; the same selection
21 criteria were used as for the previous experiments. We retained the data for the 203
22 participants (88 female; average age 35.98) who indicated that English was their first
23 language.

1 **Materials.** There were 3 practice items (see Procedure) and 20 test items.
2 These can be seen in a copy of the Qualtrics survey on the Open Science Framework
3 at osf.io/r9cna. Each test item described a scenario, gave information about a source,
4 and set up an expectation. The Null condition comprised only this information. In the
5 Unhedged, Propositional-Attitude and Double-Hedged conditions the source also
6 made a low expectedness claim (given the context). In the hedged conditions, the
7 claim was modified by a propositional attitude ('I suspect that...') or a double hedge ('I
8 might be wrong, but I think that...'). Finally, there was a reliability question: 'How
9 reliable is [Source] about [Topic]?' The following is one of the items. In the Null
10 condition, participants saw only the context and the reliability question.

11 You are talking about Susie, a high-school student, and her work habits. Susie
12 is known for being easily distracted. Jake is Susie's brother, and they both still
13 live at home.

14 *Unhedged:* He tells you, 'Susie does her homework best sitting in front of the
15 TV.

16 *Propositional Attitude:* He tells you 'I suspect that Susie does her homework
17 best sitting in front of the TV.

18 *Double-Hedged:* He tells you, 'I might be wrong, but I think that Susie does
19 her homework based sitting in front of the TV.'

20 How reliable is Jake about Susie?

21 **Procedure.** The experiment comprised a Qualtrics web-survey, posted on
22 Mechanical Turk by the intermediary MTurkData. Participants gave informed consent
23 and were (pseudo-)randomly assigned to a condition. All participants received the
24 following instructions:

1 In this task, we would like you to rate how reliable you think certain people
2 are as sources of information. You can think of someone's reliability as their
3 credibility, trustworthiness, or expertise - as whether they tend to say the
4 truth.

5 You'll see 20 scenarios. For each one you'll be asked about someone's
6 reliability. We would like you to make your judgments using the sliding scale
7 provided. '0' means not at all reliable, '100' completely reliable. You now have
8 the chance to practice on the following items.

9 They then rated the following items using a sliding scale from 0 to 100, to
10 sensitize themselves to the scale.

11 [High Reliability] You are being told about a particular type of deep-
12 sea fish by a world expert on the subject. The expert has no reason to
13 misinform you. How reliable is the expert about this fish?

14 [Medium Reliability] You are being told by a high-school student
15 about what they learnt in class today. The student has a vivid
16 imagination, and makes things up about half the time. How reliable is
17 the student about their class?

18 [Low Reliability] You are visiting a foreign city. Another tourist is telling you
19 about the city. The tourist has only briefly looked at a guidebook in a foreign
20 language that they understand very poorly. How reliable is the tourist about
21 the city?

22 Participants were then reminded of the instructions using the following text:

23 Thank you. Let's get on with the rest of the experiment. Now you'll see 20
24 items, each about a different topic. Just as with the practice questions, you'll

1 be asked about the reliability of some people as sources of information about a
2 topic. As you've just practiced, please rate their reliability using the scale
3 provided. '0' means not at all reliable. '100' means completely reliable.


4 Each participant then rated 20 items. The web survey randomized the order of
5 presentation. Finally, participants received debriefing information, and received a fee
6 of \$1.50, which was chosen to exceed the US minimum wage per minute.

7 **Results**

8 Copies of the data files and analysis scripts can be found on the Open Science
9 Framework at osf.io/r9cna. The analysis comprised a Bayesian zero-one inflated beta
10 regression run in the 'brms' package (Bürkner, 2017). Conclusions, here, are based on
11 the estimated marginal means and the 95% Highest Posterior Density intervals,
12 calculated with the emmeans package (Lenth, 2018). A survey error led to ambiguities
13 in two items (Items 11 ("Eric") and 12 ("Alex") in the survey; see the Open Science
14 project) that may have decreased the effect of hedging. Accordingly we excluded
15 these items, and ran all analyses on the remaining 18 items.

16 Fig. 4 shows the descriptive data and was produced using the 'pirateplot'
17 function of the 'yarr package' (Phillips, 2017). The descriptive data suggest a clear
18 drop in perceived source reliability between the Null and remaining conditions. There
19 are small differences between the remaining conditions with the Propositional

- 1 Attitude (PropAtt) condition rated slightly worse than the Unhedged and Double-
- 2 Hedged (Double) condition.



3

4 *Figure 4.* The figure shows the horizontally jittered raw data and smoothed density
 5 distributions by condition. Black bars represent the mean for each condition; the white boxes
 6 represent the 95% high density interval (see Kruschke, 2013).

7 We analyzed the data with a zero-one-inflated beta regression with perceived
 8 reliability (rescaled to the range 0-1) predicted by Condition (Null, Unhedged,
 9 Propositional Attitude, Double-Hedging). The model included random slopes and
 10 intercepts for items and random intercepts for participants: that is, in R syntax, the
 11 model included the terms '(Condition|Item)' and '(1|ID)', where 'ID' stands for
 12 'participant'.

13 We used the default uninformative, or weakly informative, priors from the
 14 brms package. For the beta coefficients, the prior for the intercept was a 0-centered

1 Student- t distribution with 3 degrees of freedom and scale of 31; the prior for the
 2 remaining coefficients was an improper uniform prior across the real numbers. For the
 3 standard deviations of the random effects, the prior was a 0-centered Student- t
 4 distribution with 3 degrees of freedom and scale of 31); for the correlation between
 5 random effects, the prior, ('lkj(1)'), set all correlation matrices to be equally likely. For
 6 the conditional-one inflation (coi) and zero-one inflation probabilities (zoi), the priors
 7 were uniform (Beta($\alpha = 1$, $\beta = 1$)). For phi, the prior was a gamma distribution with
 8 shape and spread of .01. We ran 3 chains for 10,000 iterations. These settings
 9 achieved good convergence for all chains. All Rhat values were 1.00; and effective
 10 sample sizes were above 3,800.

11 Table 5 reports the parameter estimates.

12 *Table 5.* Parameter estimates (excluding group-level parameters) and 95% Credible Intervals.
 13 Note that estimates are on the logit scale, and that the independent variable was treatment-
 14 coded.

Parameter	Parameter Estimate	95% Credible Interval
Null (Intercept)	1.05	.82, 1.27
Unhedged	-1.03	-1.31, -.75
Propositional Attitude	-1.10	-1.38, -.83
Double	-1.09	-1.36, -.83
phi	4.40	4.20, 4.61
zoi	.06	0.05, 0.07
coi	.65	.58, .71

15

16 Table 6 reports the estimated marginal means, which are on the response scale.

1 *Table 6.* Estimated marginal means and 95% HPDIs. Note that all estimates are on the
 2 original response scale.

Parameter	Estimated marginal mean	95% HPDI
Null (Intercept)	74.13	69.62, 78.25
Unhedged	50.55	43.38, 57.89
Propositional Attitude	48.87	42.07, 55.78
Double	49.05	42.75, 55.60

3

4 The estimated marginal means show that the experimental (non-null)
 5 conditions are clearly reliably lower than the Null Condition, but are themselves close
 6 together with substantial overlap in the 95% highest posterior density intervals. We
 7 consider the evidence for each hypothesis in turn.

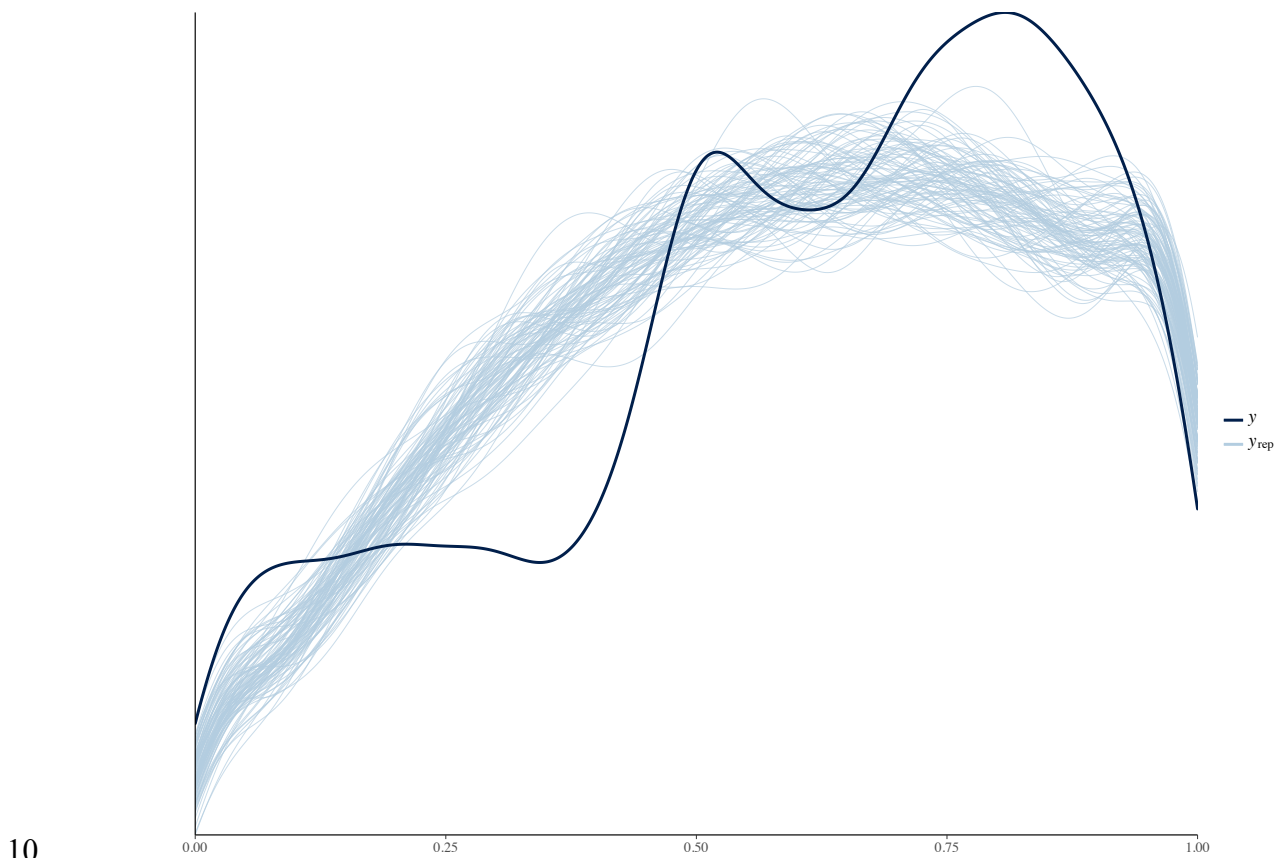
8 ***Reputational Hit.*** The data replicate the basic reputational hit seen in
 9 Experiments 1 to 4: ratings are reliably lower in the Unhedged Condition than in the
 10 Null Condition.

11 ***Weak Vigilance.*** The data also replicate the finding that ratings are reliably
 12 lower in both the Propositional Attitude condition and the Double Hedged condition
 13 than in the Null Condition.

14 ***Weak Hedging.*** The data show no clear evidence for an effect of hedging: the
 15 hedged conditions do not have ratings reliably higher than the Unhedged Condition:
 16 there is almost complete overlap in the HPD intervals, and the mean ratings were
 17 slightly lower for hedged conditions than for the Unhedged Condition.

1 **Shielding.** The data show no evidence of shielding: the Double-Hedged
2 condition did not have reliability ratings reliably higher than the Propositional
3 Attitude Condition. Indeed, there was almost complete overlap of 95% HPD intervals,
4 for the two types of hedging: the Double-Hedged Condition extended only .44 points
5 on the reliability scale above the Propositional Attitude Condition.

6 We tested model fit with posterior predictive checks, shown in Fig. 5. These
7 checks suggest adequate fit, though note the mismatch between approximately .25
8 and .50 (25 and 50 on the response scale): the model predicts too many ratings in this
9 interval.



10

11 *Figure 5.* Posterior predictive checks for the model, showing the data (y) against 100
12 predictions from the model (yrep).

1

2 In summary, the data for Experiment 5 only offer support for *Weak Vigilance*.
3 They suggest, further, that ratings were highly similar for the Unhedged, Propositional
4 Attitude, and Double Hedged Conditions.

5 **Discussion**

6 A fifth experiment has now replicated the reputational hit from low
7 expectedness claims and shown that it can occur even when the claims are hedged.
8 This experiment used new items, a new design, and a more tightly defined dependent
9 variable. Nevertheless, the data replicate the basic pattern seen in Experiments 1 to 4:
10 support for *Weak Vigilance* and no support for *Weak Hedging*. Experiment 5 adds a
11 direct comparison between hedging with the propositional attitude 'I suspect that' and
12 double-hedging with 'I might be wrong, but I think...'. There was no reliable difference
13 between these types of hedging.

14 At this point, however, we must reconsider our auxiliary assumptions and, in
15 particular, the assumption that expectations are a fair test case for reputation
16 management. The assumption seems reasonable: low expectedness claims seem to
17 provide a clear case of reputation damage and should, on McCready's (2015) account,
18 threaten cooperation. As we have already discussed, expectations seem to set the bar
19 somewhat lower. And many conversations will be about things whose truth we cannot
20 - at least not immediately - establish. Our approach, then, seems to treat an important
21 part of the picture.

22 Nevertheless, McCready does not treat expectations, but cases where the
23 outcome is known: to repeat an earlier quotation, she focuses on the case where
24 "hedges shield the speaker from blame [reputational damage] if it turns out that her

1 assertion fails to represent the facts correctly"(McCready, 2015, p. 3). It might be a
2 fairer test of McCready's account to include outcomes. Outcomes feature in our final
3 experiment, to which we now turn.

4 **Experiment 6**

5 This experiment used the same design as Experiment 5 with only a single
6 alteration. We added, to each item, the phrase 'Later you learn that what [source] told
7 you was wrong'. We repeat the predictions from above:

8 *Reputational Hit*: the Unhedged Condition should have reliability ratings
9 reliably lower than the Null condition.

10 *Weak Vigilance*: reliability ratings should be reliably lower than the Null
11 Condition in all other conditions.

12 *Weak Hedging*: the hedged conditions (Propositional Attitude and Double-
13 Hedged Conditions) should have reliability ratings reliably higher than the
14 Unhedged Condition.

15 *Shielding*: the Double-Hedged condition should be rated reliably higher than
16 the Propositional Attitude condition.

17

18 **Methods**

19 *Participants*. 205 participants completed a web survey; the same selection
20 criteria were used as in all previous experiments. We retained the data for the 204
21 participants (86 female; average age = 36.86) who listed English as their first
22 language.

23 *Materials*. The materials were the same as for Experiment 5 (the errors in
24 items 11 and 12 having been corrected). For each item, an extra line was added: 'Later


1 you learn that what [source] told you was wrong.' A copy of the Qualtrics survey is
2 available on the Open Science Framework at osf.io/r9cna.

3 **Procedure.** The procedure was identical to Experiment 5.

4 **Results**

5 Copies of the data files and analysis scripts are available on the Open Science
6 Framework at osf.io/r9cna. As with Experiment 5, the analysis comprised Bayesian
7 zero-one inflated beta regression run in the 'brms' package, and conclusions are based
8 on the estimated marginal means and 95% HPDs.

9 Fig. 6 shows the descriptive data.



10

11 *Figure 6.* The figure shows the horizontally jittered raw data and smoothed density
12 distributions by condition. Black bars represent the mean for each condition; the white boxes
13 represent the 95% high density interval (see Kruschke, 2013).

1

2 The data suggest that, when low expectedness information proves false, there
3 is a large hit to reliability. The pattern differs from that seen in Experiment 5: the
4 Double-Hedging Condition has a noticeably higher mean than the Unhedged and
5 Propositional Attitude Conditions.

6 We analyzed the data with a zero-one-inflated beta regression with perceived
7 reliability (rescaled to the range 0-1) predicted by Condition (Null, Unhedged,
8 Propositional Attitude, Double-Hedging). The model included random slopes and
9 intercepts for items and random intercepts for participants: that is, in R syntax, the
10 model included the terms '(Condition|Item)' and '(1|ID)', where 'ID' stands for
11 'participant'. We used the same priors as for Experiment 5, and once again ran 3
12 chains for 10,000 iterations, achieving good convergence. All Rhat values were 1.00;
13 effective sample sizes were all above 1,100.

14 Table 7 reports the parameter estimates.

15 *Table 7.* Parameter estimates (excluding group-level parameters). Note that estimates are on
16 the logit scale, and that the independent variable was treatment-coded.

Parameter	Parameter Estimate	95% Credible Interval
Null (Intercept)	1.22	.95, 1.48
Unhedged	-2.31	-2.65, -1.97
Propositional Attitude	-2.49	-2.82, -2.15
Double	-2.06	-2.38, -1.74
phi	7.78	7.42, 8.14
zoi	.11	.10, .12

coi .10 .07, .13

1

2 Table 8 reports the estimated marginal means, which are on the response scale.

3

4

5

6 *Table 8.* Estimated marginal means and 95% HPDIs. Note that all estimates are on the
7 original response scale.

Parameter	Estimated marginal mean	95% HPD
Null (Intercept)	77.15	71.20, 81.24
Unhedged	25.13	20.73, 30.53
Propositional Attitude	21.82	17.92, 26.56
Double	29.97	25.01, 35.70

8 The estimated marginal means show that the experimental (non-null)

9 conditions are clearly reliably lower than the Null Condition. A new finding, here, is

10 that the Double Hedging Condition has higher ratings than either the Unhedged or

11 Propositional Attitude Conditions. However, there is overlap in the 95% HPD

12 intervals, if rather minimal between the hedged conditions. Unsurprisingly, ratings in

13 the non-null conditions are considerably, and reliably, lower in this experiment than in

14 Experiment 5.

15 ***Reputational Hit.*** The data replicate the basic reputational hit seen in

16 Experiments 1 to 4: ratings are reliably lower in the Unhedged Condition than in the

17 Null Condition.


1 **Weak Vigilance.** The data also replicate the finding that ratings are reliably
2 lower in both the Propositional Attitude condition and the Double Hedged condition.

3 **Weak Hedging.** The data do not show clear evidence of hedging. The
4 Propositional-Attitude Condition once again has ratings lower than the Unhedged
5 Condition, though these ratings are, once again, not reliably different. The Double
6 Hedging Condition, this time, has higher ratings than the Unhedged Condition, but the
7 95% HPD intervals overlap by some 4.28 reliability points.

8 **Shielding.** The data hint at an effect of shielding. The Double-Hedged
9 Condition had higher ratings than the Propositional Attitude Condition, but the 95%
10 HPD intervals overlap by some 1.55 reliability points. There is, therefore, insufficient
11 evidence to conclude in favour of an effect of shielding.

12 We tested model fit with posterior predictive checks, shown in Fig. 7. The
13 checks suggest reasonable fit, though some over-prediction in the .75 to 1 range (75%

1 to 100% on the response scale) and some under prediction in approximately the .15
 2 to .40 (15 to 40) range.



4 *Figure 7.* Posterior predictive checks, showing the data (y) against 100 simulations from the
 5 model (y_{rep}).
 6

7 In summary, the headline finding is that there is, once again, no clear evidence
 8 that hedging protects reliability. This headline should not, however, obscure the
 9 differences between the data sets. While Experiment 5 offered good evidence for the
 10 similarity of ratings in the Unhedged, Propositional Attitude, and Double Hedged
 11 Conditions, Experiment 6 offers tentative evidence for differences: in particular,
 12 between the Unhedged and Double Hedged Conditions - where the overlap in 95%

1 HPDIs is now much less - and between the Propositional Attitude and Double Hedged
2 Conditions, where there is only small overlap.

3 We explored whether the above differences were statistically reliable by
4 compiling the data sets for Experiments 5 and 6 and running a Bayesian mixed-effects
5 zero-one-inflated beta regression. For this analysis, we excluded items 11 and 12 from
6 Experiment 6 to match the exclusions from Experiment 5. The model included terms
7 for Condition (Null, Unhedged, Propositional Attitude, and Double) and Experiment
8 (Expectation, Outcome; where 'Expectation' stands for Experiment 5 and 'Outcome'
9 for Experiment 6). We allowed the interaction and main effects to vary across items
10 (i.e., in R syntax (Condition*Experiment|Item) and included random intercepts for
11 items and participants. As above, we used the default priors in the brms package. We
12 ran 3 chains for 10,000 iterations. These settings resulted in good convergence. All
13 Rhat values were 1.00, and effective sample sizes were above 3,100.

14 Table 9 reports the parameter estimates.

15

16

17

18

19

20

21 *Table 9.* Parameter estimates (excluding the group-level parameters). Note that estimates are
22 on the logit scale, and that both independent variables were treatment-coded.

Parameter	Parameter	95% Credible Interval
	Estimate	

Null (Intercept)	1.08	.89, 1.27
Unhedged	-.98	-1.22, -.75
Propositional Attitude	-1.17	-1.40, -.95
Double	-1.12	-1.36, -.89
Experiment (Outcome)	.08	-.03, .18
Unhedged: Outcome	-.99	-1.18, -.82
PropAtt: Outcome	-1.10	-1.27, -.93
Double: Outcome	-.73	-.88, -.57
phi	4.30	4.16, 4.45
zoi	.08	.08, .09
coi	.30	.27, .34

1

2

Crucial, here, are the interaction terms (Unhedged: Outcome, PropAtt:

3

Outcome, and Double: Outcome). As we now expect, there is a reputational hit in all

4

the experimental (non-null) conditions. This effect is reliably larger when outcomes

5

are known than when there are merely expectations. But, notably, the interaction term

6

'Double: Outcome' is the smallest, and is reliably smaller than the term 'PropAtt:

7

Outcome', as shown by non-overlapping 95% credible intervals. In other words, there

8

is less of an increase in reputational hit with Double Hedging than with Propositional

9

Attitude Hedging. These data support the conclusion that there is a different pattern

10

among the experimental (non-null) conditions in Experiment 6 than in Experiment 5.


11

We tested model fit with posterior predictive checks, as shown in Fig. 8, which

12

reveal over-prediction of ratings in approximately the .15 to .50 range (15 to 50 on the

1 response scale) and under-prediction of ratings in approximately the .65 to .90 range
 2 (65 to 90 on the response scale).



3
 4 *Figure 8.* Posterior predictive checks for the model, showing the data (y) against 100
 5 predictions from the model (y_{rep}).

6

7 **Discussion**

8 Experiment 6 corroborates the findings of the previous experiments. Once
 9 outcomes are known, there is still good evidence for the reputational hit and for *Weak*
 10 *Vigilance*: unsurprisingly, there is a considerably larger reputational hit when
 11 outcomes are known. Moreover, telling participants the outcome does not reveal
 12 evidence to endorse hedging: there is still considerable overlap between Unhedged
 13 and Hedged conditions, though with somewhat reduced overlap between Unhedged

1 and Double Hedged Conditions relative to Experiment 5. There is still, then, no clear
2 evidence for a protective effect of hedging. There is, however, some evidence for an
3 effect of shielding, that is, for a difference between hedging with propositional
4 attitudes and hedging with double hedges. We can see this in the small degree of
5 overlap in the respective 95% HPD intervals in Experiment 6. There also appears to
6 be a somewhat different pattern of results in Experiments 5 and 6. When outcomes are
7 known, there is always an increase in reputational hit over pure expectations, but this
8 increase is smallest in the Double Hedging Condition and is reliably smaller than in
9 the Propositional Attitude Condition. We take this to be weak evidence of a difference
10 between hedging with propositional attitudes and double hedging.

11 In summary, while introducing outcomes does affect ratings, it does not
12 fundamentally alter our conclusions about the evidence: there remains good evidence
13 for vigilance and no clear evidence for a protective effect of hedging on perceived
14 reliability.

15 **General Discussion**

16 In six experiments, we have explored whether speakers can protect their
17 reputations by hedging claims with evidential language. We have tested hedging
18 against expectation-based and outcome-based updating. We have found no clear
19 evidence for a protective effect of hedging in these experiments, and strong evidence
20 of a reputational hit even when claims are hedged. Our data are unpromising for the
21 view that evidential hedging serves to protect speakers' reputations, whether with a
22 propositional attitude or a double hedge. Successful hedging would have been a way
23 to escape both the confirmation bias seen with expectation-based updating and the

1 dilemma for game-theoretic accounts of pragmatics. But we find no evidence for
2 successful hedging.

3 Although our data do not look promising for reputation management with
4 evidential hedges, we do not claim to have falsified it. We have made a number of
5 assumptions to test for such hedging (on auxiliary assumptions and falsification, see,
6 e.g., Earp & Trafimow, 2015; Lakatos, 1978; Meehl, 1990; Trafimow, 2009, 2012).
7 These assumptions have produced a distinctive set of tasks in which we failed to find
8 successful hedging. But varying these assumptions might produce more promising
9 results; we see this paper as part of a broader debate. To motivate future work, we
10 discuss our assumptions, attendant limitations, and future directions. We then consider
11 the possible implications of our data for the literature.

12

13

14

15 **Expectations**

16 There are a number of points to be made about expectations. Firstly, one key
17 assumption was that expectations are a fair test case for reputation management. This
18 assumption is clearest in Experiments 1 to 5, where the materials only manipulated
19 expectations and never specified the outcomes. But the assumption was carried over
20 into Experiment 6: to keep the designs and results as comparable as possible, we fixed
21 the claims in Experiment 6 as low expectedness, and specified the outcome - that the
22 claims proved false.

23 We have justified the use of expectations on various grounds. We have linked
24 expectations and known outcomes: we see known outcomes as an extreme case of

1 expectations, and see no reason to sharply distinguish the two from the recipient's
2 perspective (though see the Introduction for a distinction in accuracy). Moreover,
3 expectations are the more general case, as they are implicated in a great many
4 conversations about things whose truth we cannot (immediately) establish. Finally,
5 expectation-based updating seems to pose as much of a challenge to cooperation as
6 outcome-based updating.

7 But we must acknowledge a limitation with respect to testing McCready's
8 theory (2015). Since McCready's (2015) theory focuses on cases when outcomes are
9 known, more data are needed from experiments that specify outcomes. We found no
10 clear evidence of hedging when recipients learnt that a low expectedness claim proved
11 false. But hedging might prove more effective when claims have medium or high
12 prior probability and yet prove false. For example, if a sports journalist predicts that a
13 horse with an excellent track record will win when it, in fact, loses, then that journalist
14 may suffer reputational damage. But if the journalist hedges the prediction, they may
15 plausibly evade the reputational damage. Future experiments could usefully vary
16 expectations to cover more of the probability scale. We question, though, whether
17 claims that prove false will actually lead to reputational damage at all, hedged or not,
18 when the claim is highly plausible.

19 Varying expectations brings us to a second key point. Our data suggest that,
20 when speakers make low expectedness claims, they undermine their reputation; and
21 by "low expectedness" we have meant low prior probability. But if speakers are to be
22 highly informative, they must communicate true information that is not just unknown
23 to the hearer but also surprising to them: that has a low prior probability. It might,
24 then, be impossible to be informative without the risk of reputational damage.

1 What we have not tested, however, is the extent to which that reputational
2 damage scales with the level of probability in question. Our methods could be adapted
3 to vary the expectedness (prior probabilities) alongside knowledge of the outcomes, to
4 explore the full range of the probability scale and identify where, and at what level,
5 reputational hits occur. We suspect that a key issue will be whether speakers provide
6 any supporting information. If speakers make surprising claims without supporting
7 evidence - as characters in our experimental materials did - they may suffer a
8 reputational hit. But perhaps a hit is defused if speakers rely on more complex
9 utterances to make low expectedness claims: claims together with supporting
10 evidence, or with some acknowledgement that the claim is surprising. At the same
11 time, it will be important to examine in future work how reputation recovers when
12 low probability claims ultimately turn out to be true, in particular to what extent they
13 return reputation to a higher level than before.

14 **Hedging**

15 Across all of the experiments, we operationalized hedging in just two ways:
16 with the propositional attitude "suspect" and the double hedge "I might be wrong, but
17 I think...." We have assumed that these hedges fairly represent their respective
18 categories. There is certainly evidence that people consider these hedges weaker than
19 unhedged claims: Collins (2017) reports manipulation checks conducted with
20 Experiments 3 and 4 that show that both hedges were rated as expressing significantly
21 weaker claims than unhedged claims. But there is, of course, a wide range of
22 alternative hedges available. For instance, there are numerous other propositional
23 attitudes of different strengths, such as 'suppose', 'feel', 'intuit', 'reckon', and many
24 more.

1 It is possible that other hedges might be more effective. But since exploring a
2 wide range hedges comes with a risk of false-positive results (see, e.g., Simmons,
3 Nelson, & Simonsohn, 2011; Wicherts et al., 2016), any such exploration should be
4 underpinned by theory to explain why hedges differ. We suggest that future work
5 should draw more deeply on semantics than we have done in this paper. We have
6 taken a similar path to research on verbal probability expressions, and assumed that
7 the probability (certainty) expression maps onto the probability scale, without
8 articulating a clear semantic (or pragmatic) theory (for discussion, see Collins &
9 Hahn, 2018). A clear theory would be invaluable in predicting which hedges should
10 work and in which contexts. We will suggest some possible ingredients of successful
11 hedging later in this section.

12

13 **Reliability**

14 Another set of assumptions concerned reliability. Firstly, we implicitly
15 assumed that hedging should work across the range of source reliability, and fixed our
16 sources as reliable across all the experiments to avoid floor effects. This assumption is
17 open to question. There are, in fact, data suggesting that people expect high-expertise
18 (i.e. reliable) sources to be speak with authority or certainty (Longman et al., 2012).
19 Thus, if high-expertise sources communicate uncertainty through hedging, they may
20 actually damage their own reputation rather than protecting it (though see Karmarkar
21 & Tormala, 2010, for counterintuitive effects on persuasion). If hedging made
22 reputations worse in our tasks, we would expect to see hedged conditions rated
23 reliably lower than unhedged conditions, a pattern we did not see. It could
24 nevertheless be informative to include lower-reliability sources.

1 While our design could have detected a quantitative effect like this, it could
2 not have detected a qualitative shift. We have assumed that people interpret reliability
3 in the same way across conditions: that "reliability" means the same when there are
4 unhedged claims as when there are hedged ones. In Experiments 5 and 6, we included
5 a definition to try to achieve this consistency: a person's "credibility, trustworthiness,
6 or expertise - as whether they tend to say the truth". We consider this definition
7 faithful to the sense of reliability relevant for McCready (2015) and for Bayesian
8 models of testimony. But it does allow some flexibility in interpretation.

9 It is conceivable that participants focused on different aspects of reliability in
10 the different conditions. For instance, without a hedge, participants may have focused
11 strictly on the source's accuracy; but with a hedge, they may have focused on
12 expertise in a different sense, the source's perceived authoritativeness. After all, the
13 source who hedges may be accurately reporting their mental state, irrespective of the
14 true state of affairs in the world: whether or not valium is a good treatment for severe
15 coughs, the speaker *suspects* that it is. A reputational hit could, then, mean different
16 things in different conditions, with ratings reflecting different aspects of reliability. To
17 detect such qualitative shifts, future experiments could unpack reliability into
18 different scales (as is often done in the persuasion literature; see, e.g., McCroskey &
19 Teven, 1999). More theoretical work would also be needed to identify which aspect of
20 reliability is most important to cooperation.

21 A second assumption is embodied in the way we conveyed the sources'
22 expertise. We decided against giving participants a sequence of interactions with a
23 particular source from which they could induce the source's reliability. We instead
24 used background knowledge and context to suggest the source's reliability, which

1 minimized demands on memory and inference. This simplification seems reasonable,
2 as does the assumption that evidential hedging should still apply, but future work
3 could give participants direct experience with a source over multiple trials. Plausibly,
4 the way information is conveyed - by description or experience - could affect how
5 firmly participants fix their initial belief about the source's reliability, hence how
6 prepared they are to shift it based on the low expectedness claim.¹⁴

7 Finally, we assumed that reliability could be measured on ratings scales. Such
8 scales are a simple, direct way of operationalizing reliability. These, and similar,
9 scales are widely used in research on persuasion (for reviews of resulting findings, see
10 Briñol & Petty, 2009; Petty & Briñol, 2008), argumentation (e.g. Bhatia & Oaksford,
11 2015; Corner & Hahn, 2009; Eemeren, Garssen, & Meuffels, 2009; Hahn &
12 Oaksford, 2007), and testimony (e.g. Collins et al., 2018; Harris, Hahn, Madsen, &
13 Hsu, 2016). Other researchers might prefer an alternative dependent measure. For
14 instance, a character's reliability might be inferred from a participant's willingness to
15 cooperate with them in a cooperation game.

16

17

18

19 **Implications**

20 These assumptions suggest the need for a larger research program. But let us
21 assume, for now, that our results stand with the alternative methods and measures
22 mentioned above. What would the implications be? Hedging of some kind is

¹⁴ We thank an anonymous reviewer for this suggestion.

1 intuitively appealing. We still find it plausible that speakers can somehow protect their
2 reputations with hedging. One possibility is that hedging with propositional attitudes,
3 shield hedges, or double hedging works when there is a clear contrast. In our
4 experiments, even when participants were prompted to attend to the hedges, they had
5 nothing to compare the hedges to. But successful hedging may, in fact, rely on explicit
6 contrasts. Take, for instance, the following exchange:

7 Bob: Looks like it'll rain tomorrow.

8 Sandy: Are you sure?

9 Bob: Well, I suspect it will.

10 In this context, Bob may be understood as meaning that he is not sure: that he only
11 has sufficient confidence to use 'suspect' (Horn, 1989; Levinson, 2000; Van Der
12 Auwera, 1996; Verstraete, 2005). Hedging might work with such clear contrasts. But
13 McCready's (2015) account is so appealing in large part because hedging provides a
14 way to sustain cooperation. If hedging only works through such pragmatic inferences,
15 it presumably plays less of a role in justifying and maintaining cooperation, since
16 pragmatics is typically taken to be the result of cooperation, not a precondition for it¹⁵.
17 Nevertheless, one future direction is to develop materials such as the example
18 dialogue above, where a demand for certainty is met with a hedged claim.

19 Another potential factor can be seen in a parallel with verbal probability
20 expressions: modal adjectives and adverbs, such as 'impossible', 'possibly', 'likely'
21 and 'certain'. These terms feature in standardized lexicons for expressing risk, for

¹⁵ We take McCready (2015) to hold a more nuanced view of her proposed mechanism, one in which the mechanism sits at the boundary between semantics and pragmatics.

1 instance, the risks of climate change (e.g. Intergovernmental Panel on Climate
2 Change, 2005). One feature of verbal probability expressions is that, independently of
3 the probability they convey, some suggest that the described event will occur and
4 some that it will not. That is, they have directionality. For example, although
5 participants give ‘some possibility’ and ‘quite uncertain’ similar numerical
6 interpretations, significantly more participants prefer an operation described as having
7 ‘some possibility’ of success to one whose success is ‘quite uncertain’ (Teigen &
8 Brun, 1999; for more general discussion of scales of alternatives, see Geurts, 2013).

9 Thanks to directionality, verbal probability expressions can have subtle effects
10 on reputation. Teigen (1988) presented participants with predictions couched in verbal
11 probability expressions with different directionality: one expert predicted a rise in
12 crude oil prices by saying "It is possible that oil prices will reach \$20 in October",
13 another by saying "It is not quite certain that oil prices will reach \$20 in October".
14 Participants were then told that oil reached \$20. Even though participants thought the
15 first expert had a lower probability in mind than the second, they decided that the first
16 was more right (for discussion, see Teigen & Brun, 1999).

17 The question naturally arises of whether directionality is essential to hedging.
18 Certainly, the present studies used a propositional attitude with positive directionality.
19 One test for directionality is to think of (or ask participants for) continuations: so, for
20 example, to complete the sentence "It's possible that X because..." we would add a
21 reason for occurrence; but to complete the sentence "I'm not completely certain that X
22 because...." we would add a reason for non-occurrence. ‘Suspect’ suggests
23 occurrence: "I suspect that X" naturally invites continuations that justify X being the
24 case. The double hedge is more complicated: the sentence "I might be wrong,

1 because..." invites reasons for wrongness, but the word 'but' seems to override this
2 directionality. Future work could manipulate directionality in hedging with
3 propositional attitudes. However, if such hedging is to play a role in justifying
4 cooperation, an account would also need to show that directionality is a semantic,
5 rather than pragmatic, phenomenon. Geurts (2013) offers hope here, outlining a
6 semantic account of scales of alternatives for various types of expression, including
7 probability statements.

8 Our data also relate to research on plausible deniability and indirect speech
9 (Lee & Pinker, 2010; Pinker, Nowak, & Lee, 2008). According to this research, a
10 strategic speaker can choose to make indirect speech acts when it is unclear whether a
11 context involves cooperation or conflict. In particular, a strategic speaker can enable a
12 cooperative hearer to act favourably and prevent an uncooperative hearer from acting
13 antagonistically. Imagine, for example, a motorist stopped by a police officer and
14 given a ticket. The motorist wishes to avoid a ticket by offering a bribe. A bribe would
15 be accepted by a dishonest police officer, but would lead to arrest by an honest police
16 officer. One option is to offer a bribe through an indirect speech act, such as "Gee,
17 officer, is there some way we could take care of the ticket here?" (Lee & Pinker, 2010;
18 Pinker, Nowak, & Lee, 2008). This utterance would allow a dishonest officer to
19 recognize and accept the bribe, but prevents an honest officer from having clear
20 evidence for an arrest. There is experimental evidence that indirectness is sensitive to
21 such pay-off structures (Lee & Pinker, 2010).

22 Strategic indirectness might seem like a case of successful hedging, and hence
23 might seem to conflict with our results. There is, however, no true conflict. Strategic
24 indirectness may well allow plausible deniability, but its effects do not depend on

1 reputation management. Take the case of the motorist above. While the motorist may
2 well have done enough to avoid being arrested, they are likely to suffer reputational
3 damage. Indeed, such manoeuvring can be transparent and can itself cause
4 reputational damage without undermining its strategic objective.

5 As we have seen, evidential language might yet allow reputation management
6 but through a pragmatic mechanism: direct comparison with a set of alternative
7 expressions, giving rise to a scalar implicature. But what function might the evidential
8 language fulfil here if not reputation management? The evidential language is
9 presumably not meaningless. We can again draw a comparison with verbal probability
10 expressions. These expressions have uses beyond conveying probabilities (for
11 discussion, see Collins & Hahn, 2018). Bonnefon and Villejoubert (2006) showed that
12 verbal probability expressions can be used to convey bad news tactfully. They had
13 participants read the sentence, "The doctor tells you, you will possibly suffer from
14 insomnia [deafness] soon" and elicited membership functions from participants.¹⁶
15 Participants understood 'possibly' as indicating higher probabilities with deafness
16 than insomnia. Moreover, some 60% of participants indicated that the doctor was
17 being tactful not uncertain. Bonnefon and Villejoubert argue that different
18 interpretations can result from participants perceiving an utterance as a face-
19 threatening act: that is, as an act that threatens someone's desire for autonomy (their
20 negative face) or their desire for connection with others (their positive face) (Brown

¹⁶ In membership-function studies, participants see a number line representing the probability range [0,1]. The line picks out numbers at even intervals across the range, different studies selecting different intervals. Participants are asked to rate how well each number corresponds to a verbal probability expression. These data allow a membership function to be calculated which can be interpreted as representing the meaning of the expression.

1 & Levinson, 1987; Holtgraves, 2002). On Bonnefon and Villejoubert's account,
2 verbal probability expressions can serve as face-management devices: they allow a
3 speaker to acknowledge, and lessen the impact of, a face-threatening act.

4 This politeness account can be extended to the present data. Participants may
5 have understood at least some of the claims as advice: as advice, say, to take a
6 medicine, buy an oven, bet on a horse, and so on. Participants might take this advice
7 as an imposition or threat to their freedom of action: in the terms of Politeness Theory
8 (Brown & Levinson, 1987), as a threat to negative face. Advice is not far removed
9 from requests, offers, compliments, and so on, which are typically taken to threaten
10 negative face (for discussion, see Holtgraves, 2002). Future work could probe this
11 possibility by asking participants why the speaker used the hedges, after Bonnefon
12 and Villejoubert (2006). Although a polite source could seem cooperative, helpful,
13 and so on – hence, in some sense, reliable - they may nevertheless not appear reliable
14 in the sense of being a source of true information. It could be instructive, therefore, to
15 manipulate the social setting and, with it, the plausibility of face-management
16 strategies.

17 As we have seen, much remains to be explored about reputation management.
18 This work should be part of a broader program of experimental work on the
19 relationship between trust (perceived reliability), cooperation, and pragmatics,
20 addressing the question of to what extent pragmatics can proceed without trust. Much
21 as McCready (2015) seems, to us, to have identified a genuine threat to pragmatics,
22 there are clearly contexts in which pragmatics can proceed with incomplete trust or
23 cooperation. Pragmatics does not stop in arguments or in adversarial situations such
24 as court rooms (Asher & Lascarides, 2013; Goodwin, 2001). And even when a

1 speaker is, in a sense, as unreliable as possible - is a compulsive liar - their utterances
2 may be informative. In some contexts, the hearer could reverse the content of the
3 message or decrease their belief in what the compulsive liar says. More data are
4 needed on when, and precisely how, distrust undermines pragmatics. Part of the
5 picture will be speakers' attempts to manage their reputations.

6

7 **Conclusions**

8 Our data suggest that reputation damage persists despite hedging with
9 propositional attitudes and double hedges. This damage occurs both with expectations
10 and known outcomes. The data also cast doubt on the view that evidential hedging
11 serves a strategy for speakers to manage their reputations and protect their reliability
12 as information sources and conversation partners. These data do not support a way to
13 defuse the damaging consequences of expectation-based or outcome-based updating.
14 We have suggested a number of avenues for future work. It remains unclear, though,
15 whether the options we suggest could sustain cooperation, and whether they can offer
16 an escape from the dilemma that faces game-theoretic accounts of pragmatics.

17

18 **Acknowledgements**

19 Peter Collins was supported by ESRC grant (ES/J500021/1) from the
20 Bloomsbury Doctoral Training Centre, held at Birkbeck, University of London. This
21 grant also funded Experiments 1 to 4. He was also supported by a postdoctoral
22 fellowship at the Munich Center for Mathematical Philosophy, funded by an
23 Alexander von Humboldt Professorship to Professor Stephan Hartmann. Ulrike Hahn
24 was supported by an Annaliese Maier Research Award from the Alexander von

1 Humboldt Foundation. This grant also funded Experiments 5 and 6. The funders had
2 no involvement in study design; in collection, analysis and interpretation of data; in
3 writing the report; or in the decision to submit the article for publication.

4

5

1 **References**

2 Adler, J. (2015). *Epistemological problems of testimony* (Summer 2015; E. N.
3 Zalta, Ed.). Retrieved from
4 <https://plato.stanford.edu/archives/sum2015/entries/testimony-episprob/>

5 Asher, N., & Lascarides, A. (2013). Strategic conversation. *Semantics and*
6 *Pragmatics*, 6(0), 2-1–62. <https://doi.org/10.3765/sp.6.2>

7 Aust, F., Diederhofen, B., Ullrich, S., & Musch, J. (2013). Seriousness checks
8 are useful to improve data validity in online research. *Behavior Research Methods*,
9 45(2), 527–535. <https://doi.org/10.3758/s13428-012-0265-2>

10 Bhatia, J.-S., & Oaksford, M. (2015). Discounting testimony with the
11 argument ad hominem and a Bayesian congruent prior model. *Journal of*
12 *Experimental Psychology: Learning, Memory, and Cognition*, 41(5), 1548–1559.
13 <https://doi.org/10.1037/xlm0000151>

14 Bonnefon, J.-F., & Villejoubert, G. (2006). Tactful or Doubtful? Expectations
15 of Politeness Explain the Severity Bias in the Interpretation of Probability Phrases.
16 *Psychological Science*, 17(9), 747–751. [https://doi.org/10.1111/j.1467-](https://doi.org/10.1111/j.1467-9280.2006.01776.x)
17 [9280.2006.01776.x](https://doi.org/10.1111/j.1467-9280.2006.01776.x)

18 Bovens, L., & Hartmann, S. (2003). *Bayesian Epistemology* [OUP Catalogue].
19 Retrieved from Oxford University Press website:
20 <https://ideas.repec.org/b/oxp/obooks/9780199270408.html>

21 Briñol, P., & Petty, R. E. (2009). Source factors in persuasion: A self-
22 validation approach. *European Review of Social Psychology*, 20(1), 49–96.
23 <https://doi.org/10.1080/10463280802643640>

- 1 Brown, P., & Levinson, S. C. (1987). *Politeness: Some Universals in*
2 *Language Usage*. Cambridge, England: Cambridge University Press.
- 3 Budescu, D. V., & Wallsten, T. S. (1985). Consistency in interpretation of
4 probabilistic phrases. *Organizational Behavior and Human Decision Processes*,
5 36(3), 391–405. [https://doi.org/10.1016/0749-5978\(85\)90007-X](https://doi.org/10.1016/0749-5978(85)90007-X)
- 6 Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models
7 Using Stan. *Journal of Statistical Software*, 80(1).
8 <https://doi.org/10.18637/jss.v080.i01>
- 9 Chance, Z., Norton, M. I., Gino, F., & Ariely, D. (2011). Temporal view of the
10 costs and benefits of self-deception. *Proceedings of the National Academy of*
11 *Sciences*, 108(Supplement 3), 15655–15659.
12 <https://doi.org/10.1073/pnas.1010658108>
- 13 Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of
14 language statistics in psychological research. *Journal of Verbal Learning and Verbal*
15 *Behavior*, 12(4), 335–359. [https://doi.org/10.1016/S0022-5371\(73\)80014-3](https://doi.org/10.1016/S0022-5371(73)80014-3)
- 16 Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford, England:
17 Oxford University Press.
- 18 Collins, P. J. (2017). *Rationality, pragmatics, and sources*. (Unpublished
19 doctoral dissertation, Birkbeck, University of London). Retrieved from
20 <http://bbktheses.da.ulcc.ac.uk/284/>
- 21 Collins, P. J., & Hahn, U. (2018). Communicating and reasoning with verbal
22 probability expressions. In *Psychology of Learning and Motivation* (Vol. 69, pp. 67–
23 105).

1 Collins, P. J., Hahn, U., von Gerber, Y., & Olsson, E. J. (2015). The bi-
2 directional relationship between source characteristics and message content. In D. C.
3 Noelle, R. Dale, S. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P.
4 Maglio (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science*
5 *Society* (pp. 423–428). Austin, TX: Cognitive Science Society.

6 Collins, P. J., Hahn, U., von Gerber, Y., & Olsson, E. J. (2018). The Bi-
7 directional Relationship between Source Characteristics and Message Content.
8 *Frontiers in Psychology, 9*. <https://doi.org/10.3389/fpsyg.2018.00018>

9 Corner, A., & Hahn, U. (2009). Evaluating science arguments: Evidence,
10 uncertainty, and argument strength. *Journal of Experimental Psychology. Applied,*
11 *15(3)*, 199–212. <https://doi.org/10.1037/a0016533>

12 Earp, B. D., & Trafimow, D. (2015). Replication, falsification, and the crisis of
13 confidence in social psychology. *Frontiers in Psychology, 6*.
14 <https://doi.org/10.3389/fpsyg.2015.00621>

15 Eemeren, F. H. van, Garssen, B., & Meuffels, B. (2009). *Fallacies and*
16 *Judgments of Reasonableness: Empirical Research Concerning the Pragma-*
17 *Dialectical Discussion Rules*. Dordrecht: Springer.

18 Ganea, P. A., Koenig, M. A., & Millett, K. G. (2011). Changing your mind
19 about things unseen: Toddlers' sensitivity to prior reliability. *Journal of Experimental*
20 *Child Psychology, 109(4)*, 445–453. <https://doi.org/10.1016/j.jecp.2011.02.011>

21 Geurts, B. (2013). Alternatives in Framing and Decision Making. *Mind &*
22 *Language, 28(1)*, 1–19. <https://doi.org/10.1111/mila.12005>

23 Gilbert, D. T., Krull, D. S., & Malone, P. S. (1990). Unbelieving the
24 unbelievable: Some problems in the rejection of false information. *Journal of*

1 *Personality and Social Psychology*, 59(4), 601–613. [https://doi.org/10.1037/0022-](https://doi.org/10.1037/0022-3514.59.4.601)
2 3514.59.4.601

3 Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe
4 everything you read. *Journal of Personality and Social Psychology*, 65(2), 221–233.
5 <https://doi.org/10.1037/0022-3514.65.2.221>

6 Goodman, N. D., & Frank, M. C. (2016). Pragmatic Language Interpretation
7 as Probabilistic Inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
8 <https://doi.org/10.1016/j.tics.2016.08.005>

9 Goodwin, J. (2001). The non-cooperative pragmatics of arguing. In E. Nemeth
10 (Ed.), *Pragmatics in 2000: Selected papers from the 7th International Pragmatics*
11 *Conference*. (Vol. 2, pp. 263–277). Antwerp, Belgium: International Pragmatics
12 Association.

13 Grice, H. P. (1975). Logic and Conversation. In Davidson, D. & (second)
14 Harman, G. (Eds.), *The logic of grammar* (pp. 64–75). Encino, CA: Dickenson.

15 Hahn, U., Merdes, C., & von Sydow, M. (2018). How good is your evidence
16 and how would you know? *Topics in Cognitive Science*, 10(4), 660–678.

17 Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation:
18 A Bayesian approach to reasoning fallacies. *Psychological Review*, 114(3), 704–732.
19 <https://doi.org/10.1037/0033-295X.114.3.704>

20 Hahn, U., Oaksford, M., & Harris, A. J. L. (2012). Testimony and Argument:
21 A Bayespian Perspective. In F. Zenker (Ed.), *Bayesian Argumentation* (pp. 15–38).
22 New York, NY: Springer.

1 Harris, A. J. L., Hahn, U., Madsen, J. K., & Hsu, A. S. (2016). The Appeal to
2 Expert Opinion: Quantitative Support for a Bayesian Network Approach. *Cognitive*
3 *Science*, 1496–1533. <https://doi.org/10.1111/cogs.12276>

4 Hasson, U., Simmons, J. P., & Todorov, A. (2005). Believe It or Not On the
5 Possibility of Suspending Belief. *Psychological Science*, 16(7), 566–571.
6 <https://doi.org/10.1111/j.0956-7976.2005.01576.x>

7 Holtgraves, T. M. (2002). *Language As Social Action: Social Psychology and*
8 *Language Use*. Mahwah, New Jersey: Lawrence Erlbaum Associates.

9 Horn, L. R. (1989). *A natural history of negation*. Chicago, IL: University of
10 Chicago Press.

11 Intergovernmental Panel on Climate Change. (2005). *Guidance notes for lead*
12 *authors of the IPCC Fourth Assessment Report on Addressing Uncertainties*.
13 Retrieved from [http://www.ipcc-wg2.awi.de/guidancepaper/uncertainty-guidance-](http://www.ipcc-wg2.awi.de/guidancepaper/uncertainty-guidance-note.pdf)
14 [note.pdf](http://www.ipcc-wg2.awi.de/guidancepaper/uncertainty-guidance-note.pdf)

15 Jensen, J. D. (2008). Scientific Uncertainty in News Coverage of Cancer
16 Research: Effects of Hedging on Scientists' and Journalists' Credibility. *Human*
17 *Communication Research*, 34(3), 347–369. [https://doi.org/10.1111/j.1468-](https://doi.org/10.1111/j.1468-2958.2008.00324.x)
18 [2958.2008.00324.x](https://doi.org/10.1111/j.1468-2958.2008.00324.x)

19 Jussim, L., Crawford, J. T., & Rubinstein, R. S. (2015). Stereotype
20 (In)Accuracy in Perceptions of Groups and Individuals. *Current Directions in*
21 *Psychological Science*, 24(6), 490–497. <https://doi.org/10.1177/0963721415605257>

22 Karelitz, T. M., & Budescu, D. V. (2004). You Say “Probable” and I Say
23 “Likely”: Improving Interpersonal Communication With Verbal Probability Phrases.

1 *Journal of Experimental Psychology: Applied*, 10(1), 25–41.

2 <https://doi.org/10.1037/1076-898X.10.1.25>

3 Karmarkar, U. R., & Tormala, Z. L. (2010). Believe Me, I Have No Idea What
4 I'm Talking About: The Effects of Source Certainty on Consumer Involvement and
5 Persuasion. *Journal of Consumer Research*, 36(6), 1033–1049.

6 <https://doi.org/10.1086/648381>

7 Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of*
8 *Experimental Psychology: General*, 142(2), 573–603.

9 <https://doi.org/10.1037/a0029146>

10 Kuhn, D. (1991). *The Skills of Argument*. Cambridge University Press.

11 Lakatos, I. (1978). *The methodology of scientific research programmes*.
12 Cambridge, England: Cambridge University Press.

13 Lassiter, D. (2010). Gradable epistemic modals, probability, and scale
14 structure. *Semantics and Linguistic Theory*, 20, 197.

15 <https://doi.org/10.3765/salt.v20i0.2557>

16 Lassiter, D. (2017). *Graded Modality: Qualitative and Quantitative*
17 *Perspectives*. Oxford University Press.

18 Lee, J. J., & Pinker, S. (2010). Rationales for indirect speech: The theory of
19 the strategic speaker. *Psychological Review*, 117(3), 785–807.

20 <https://doi.org/10.1037/a0019688>

21 Lenth, R. (2018). emmeans: Estimated Marginal Means, aka Least-Squares
22 Means. (Version R packages version 1.2.3). Retrieved from [https://CRAN.R-](https://CRAN.R-project.org/package=emmeans)
23 [project.org/package=emmeans](https://CRAN.R-project.org/package=emmeans)

- 1 Levine, G., & Parkinson, S. (1994). *Experimental Methods in Psychology*.
2 Mahwah, New Jersey: Lawrence Erlbaum Associates.
- 3 Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized*
4 *conversational implicature*. Cambridge, MA: MIT Press.
- 5 Loftus, E. F. (1975). Leading questions and the eyewitness report. *Cognitive*
6 *Psychology*, 7(4), 560–572. [https://doi.org/10.1016/0010-0285\(75\)90023-7](https://doi.org/10.1016/0010-0285(75)90023-7)
- 7 Longman, T., Turner, R. M., King, M., & McCaffery, K. J. (2012). The effects
8 of communicating uncertainty in quantitative health risk estimates. *Patient Education*
9 *and Counseling*, 89(2), 252–259. <https://doi.org/10.1016/j.pec.2012.07.010>
- 10 Mandelbaum, E. (2014). Thinking is Believing. *Inquiry*, 57(1), 55–96.
11 <https://doi.org/10.1080/0020174X.2014.858417>
- 12 Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading
13 components of children’s vigilance towards deception. *Cognition*, 112(3), 367–380.
14 <https://doi.org/10.1016/j.cognition.2009.05.012>
- 15 Mazar, N., Amir, O., & Ariely, D. (2008). The Dishonesty of Honest People: A
16 Theory of Self-Concept Maintenance. *Journal of Marketing Research*, 45(6), 633–
17 644. <https://doi.org/10.1509/jmkr.45.6.633>
- 18 McCready, E. (2015). *Reliability in Pragmatics*. Oxford, England: Oxford
19 University Press.
- 20 McCroskey, J. C., & Teven, J. J. (1999). Goodwill: A reexamination of the
21 construct and its measurement. *Communication Monographs*, 66(1), 90–103.
22 <https://doi.org/10.1080/03637759909376464>

- 1 Meehl, P. E. (1990). Appraising and Amending Theories: The Strategy of
2 Lakatosian Defense and Two Principles that Warrant It. *Psychological Inquiry*, 1(2),
3 108–141. https://doi.org/10.1207/s15327965pli0102_1
- 4 Meredith, M., & Kruschke, J. K. (2013). BEST: Bayesian Estimation
5 Supersedes the t-Test (Version .2.0). Retrieved from [http://cran.r-](http://cran.r-project.org/web/packages/BEST)
6 [project.org/web/packages/BEST](http://cran.r-project.org/web/packages/BEST)
- 7 Mills, C. M., & Keil, F. C. (2005). The Development of Cynicism.
8 *Psychological Science*, 16(5), 385–390. [https://doi.org/10.1111/j.0956-](https://doi.org/10.1111/j.0956-7976.2005.01545.x)
9 [7976.2005.01545.x](https://doi.org/10.1111/j.0956-7976.2005.01545.x)
- 10 Mills, C. M., & Keil, F. C. (2008). Children’s developing notions of
11 (im)partiality. *Cognition*, 107(2), 528–551.
12 <https://doi.org/10.1016/j.cognition.2007.11.003>
- 13 Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes
14 Factors for common designs. (Version R package version 0.9.12-2). Retrieved from
15 <https://CRAN.R-project.org/package=BayesFactor>
- 16 Olsson, E. J. (2011). A Simulation Approach to Veritistic Social Epistemology.
17 *Episteme*, 8(02), 127–143. <https://doi.org/10.3366/epi.2011.0012>
- 18 Olsson, E. J., & Vallinder, A. (2013). Norms of assertion and communication
19 in social networks. *Synthese*, 190(13), 2557–2571. [https://doi.org/10.1007/s11229-](https://doi.org/10.1007/s11229-013-0313-1)
20 [013-0313-1](https://doi.org/10.1007/s11229-013-0313-1)
- 21 Petty, R. E., & Briñol, P. (2008). Persuasion: From Single to Multiple to
22 Metacognitive Processes. *Perspectives on Psychological Science*, 3(2), 137–147.
23 <https://doi.org/10.1111/j.1745-6916.2008.00071.x>

1 Phillips, N. (2017). Yarr: A companion to the e-Book “YaRrr!: The Pirate’s
2 Guide to R”. (Version R package version 0.1.5). Retrieved from [https://CRAN.R-](https://CRAN.R-project.org/package=yarr)
3 [project.org/package=yarr](https://CRAN.R-project.org/package=yarr)

4 Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech.
5 *Proceedings of the National Academy of Sciences*, 105(3), 833–838.
6 <https://doi.org/10.1073/pnas.0707192105>

7 Poulin-Dubois, D., Brooker, I., & Polonia, A. (2011). Infants prefer to imitate
8 a reliable person. *Infant Behavior and Development*, 34(2), 303–309.
9 <https://doi.org/10.1016/j.infbeh.2011.01.006>

10 Poulin-Dubois, D., & Chow, V. (2009). The effect of a looker’s past reliability
11 on infants’ reasoning about beliefs. *Developmental Psychology*, 45(6), 1576–1582.
12 <https://doi.org/10.1037/a0016715>

13 R Core Team. (2016). *R: A language and environment for statistical*
14 *computing*. Retrieved from <https://www.R-project.org/>

15 Robinson, E. J., Champion, H., & Mitchell, P. (1999). Children’s ability to
16 infer utterance veracity from speaker informedness. *Developmental Psychology*,
17 35(2), 535–546. <https://doi.org/10.1037/0012-1649.35.2.535>

18 Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from
19 knowledgeable versus ignorant speakers: Links between preschoolers’ theory of mind
20 and semantic development. *Child Development*, 72(4), 1054–1070.

21 Schwarzer, G. (2007). meta: An R package for meta-analysis. *R News*, 7(3),
22 40–45.

23 Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive
24 Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows

- 1 Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366.
2 <https://doi.org/10.1177/0956797611417632>
- 3 Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., &
4 Wilson, D. (2010). Epistemic Vigilance. *Mind & Language*, 25(4), 359–393.
5 <https://doi.org/10.1111/j.1468-0017.2010.01394.x>
- 6 Teigen, K. H. (1988). The language of uncertainty. *Acta Psychologica*, 68(1–
7 3), 27–38. [https://doi.org/10.1016/0001-6918\(88\)90043-1](https://doi.org/10.1016/0001-6918(88)90043-1)
- 8 Teigen, K. H., & Brun, W. (1999). The Directionality of Verbal Probability
9 Expressions: Effects on Decisions, Predictions, and Probabilistic Reasoning.
10 *Organizational Behavior and Human Decision Processes*, 80(2), 155–190.
11 <https://doi.org/10.1006/obhd.1999.2857>
- 12 Toulmin, S., Rieke, R., & Janik, A. (1979). *An introduction to reasoning*. New
13 York NY: Macmillan.
- 14 Trafimow, D. (2009). The Theory of Reasoned Action: A Case Study of
15 Falsification in Psychology. *Theory & Psychology*, 19(4), 501–518.
16 <https://doi.org/10.1177/0959354309336319>
- 17 Trafimow, D. (2012). The role of auxiliary assumptions for the validity of
18 manipulations and measures. *Theory & Psychology*, 22(4), 486–498.
19 <https://doi.org/10.1177/0959354311429996>
- 20 Van Der Auwera, J. (1996). Modality: The Three-layered Scalar Square.
21 *Journal of Semantics*, 13(3), 181–195. <https://doi.org/10.1093/jos/13.3.181>
- 22 Verstraete, J.-C. (2005). Scalar quantity implicatures and the interpretation of
23 modality. *Journal of Pragmatics*, 37(9), 1401–1418.
24 <https://doi.org/10.1016/j.pragma.2005.02.003>

1 Wallsten, T. S., & Budescu, D. V. (1995). A review of human linguistic
2 probability processing: General principles and empirical evidence. *The Knowledge*
3 *Engineering Review*, 10(01), 43–62. <https://doi.org/10.1017/S0269888900007256>

4 Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van
5 Aert, R. C. M., & van Assen, M. A. L. M. (2016). Degrees of Freedom in Planning,
6 Running, Analyzing, and Reporting Psychological Studies: A Checklist to Avoid p-
7 Hacking. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.01832>

8

9

1

2