Hedged testimony

Peter van Elswyk, University of Wisconsin–Milwaukee

Penultimate draft · Forthcoming in *Noûs*

Speakers offer testimony. They also hedge. This essay offers an account of how hedging makes a difference to testimony. Two components of testimony are considered: how testimony warrants a hearer's attitude, and how testimony changes a speaker's responsibilities. Starting with a norm-based approach to testimony where hearer's beliefs are warranted because of social norms and speakers acquire responsibility from these same norms, I argue that hedging alters both components simultaneously. It changes which attitudes a hearer is prima facie warranted in forming in response to testimony, and reduces how much responsibility a speaker undertakes in testifying. A consequence of this account is that speakers who hedge for strategic purposes deprive their hearers of warrant for stronger doxastic attitudes.

1 Introduction

Testimony happens when a speaker presents a content to a hearer that they stand behind. When a hearer has comprehended what the speaker presented, doxastic responses to testimony are prima facie warranted for them in the absence of certain complications.[1] By DOXASTIC RESPONSE, I mean any of the ways a hearer may change their mind towards the content presented (*e.g.* updating credences, forming beliefs). A speaker commonly provides testimony by using a declarative to present a content determined by the meaning of the sentence.

But declaratives comes in at least two varieties. A speaker either uses a bare declarative like (1), or hedged declarative such as (2).[2] The difference between them is the strength or confidence with which the speaker stands behind the content. The presence of the epistemic term—in this case, the attitude verb *think* in a parenthetical position—weakens how forcefully the speaker backs the content presented.

    (1) The corgi splooted.

    (2) The corgi splooted, I think.

Hedging is a regular part of conversation. In some settings where its frequency has been measured, it occurs every fifteen seconds as people share information (Prince

---

[1] In what follows, I will omit these qualifications. In particular, I will omit the qualification that the hearer must understand or comprehend what the speaker said in what follows. Though comprehension is epistemically significant, it will not play a central role in this essay. Graham (2010) and Peet (2016, 2018) explore this issue. I will also omit the qualification about complications. The most obvious complication is the hearer having a defeater, but other complications exist. For example, see MacFarlane (2005) and Lackey (2008).

[2] The hedging discussed in this essay is sometimes called ILLOCUTIONARY HEDGING (Brown and Levinson, 1987; Fraser, 2010), or RELATIONAL HEDGING (Prokofieva and Hirschberg, 2014). Such hedging changes the strength with which a content is presented as opposed to changing the content itself. It contrasts with modifiers like *sorta* or *kinda* as in *The corgi sorta splooted*. Though these terms are sometimes called hedges because they also weaken, they weaken by fudging the content presented by a sentence.

et al., 1982). But hedging has been overlooked in historical and contemporary discussions of testimony. Though passing mention is occasionally made to hedged testimony (*viz.* Adler, 2002; Jackson, 2020), no sustained attempt has been made to detail how hedging makes a difference to the epistemology of testimony. On the contrary, testimony is regularly characterized as happening only if a speaker performs the speech act of assertion.[3] Since hedged statements like (2) are usually considered to be speech acts other than assertions such as conjectures, hedged testimony is excluded from qualifying as testimony.

This essay provides some of what's missing. My focus is on understanding how hedging makes a difference to two distinct but related aspects of testimony. The first concerns testimony as a source of evidence, or the WARRANT COMPONENT.[4] Doxastic responses to testimony are prima facie warranted for a hearer. The second component concerns what is normatively distinct about testimony, or the RESPONSIBILITY COMPONENT. Speakers undertake responsibility when they offer testimony. Other sources of evidence do not similarly bring about a normative difference. This essay develops a view where both components are simultaneously influenced: hedging change what doxastic responses are prima facie warranted for the hearer and what the speaker is responsible for.

I will not start from scratch. I adopt an approach to testimony where social norms are what underwrite both the prima facie warrant conferred to hearer's doxastic response and the responsibility incurred by the speaker (Graham, 2010, 2012, 2014, 2015, 2020a,b; Greco, 2016, 2020; Simion and Kelp, 2020; Simion, 2021). To extend a norm-based approach to hedged testimony, my central contribution is a new norm—the positional norm—that regulates testimony performed with bare and hedged declaratives alike. I will argue that one consequence of the positional norm is that hedging influences testimonial warrant by determining which doxastic responses are warranted. Another outcome is that hedging influences responsibility by determining what epistemic position the speaker is responsible for. In contrast to theories of testimony such as Moran (2005, 2018) that base the prima facie warrant conferred in a speaker undertaking responsibility, my norm-based approach will treat warrant-conferring and responsibility-undertaking as separate consequences of the same norm.

The essay is organized as follows. I begin by clarifying hedging as a linguistic phenomenon (§2). From there, I consider the etiological function of hedging, or what purpose hedging has (§3). With hedging better understood, I turn to how it influences testimonial warrant (§4). After that, I consider the responsibility component (§5). Before concluding, I discuss the epistemic consequences of strategic hedging (§6).

---

[3] See Reynolds (2002), Moran (2005, 2018), Hinchman (2005, 2020), Owens (2006), and Simion (2021). Goldberg (2010) has related criticism.

[4] I use *warrant* to name the truth-conducive support provided by testimony. Following Burge (1993) and others, warrant divides into justification and entitlement. I remain neutral on how to distinguish the two and which kind is provided by testimony. The view advanced here is compatible with various perspectives.

2 Hedging

Epistemic terms indicate an agent's evidence or attitudes. They come from various syntactic categories. Examples include verbs like *think* and *guess*, adverbs such as *probably* and *maybe*, auxiliaries like *might* and *can*, and adjectives such as *possible* and *alleged*. Such terms are typically theorized as relating an underlying content to a body of information. That underlying content is known as the PREJACENT. So the standard interpretation of an epistemic term is one where it contributes to the meaning of the declarative sentence by indicating a way in which the prejacent is related to a body of information.

Sometimes epistemic terms receive a non-standard, hedging interpretation. When terms hedge, their prejacent is what is relevant in conversation and the term is interpreted as weakening the strength with which the prejacent is presented. Consider (4b).

(3) (a) Which dog splooted?
     (b) The corgi splooted.

(4) (a) Which dog splooted?
     (b) I think the corgi splooted.

Both (3b) and (4b) respond to the same question, and both intuitively present the same content—that the corgi splooted—as an answer. However, (4b) differs by beginning with a verb indicating the speaker's attitude. That term is interpreted as a hedge. It weakens how forcefully its prejacent is offered as opposed to contributing to an answer about the speaker's thoughts.

Whether an epistemic term hedges depends partly on the broader discourse (van Elswyk, 2022). Exchange (5) demonstrates. When we keep the reply the same as (4b), but change the question being replied to, the hedging interpretation of *I think* disappears. The verb receives its standard interpretation as a contribution to the declarative's content. That content is then offered as a response to the question as opposed to the prejacent.

(5) (a) Which dog do you think splooted?
     (b) I think the corgi splooted.

As a result, (5b) is not weaker. It has a different content but the same default strength as the bare declarative (3b).

I have illustrated hedging with an attitude verb in a matrix position, but the phenomenon is general. Verbs, adverbs, adjectives, and auxiliaries can all be interpreted as hedges in a discourse (van Elswyk, 2022). The (b) through (e) replies show.

(6) (a) Which dog splooted?
     (b) The corgi splooted, I think.
     (c) Probably the corgi splooted.
     (d) the corgi might have splooted.

(e) It is possible that the corgi splooted.

There are various semantic and pragmatic explanations for how epistemic terms are interpreted as hedges.[5] Given this essay's focus, I take for granted that hedging happens to attend to why it happens.

What does merit elaboration is what the weakening effect consists in. I assume what I have defended elsewhere: that hedging consists in the suspension of the signal that the speaker knows the content conveyed by a declarative.[6] Accordingly, I further assume that the use of a bare declarative signals that the speaker knows its content.[7]

> KNOWLEDGE HYPOTHESIS (κ-hypothesis)
> For a speaker S and bare declarative d expressing primary content p in
> a context c, S's use of d signals S as knowing p in c.

Knowledge sets a threshold for hedging. An epistemic term hedges by overriding this default signal to indicate that the speaker occupies a position below the threshold. Weakness consists in lacking knowledge.

A consequence of this threshold view is that only non-factive terms can hedge. This consequence is easily confirmed. Reply (7b) features *know* in matrix position, and, for variety, (7c) features *learned* in parenthetical position. Both are factive and neither hedges.

> (7) (a) Which dog splooted?
> (b) I know the corgi splooted.
> (c) The corgi splooted, I learned.

Non-factive terms do not all have the same strength. So speakers can vary the strength with which they hedge according to which terms they hedge with. For example, *it is possible that* in (6e) hedges more than *probably* in (6b). Such a strength difference can be explained with entailment. (6b) entails (6e) but not *vice versa*.

---

[5] Some regard hedging as illocutionary force modification (Brown and Levinson, 1987; Bach, 2008; Fraser, 2010). Another explanation appeals to Gricean calculation (McCready, 2015). In certain discourses, speakers are uncooperative if they use a declarative containing an epistemic term as opposed to a bare declarative. From such uncooperativity, hearers reason to the conclusion that the speaker is hedging with the term on the presumption that they are broadly cooperative. A final explanation that has been defended appeals to how discourses are organized with coherence relations (van Elswyk, 2022). In certain configurations of such relations in a discourse, epistemic terms are argued to perform a dedicated hedging function.

[6] See Benton and van Elswyk (2020), van Elswyk and Sapir (2021), and van Elswyk (2022) for discussion of this assumption in connection to assertion, inquiry, and discourse structure.

[7] This hypothesis is widely associated with the knowledge norm of assertion (Williamson, 2000; Benton, forthcoming). But there is an important distinction to be drawn between what the use of a bare declarative represents or signals about the speaker's epistemic position, and what position a speaker is required to occupy by a norm. The knowledge norm runs these together, but they can and should be separated. See van Elswyk and Benton (2022) for more on the signal/norm distinction. Once we make this distinction, non-normative implementations of the κ-hypothesis are on offer. For example, the signal about the speaker's knowledge has been proposed to be owed to a covert operator (Meyer, 2013) or parenthetical verb (van Elswyk, 2021).

However, not all strength may be owed to entailment. In this essay, I just work with the assumption that speakers exploit strength differences when choosing how much to weaken their presentation.

In clarifying how hedging weakens the strength with which a content is shared, we clarify how to investigate the epistemology of hedged testimony. Hedging happens when speakers override the default signal that they know the content of a declarative to indicate that they occupy a weaker epistemic position. An epistemology of hedged testimony must account for how indicating a position short of knowledge impacts testimony.

3   The function of hedging

My entry into the epistemology of hedging is its function. Once we grasp what hedging is for, we can move to consider how hedging influences the warrant and responsibility components of testimony. To start, we need to distinguish SPEAKER FUNCTION from ETIOLOGICAL FUNCTION (Millikan, 1987). I work with a broad definition of the latter.[8]

ETIOLOGICAL FUNCTION (E-function)
An item β has the etiological function F in a system S iff:

(I) Previous βs had F in S.

(II) S benefited from previous βs having F.

(III) βs exist because F benefited S.

The definition has two key conditions. Condition (II) is a WELFARE CONDITION (Bedau, 1991; McLaughlin, 2012). It requires that the function of an item benefit the system in which the item is a component. Condition (III) is a connected EXPLANATION CONDITION (Wright, 1973). It states that the function's being beneficial explains why the item currently exists in the system. The heart's etiological function is a familiar example. Past hearts pumped blood. Pumping blood benefits the organism as the broader system in which the heart is a part. That benefit explains why the heart presently exists in various organisms.

The speaker function for an item is whatever purpose a speaker uses that item for in a context. In some contexts, the speaker function will be the etiological function. But it need not be. It is easy to imagine someone who is dramatically hedging everything they say in order to mock a friend for hedging too much earlier.

---

[8] This definition closely follows Graham (2014). It is broad in that it applies to organisms and artifacts by omitting biological notions (*e.g.* fitness, selection). The definition I adopt differs from a definition recently adopted by Kelp (2018). He uses a weak version of condition (III) where the beneficiality of F merely contributes to the explanation for why βs exist. For support, he cites Buller (1998). But it is not clear that Buller's selection-oriented motivations for weakening the explanation condition in a biological definition apply to a broad definition. I stick with a strong explanation condition because its demandingness is useful in comparing the plausibility of candidate functions that differ in subtle ways. In particular, it will allow us to distinguish between an E-function and common speaker functions.

The speaker function of hedging in this context is to mock. But that is not what hedging is for; mocking does not confer a benefit that explains why hedging exists. Accordingly, my focus is on hedging's ᴇ-function. Hedging is something speakers do with hedges, i.e., linguistic items receiving a special interpretation in a discourse. In what follows, I will therefore alternate between talking about the etiological function of hedging and hedges.

A proposal for hedging's ᴇ-function will satisfy two basic constraints. First, the proposed function will meet conditions (ɪ), (ɪɪ), and (ɪɪɪ). If it does not meet these conditions, the proposed function can only be a speaker function. Second, the proposed function will elucidate why we hedge the way we do. As clarified in §2, hedging happens when a non-factive epistemic term receives a special interpretation in a discourse. The ᴇ-function of hedging will clarify why hedging is accomplished with such terms as opposed to other terms or phrases. In this section, I entertain a few proposals before arguing that hedging's ᴇ-function is alerting hearers to epistemic risk.

3.1 Hedging and managing the speaker's reputation

McCready (2015) proposes that hedging helps a speaker manage their reputation. In particular, hedging is an attempt to exempt a particular act of testimony from influencing the speaker's reputation. If a speaker hedges their presentation of p, then they get no positive update to their reputation if p is true and no negative impact to their reputation if p is false.

I have no doubts that there is a link between hedging and reputation management. But managing speaker reputation cannot be the ᴇ-function. It does not meet either of the basic constraints. Start with the first. Linguistic communication is a cooperative undertaking between a speaker and hearer. To satisfy the welfare condition, the ᴇ-function of a linguistic item must benefit the speaker and the hearer alike (Sperber, 2001; Millikan, 2004).[9] Hedging is no exception. If hedging does not sufficiently benefit the speaker by prompting a preferred response in the hearer, the speaker would not have an interest in hedging. Hedges would then cease to be produced. Likewise, if hedging does not sufficiently benefit the hearer, the hearer would not be interested to respond in the way preferred by the speaker. Again, hedges would not be produced.

However, managing a speaker's reputation does not benefit the hearer. The hearer does not gain anything if the speaker refrains from investing their reputation in what they shared. On the contrary, they lose. Since hearers are also speakers, they know the speaker's motivations. Either the speaker is deceiving them about p and wants to insulate themselves from any hit to their reputation, or they are hesitant enough about p that they do not want to take a chance on improving their reputation by giving unhedgeg testimony. Consequently, deception becomes

---

[9] This follows from a more general principle about evolutionarily stable communication. Signals must benefit senders and receivers. See Dawkins and Krebs (1978), Smith and Harper (2003), and Scott-Phillips (2008).

a likely explanation for why the speaker hedged. In contrast to unhedged testimony where the hearer has no special reason to believe the speaker is deceiving, hedging does give the hearer such a reason.

In other words, hedging furnishes the hearer with an UNDERCUTTING DEFEATER (Pollock and Cruz, 1999). An undercutting defeater eliminates the warrant provided by a source of evidence for content p by providing auxiliary warrant that the source does not support p. In the present case, the speaker's hedge would defeat the warrant for p that their testifying p otherwise provides. Defeated testimony has no evidential value. So hedged testimony has no evidential value for the hearer when it just serves to protect the speaker's reputation.

Turn next to the second constraint. Epistemic terms indicate an individual's evidence or attitudes. They do not not have meanings that directly or indirectly relate to an individual's reputation in a group. Of interest is that we do have linguistic expressions whose function is exclusively to manage how testimony impacts reputation.

    (8) You didn't hear this from me, but the corgi splooted.

    (9) This is off the record: the corgi splooted.

The prefaces in (8) and (9) illustrate. McCready's proposal plausibly explains them as linguistic devices that shield a speaker's reputation from future damage. But, tellingly, the prefaces are not hedges. They do not weaken the strength with which a content is presented by overriding the default signal that the speaker knows. So the proposal is disconnected from the fact that we hedge with non-factive epistemic terms.

3.2 Hedging and managing speaker responsibility

Hedging is often characterized as diminishing the amount of responsibility a speaker incurs when presenting a content (Brown and Levinson, 1987; Fraser, 2010; McCready, 2015; Benton and van Elswyk, 2020). I will argue as much in §5. But some might be tempted to regard this diminishing effect as hedging's E-function. To hedge is to take less responsibility than one otherwise would if a bare declarative was used to present that same content.

Nevertheless, diminishing speaker responsibility cannot be the E-function. If diminishing was all there was to hedging, the welfare condition would not be fulfilled because the hearer would not benefit. A hearer does not gain anything by receiving testimony that the speaker is taking less responsibility for. They either lose the ability to hold speakers accountable or have that ability significantly diminished. Since accountability is the recourse hearers have if a content turns out to be false or unknown by the speaker, hearers lack meaningful recourse for hedged testimony that is learned to be defective.

The loss of recourse snowballs into another loss for the hearer similar to the loss mentioned above for the reputation management proposal. Deception constitutes

a likely explanation for the speaker is trying to shirk responsibility. That provides the hearer with a reason to not form a doxastic response in response to the speaker's testimony. That reason may not constitute a total defeater for the warrant provided by the speaker's testimony. Perhaps it is partial; perhaps that reason just gives cause to regard the amount or kind of testimonial warrant provided as lower. Even still, hedged testimony has significantly minimal evidential value for the hearer when it just weakens the speaker's responsibility.

The proposal under consideration also fails to meet the second constraint. It does not elucidate why we hedge with non-factive epistemic terms. Such terms do not directly or indirectly relate to the responsibility that speakers incur in communication. Epistemic terms like *probably* or *I guess* contrast with the prefaces in (10) and (11).

(10) Don't hold me to this, but the corgi splooted.

(11) I don't want to be blamed for saying this, but the corgi splooted.

These prefaces clearly attempt to alter how much responsibility, if any, the speaker incurs. But, like before, they fail to be hedges. They do not weaken the strength with which a content is presented by overriding the default signal that the speaker knows.

3.3 Hedging and alerting the hearer to positional risk

The previous proposals started with effects hedging has, and took those to be what explains why hedging continues to happen. My starting point is different. I submit that the ᴇ-function of hedging subserves the ᴇ-function of the declarative clause. For an analogy, consider the heart's right verticle. It is a chamber that relaxes to fill with blood from an atrium and contracts to send that blood into the pulmonary artery. Its function is to move blood between two places, but it subserves the broader function of the heart in pumping blood to an organism. Whatever the function of hedging is, it analogously contributes to the clause's purpose. Its own purpose is neither orthogonal nor incidental to the clause's.

I follow Millikan (1987) who proposes that the declarative's function is *producing true beliefs in hearers*. A speaker has an interest in being believed. It empowers them to influence the thoughts and actions of others. But a hearer is not merely interested in believing what they are told. The are interested in forming accurate beliefs. Accordingly, Millikan's proposal is that uses of declaratives create enough accurate beliefs for hearers to believe enough of what speakers tell them so that speakers continue to produce declaratives.[10]

---

[10] For a recent defense of Millikan's ᴇ-function, see Graham (2010, 2020a). Recent alternatives agree that the function of the clause is to produce an epistemic position but either downgrade or upgrade the position produced. For example, Murray and Starr (2018) suggest that the ᴇ-function is *producing mutual assumptions between speakers and hearers*. But they do not defend this suggestion in detail. I find it difficult to see what benefit a hearer would get from having their assumptions changed that would be enough to satisfy the welfare condition. Note that if the benefit is related to accuracy, their

For testimony to produce a true belief, epistemic risk must be managed. Epistemic risk is the risk of an inaccurate doxastic response. A false belief is an inaccurate response, and so is an updated credence that is not in accord with the truth.[11] Testimony as a source of evidence comes with two distinctive kinds of epistemic risk. The first is DECEPTION RISK. The speaker might be deceiving the hearer: if they are, a doxastic response will be based on an unreliable instance of testimonial evidence. That unreliability is substantial. Were a hearer to learn the speaker was deceptive, they would acquire an undercutting defeater for any information provided by the speaker.

The second kind of risk is what I call POSITIONAL RISK. The speaker might have a weak epistemic position: if they do, a doxastic response may have a less reliable source. Such unreliability may not always substantial enough to yield a defeater were the hearer to learn that the speaker had a weak position. But it still puts the hearer at greater epistemic risk. To illustrate, compare a doxastic response to testimony originating from a speaker who knows in comparison to a speaker who guesses. Knowledge is factive. So if the speaker knows, a doxastic response based on their testimony will be accurate. But guessing does not guarantee accuracy like knowledge. Guessing p is also compatible with p being improbable (Kahneman and Tversky, 1982; Dorst and Mandelkern, forthcoming). With speakers who guess, a doxastic response based on their testimony may or may not be accurate. That doxastic response is a riskier one.

Positional risk is importantly distinct from what we might call COMPETENCE RISK. The speaker may be unskilled at taking up an epistemic position towards p that is warranted. If they are, a hearer's doxastic response will have a less reliable source. But not because of what position is taken by the speaker. Compare a speaker sharing their guess that p when they know p is false with a speaker sharing their guess p even if p is improbable for them. Both speakers subject the hearer to positional risk by sharing guesses, but only the initial speaker puts the hearer at greater epistemic risk because of incompetence.

My proposal is that the E-function of hedging is *alerting hearers to positional risk*. Hedging enables speakers to share information while being in a variety of different epistemic positions towards that information. When speakers do not know, they can use a non-factive epistemic term to override the default signal that they know and indicate what position they do occupy. As a result, hedging ensures that hearers receive information that is accompanied with a risk assessment. They learn from the speaker what position the speaker backs the information with, and thereby learn the positional risk of forming a doxastic response. What hedging does, in other

---

account may not be meaningfully different from Millikan's. Kelp (2018) proposes that the E-function of assertion is *producing knowledge in hearers*. But this is not a true alternative. Millikan's proposal is about the declarative clause as opposed to assertion. Indeed, it is compatible with Millikan's proposal that assertions produce true beliefs in hearers by producing knowledge whereas other uses of the clause produce true beliefs differently.

[11] The accuracy of a credence depends upon what is known as a SCORING RULE. For recent discussion, see Joyce (1998, 2009), Predd et al. (2009), Pettigrew (2016), and Horowitz (2019). In what follows, I will continue to talk about accuracy generally.

words, is help the hearer form accurate responses from uses of declaratives by not subjecting them to avoidable risk.

To motivate the proposal, let's begin with a hypothetical. Consider the situation a speaker would be in if they were not able to hedge. Given the earlier assumption of the κ-hypothesis that the use of a bare declarative signals that the speaker knows (§2), the speaker would have two options when they did not know the content of a declarative. Either use the declarative and deceive the hearer by signaling that they know when they do not, or remain silent. Being limited to these options is what I call the SILENCE/DECEPTION DILEMMA. Neither is beneficial. Silence does not enable the speaker to access the benefit of changing the hearer's mind. Deception does, but with considerable social costs if the speaker is caught. I detail some of these costs and their significance later in §4.

The hearer would not benefit in this situation either. They receive silence or deception. Silence precludes them from accessing testimony that is still of evidential value. Though that value is below that had by unhedged testimony because it is riskier, hedged testimony warrant some doxastic responses. Nor is it beneficial for the hearer to be deceived. In this hypothetical, the hearer would be unable to distinguish testimony based on knowledge from testimony based on epistemic positions that fall short of knowledge. So they would not able to tailor their doxastic response to the speaker's position. Instead, the hearer would proceed as if the speaker knew even though they did not. In doing so, they would be at greater risk of forming an inaccurate response.

The value of hedging is that it solves the silence/deception dilemma. A speaker who does not know the content of a declarative can speak up by hedging. By hedging with the epistemic term that correctly indicates the position they occupy, the speaker does not deceive the hearer about their epistemic position. They have overridden the default signal that they know and indicated their true position. As a consequence of navigating the dilemma, both the speaker and the hearer benefit.[12] The speaker benefits by prompting a doxastic response from the hearer. That response may be weaker than being believed. But it is still the same kind of benefit a speaker receives from using a declarative that is tied to the clause's etiological function. That benefit is attained without becoming liable to the social costs that deception incurs liability for.

The benefit hearers receive in turn is testimony that has evidential value. That value will be lower than what they receive from unhedged testimony. But it still

---

[12] It might be objected that speakers mostly lose in hedging by appearing less confident or authoritative in what the say (Hosman, 1989). But social science research does not support this conclusion. Repeated studies show that hearers believe speakers more when the speakers are perceived to be calibrated, i.e., when speakers express hesitation when inaccurate and confidence when accurate (Tenney et al., 2007, 2008, 2019; Sah et al., 2013; Vullioud et al., 2017). For example, Tenney et al. (2007) found that participants judge speakers who testified inaccurately to be more believable if they expressed hesitation. In contrast, speakers who testified inaccurately but were confident were judged to be less believable. Hedging contributes to a speaker being calibrated or regarded as such. It is the means by which a speaker expresses hesitation by overriding the confident signal that they know.

provides them with warrant for a doxastic response; it still provides them with a rational basis on which to change their mind about how the world is. Though such testimony carries greater positional risk, the hearer is alerted to that risk. They are able to distinguish what epistemic position underwrites the testimony because the speaker indicated as much with the hedging term. As a result, the hearer can tailor their doxastic response accordingly. They do not have to riskily proceed as if the speaker knew what they said.

Important to the plausibility of the risk alerting proposal is that it meets the two constraints for adequacy that the previous proposals did not. As the preceding discussion makes clear, managing positional risk for the hearer fulfills the welfare condition by benefiting both speaker and hearer alike. It does by being a hearer-centric function that subserves the etiological function of the declarative clause. The risk alerting proposal also straightforwardly explains why we hedge with non-factive epistemic terms. We hedge with such terms because they are the means by which a speaker alerts the hearer to the positional risk of their testimony. For example, a speaker hedges by saying *The corgi splooted, I guess* because the parenthetical verb enables them to signal they guess as opposed to know that the corgi splooted. The hearer can then choose what doxastic response to adopt given the positional risk that comes with guesses.

As a hypothesis about hedging's E-function, the risk alerting proposal is compatible with speakers hedging for other purposes. It is even compatible with reputation and responsibility management being regular functions speakers use hedging for (§6). But what's central to the hypothesis is that the risk alerting function explains why hedging continues to happen in conversation. It exists for speakers to help produce accurate responses in hearers by helping them navigate the epistemic risk associated with testimony because it is mediated through a speaker's epistemic position.

4   Hedging and the warrant component of testimony

We turn now to the warrant component of testimony. As advertised, my starting point is a norm-based approach to testimony. When it comes to warrant, such an approach is ANTI-REDUCTIONIST in that it does not explain testimonial warrant exclusively via facts about the hearer and their evidence (*e.g.* perception, memory, inference).[13] Instead, social norms are what underwrite the prima facie warrant had by a hearer's doxastic response to testimony. Norm-based theories mostly differ over what additional requirement, if any, the hearer must fulfill for their belief to be

---

[13] In addition to those cited who defend a norm-based approach, general defenses of anti-reductionism or what Pritchard (2004) calls CREDULISM can be found in Coady (1992), Burge (1993), Fricker (2007), Perrine (2014), and Zollman (2015). Nascent forms of anti-reductionism can occasionally be found in the social science literature on social norms and communication. See Scott-Phillips (2010). It should also be noted, as a referee reminds, that a norm-based approach can be implemented in a reductionist way too. For example, the hearer's awareness of social norms could be what provides prima facie warrant. I proceed with a anti-reductionist implementation in what follows.

warranted. I paper over this issue by assuming that norms are sufficient for doxastic responses to testimony to possess prima facie warrant. Extra requirements, if there are any, can be added to the view of hedging and testimonial warrant developed here without difficulty.

Testimony cannot prima facie warrant a doxastic response if the epistemic risk of forming an attitude is too high. In §3.1, two forms of risk were identified: deception and positional risk. Both risks need to be sufficiently minimized or prevented. Otherwise, doxastic responses will be at heightened risk of being inaccurate either because there is a good chance the speaker is engaged in deception, or that their epistemic position is too weak for a hearer's doxastic response to be prima facie warranted. The approach I develop will explain how social norms minimize both kinds of risk. It has been previously appreciated that norms minimize deception risk. That norms minimize positional risk will be a new upshot of a norm-based approach once room is made for hedged testimony.

I begin by explaining how social norms ensure that the declarative clause reliably performs its ᴇ-function (§4.1). Then I turn to how norms help hedging reliably fulfill its ᴇ-function of alerting hearers to positional risk (§4.2). My contribution is a new norm—the positional norm—that regulates testimony performed with bare and hedged declaratives alike. I argue that a consequence of the positional norm is that a hearer is prima facie warranted in coming to occupy at most the epistemic position that the speaker signaled that they occupy (§4.3). As a result, hedging influences testimonial warrant by determining which doxastic responses are warranted for the hearer.

4.1 Function and warrant

An item's ᴇ-function explains its existence in a system only if that function is fulfilled enough. When it is fulfilled enough, the benefit conferred on the system by that function ensures that the item persists. What counts as *enough* varies between functions. Graham (2014) distinguishes between functions where enoughness requires reliable fulfillment and functions where it does not. Let's call the former ʀᴇʟɪᴀʙʟᴇ ꜰᴜɴᴄᴛɪᴏɴꜱ or ʀ-functions. The heart's ᴇ-function is an ʀ-function. If the heart does not pump blood reliably, the benefit conferred on the organism would not be enough. In contrast, a sperm's function does not require reliable fulfillment. Very few sperm fertilize eggs, and yet, if one in a hundred million does, the benefit conferred is enough for sperm to exist.

The ᴇ-function of the declarative clause is an ʀ-function (Graham, 2010, 2020a). It is not enough for a small percentage of uttered declaratives to produce true beliefs in hearers. Otherwise speakers would benefit in such a lopsided way from the use of a declarative that hearers would cease having doxastic responses to them. Uses of clauses must produce true beliefs reliably or the benefit conferred from their use is not enough for the clause to persist.

For a communicative item to fulfill its function enough, an external mechanism

is often required to stabilize it. Consider the red deer stag's roar. Its function is to communicate the stag's social dominance. The louder the roar, the more dominance is communicated. Roar capacity is directly correlated with the stag's size. So the roar's function is reliably fulfilled merely by being performed (Reby and McComb, 2003). Human communication is not like that; a declarative does not produce true beliefs just because it was uttered. A mechanism outside of the language is needed for the function to be reliably fulfilled.

Social norms are that stabilizing mechanism for human communication (Scott-Phillips, 2010; Turri, 2017; Graham, 2020b).[14] One way they stabilize the function of the declarative clause is by minimizing deception. Deception happens when speakers fail to cooperate by preferring what is only in their self-interest when using declarative sentences (Faulkner, 2011; Simion, 2021). But social norms are a mechanism for bringing about cooperative behavior of group members even when that behavior is not in their immediate self-interest. They transform situations where group members would otherwise struggle to cooperate because they have partially opposed interests into situations where group members cooperate more easily because they have aligned interests (Bicchieri, 2006; Gintis, 2009). They are well-suited to minimize deception risk.

On a norm-based approach to testimony, social norms create enough reliability by minimizing deception to enable prima facie warrant. In other words, a hearer's doxastic response to testimony is prima facie warranted because a declarative's use reliably produces true beliefs. It reliably produces true belief because speakers comply with social norms to avoid deception. As Greco (2020, 76) helpfully puts it, "Just as natural laws structure the physical environment so as to determine the flow of information from world to mind, social norms structure the social environment so as to determine the flow of information from mind to mind." Social norms ensure quality control in the flow of information.

Characterizing social norms as a single mechanism for producing cooperation is somewhat misleading. Norms bring about cooperation in a variety of interrelated ways. To illustrate, I briefly highlight three. First, social norms are internalized. Internalization is a process whereby group members change their preferences to value norm-compliance as its own objective even when it may not be in their self-interest (Rozin et al., 1999; Sripada and Stich, 2006). Compliance is motivated for its own sake. Second, social norms are enforced. That enforcement might not be more than punishment in the form of criticism or blame, but enforcement motivates compliance to avoid punishment in some variety (Tomasello, 2019). Finally, social norm violation is an input to other cooperative mechanisms. Reputation provides a relevant example. When reputations are monitored, cooperation occurs between individuals because they want to have good enough reputations to unlock future

---

[14] I do not assume that norms are all that stabilize. Norms are what primarily explains cooperation between group members who are not kin. Kinship is well-known to create the conditions for cooperation (Hamilton, 1964; Smith, 1964). Avoiding deception is no exception (Fitch, 2004). So it is to be expected that additional mechanisms exist to ensure that friends and family avoid deception with each other.

benefits from third-parties (Wedekind and Milinski, 2000; Milinski et al., 2002; Nowak and Sigmund, 2005). But norm violation tarnishes a reputation. Even if no immediate punishment is faced by a person for violating a norm, potential reputational loss motivates compliance.

So does testimony provide prima facie warrant by being a type of utterance, or by being an utterance token made by a particular speaker?[15] I think the answer is *both* with a norm-based approach. Utterances or uses of declaratives that are norm-governed are a type. In virtue of being of this type, testimony offered with the use of a declarative provides prime facie warrant. But what it is to be norm-governed is for speakers to have certain preferences to avoid deception. That includes the speaker of any particular use of a declarative. That particular speakers have such anti-deception preferences contributes to why quality control results from being norm-governed.

4.2 The positional norm

Like the ᴇ-function of declaratives, the ᴇ-function of hedging is plausibly an ʀ-function. Hedging must reliably enable hearers to avoid positional risk. Otherwise hearers would not tailor their doxastic responses according to the risk associated with the position signaled by the speaker. They would either not change their mind, or just believe what was testified despite the risk they were alerted to. Without hearer's tailoring, speakers would cease to have an interest in hedging. They would benefit more from remaining silent, or signaling that they know the content of a declarative even when they do not.

An external mechanism is therefore needed to stabilize the function of hedging, i.e., to ensure that it enables hearers to reliably be alerted to positional risk. My proposal is that the same social norms that ensure that the declarative clause reliably performs its ᴇ-function similarly ensures that the ᴇ-function of hedging is reliably performed. They do double duty. To develop this proposal, let's consider what the norms are that minimize deception.

Specifying a norm is an idealized way of characterizing what the behavioral regularity is that is brought about by shared expectations and preferences. Some who defend a norm-based approach do not go into much detail about the content of the norms. For example, Graham (2020a, 741) notes that *telling the truth* is a textbook example of a social norm in the social science literature, and leaves it at that. The problem with characterizing the norm as *telling the truth* is that it fails to capture how a speaker's telling of the truth is mediated through their epistemic position. Norm compliant speakers tell what they believe or know is the truth from their perspective. Kelp (2018) and Simion (2021) avoid this problem by proposing that the norm is the ᴋ-norm.[16]

---

[15] Thanks to a referee for raising this question.

[16] It is an empirical question which social norm is operative. The ᴋ-norm enjoys a considerable amount of empirical support. See Turri (2017) for an introduction, and Turri (2018, 2021) for later findings. To my knowledge, the norms surrounding hedging have not been investigated as carefully.

KNOWLEDGE NORM (K-norm)
One must: assert that p only if one knows p.

But the K-norm is not helpful for developing a norm-based approach that applies to hedged testimony. Since the norm explicitly applies only to assertions, it does not apply to hedged statements because they are never used to perform assertions (Adler, 2002; Garcia-Carpintero, 2004).

What is required is a norm that applies to all uses of declaratives as opposed to specialized uses. Such a norm will provide a better characterization of the truth-telling behavior related to hedged and unhedged statements. I propose the following norm:

POSITIONAL NORM (P-norm)
One must: use a declarative d with content p in a context c only if one occupies the epistemic position $\mathcal{E}$ signaled by d in c.

The P-norm applies broadly. What it requires depends on what epistemic position the speaker signals themselves as occupying with the use of a declarative. Since what is signaled depends on context because whether an epistemic term hedges depends on the discourse in a context (§2), the P-norm is context-sensitive. Accordingly, the P-norm can be understood as yielding an array of similar norms once a discourse is fixed.[17] Where p is a variable for the content of a declarative and making the assumption that the relevant epistemic terms are interpreted as hedges, Table 1 illustrates.[18]

STATEMENT	SIGNAL	REQUIREMENT
p	being known	One must: know p.
p, I think.	being thought	One must: think p.
I guess that p.	being guessed	One must: guess p
Probably p.	being probable	One must: have p be probable.
It is possible that p.	being possible	One must: have p be possible.

Table 1: P-norm requirements

When a declarative carries the default signal that the speaker knows its content, the P-norm in that context is extensionally equivalent to the K-norm. It requires

---

[17] Greco (2015) also suggests that hedged statements are governed by norms requiring the speaker to occupy the position indicated by the hedge term. But he does not derive such norms and the knowledge norm from a common norm. Instead, he appeals to the notion of probabilistic knowledge found in Moss (2018). In van Elswyk (2022), I argue against assimilating hedging uses of epistemic terms to Moss's probabilistic approach to epistemic terms.

[18] One complication is hedging with terms that indicate a source of evidence for a proposition as opposed to an attitude or how the proposition is related to a body of information. Usually, such terms implicitly carry information about the speaker's doxastic attitude. Consider *conclude*. It is used to mark a proposition as inferred. But it appears to involve a belief-like attitude. This is evidenced by the infelicity of *I conclude that the corgi splooted, but I don't believe that*. In contrast, *heard* does not similarly involve belief. Discourses like *The corgi splooted, I heard, but I don't believe that* are felicitious. That terms indicating hearsay like *heard* do not involve belief appears to a robust, cross-linguistic fact (AnderBois, 2004).

knowledge. This is what the first line of Table 1 illustrates. But when the default signal is overridden because the speaker hedged, the P-norm requires the speaker to occupy whatever position is indicated by the epistemic term that was used to hedge. This is what all other lines show.

What the norm prohibits is POSITION MISREPRESENTATION in the signals sent by the use of a declarative. A speaker can misrepresent their epistemic position in two ways. First, speakers can anti-represent. They can occupy a position towards ¬p but signal a position towards p. An example is a speaker who believes that the corgi did not sploot but uses either a bare or hedged declarative to communicate that the corgi did sploot. They signal that their epistemic position supports p when their position supports its denial. Second, speakers can over-represent. They can occupy a position towards p that is weaker than the position signaled. An example is a speaker who signals with a bare declarative that they know that the corgi splooted, when, in reality, they merely guessed.[19]

It is by prohibiting these two kinds of misrepresentation that the P-norm performs the double-duty of ensuring that the declarative clause and hedges reliably perform their E-functions. Like the norms appealed to in other norm-based approaches, the P-norm ensures that declaratives reliably produce true beliefs in hearers by minimizing deception. Speakers who comply with the P-norm do not anti-represent. In other words, they do not deceive.[20] But speakers who comply with the P-norm do not over-represent their position either. When speakers share information backed by a weak epistemic position, they signal as much. So the P-norm also ensures that hedges reliably fulfill their E-function in alerting hearers to positional risk.

4.3 A warrant threshold

But which doxastic response is prima facie warranted? The limited focus on unhedged testimony offered with a bare declarative has guided extant work to only consider when a belief is prima facie warranted. But not every instance of hedged testimony warrants belief. In one of the few passing mentions of hedged testimony, Jackson (2020) notes that some instances of hedged testimony intuitively fall short of warranting outright belief. Consider (12b).

---

[19] Note that the P-norm does not uniformly prohibit under-representation. Consider a speaker who utters *The corgi splooted, I believe* even when they know. Since knowledge entails belief, the speaker complies with the P-norm by occupying the position signaled. However, there is still a sense in which the speaker has misrepresented their position by misleading the hearer. See *fn.* 27 below for more on how under-representation can mislead.

[20] Leading definitions of lying all have the consequence that lying requires anti-representation. For example, Benton (2018) and Holguín (2021) defend a knowledge-based definition on which to lie is to say p while knowing p is false. Signaling that you know p while knowing it is false is anti-representation. Consider next the definition in Davidson (1985) that is further developed by Fallis (2013) where to lie is to intentionally represent oneself as believing what one does not. Again, signaling that one believes what one does not is anti-representation. Finally, consider definitions like Saul (2012) and Stokke (2013, 2018) where a lie is to say or assert what you do not believe. Since the use of a declarative signals that one knows or believes, disbelieved assertions or statements involve anti-representation.

(12) (a) Which dog splooted?
    (b) It is possible the corgi splooted.

Testimony the corgi splooted that is backed by a speaker for whom the splooting is merely epistemically possible is weak testimony. It does not seem strong enough as evidence to warrant outright belief. In other words, a hearer who believes that the corgi splooted on the basis of (12b) alone forms a doxastic attitude that seems to outstrip what (12b) offers as evidence.

Were we to generalize from (12b), we might reach for the conclusion that hedging never prima facie warrants belief. But this would be hasty. The reply in (13b) illustrates why.

(13) (a) Which dog splooted?
    (b) The corgi splooted, I believe.

Above, the *believe*-parenthetical is interpreted as a hedge. It suspends the default signal that the speaker knows that the corgi splooted. But a hearer who believes on the basis of (13b) does seem prima facie warranted. Full belief is not a doxastic response that exceeds what (13b) offers as evidence.

What reflection on the previous discourses reveals is that hedging makes a difference to *which* doxastic responses to testimony are prima facie warranted for the hearer. Belief is warranted by all unhedged testimony, but it is only warranted by some instances of hedged testimony. The intuitive distribution of which responses to which instances of hedged testimony are prima facie warranted can be captured by this generalization.

WARRANT THRESHOLD GENERALIZATION (WT-generalization)
For all acts of testimony $\mathcal{T}$ carrying the signal that the speaker occupies epistemic position $\mathcal{E}$, $\mathcal{T}$ prima facie warrants a doxastic response that is equal to or less than $\mathcal{E}$.

Put more hearer-centrically, the WT-generalization states that hearers are prima facie warranted only in coming to occupy an epistemic position that is the same as or weaker than whatever epistemic position the speaker signaled themselves as occupying. The generalization earns its keep by fitting the examples. The reason why believing p in response to (12b) does not seem prima facie warranted is because the speaker signaled that p is merely possible for them. Since being a mere epistemic possibility is weaker than being believed, the hearer is not prima facie warranted in having stronger doxastic response. In contrast, believing p is a response to (13b) that is prima facie warranted because the speaker signaled with (13b) that they believe p. So the hearer's doxastic response stands in parity to what position was signaled by the speaker.

The WT-generalization also earns its keep by yielding accurate expectations about when doxastic responses are prima facie warranted. If we assume the principle, we will expect that there are no cases of hedged testimony where—on

the basis of such testimony alone—the hearer is prima facie warranted in having a much stronger doxastic response. No such cases are forthcoming.[21] However, it is strikingly easy to imagine cases where a hearer is warranted in having a doxastic response that is weaker than the position signaled by the speaker. For example, the hearer might judge that the speaker was communicating with them in a suspicious manner. This evidence is not strong enough to defeat the speaker's testimony, but it is enough to compel the hearer to take a weaker doxastic response. In doing as much, the hearer does not seem to do something that is closed to them epistemically. Their weaker response is prima facie warranted.

An adequate account of hedged testimony will explain the wt-generalization, i.e., why the distribution of prima facie warranted responses to testimony corresponds to when the responses are equal to or less than the strength of the position signaled. The p-norm enables such an explanation. There are two parts to the explanation. The first is explaining why a doxastic response is warranted. The second part is explaining which one is warranted.

On the hypothesis that the p-norm is operative as a social norm choreographing communicative behavior, speakers will reliably signal their position. As noted in §4.2, reliably signaling one's position minimizes deception as a kind of position misrepresentation. By minimizing deception, the p-norm makes *a* doxastic response prima facie warranted. The hearer is prima facie warranted in changing their mind in response to a speaker's testimony because the hearer can presume the speaker is not attempting to deceive them. This part of the explanation is standard to norm-based anti-reductionisms about warrant. Deception is owed to a breakdown in cooperation, and social norms cultivate cooperation.

Next, let's consider which response is prima facie warranted. The p-norm compels speakers to signal what doxastic response *they* are warranted to have. It does this by discouraging speakers from over-representing their position—from signaling that they occupy a position stronger than what is warranted for them. In reflecting on how terms like *probably* and *perhaps* are used, Toulmin (1958, 90-91) makes a similar point:

> Our probability-terms come to serve... not only to qualify assertions, promises and evaluations themselves, but also as an indication of the strength of the backing which we have for the assertion, evaluation or whatever. It is the quality of the *evidence* or *argument* at the speaker's disposal which determines what sort of qualifier he is entitled to include in his statements... By qualifying our conclusions and assertions in the ways we do, we authorise our hearers to

---

[21] Note that the wt-generalization is limited to what is prima facie warranted on the basis of testimony alone. It is compatible with a hearer conjoining the speaker's testimony with other evidence to beprima facie warranted in adopting a stronger doxastic response. To illustrate, suppose *A* hedges more than they need to by using *probably* when they know and weaker terms when *p* is just probable for them. Suppose *B* knows this about *A* on the basis of personal history. Then *B* is prima facie warranted in taking themselves to know on the basis of their prior evidence and *A*'s testimony. One way to understand what is happening in this situation is that *B* develops a speaker-specific interpretation of *probably* (Schuster and Degen, 2020). As a result, *B*'s prior evidence plays both an evidential role in warranting their doxastic response and a metasemantic role in resolving the context-sensitive meaning of *probably*.

> pur more or less faith in the assertions or conclusions, to bank on them, to rely on them, treat them as correspondingly more or less trustworthy.

As such, hearers who adopt a stronger position than what the speaker signaled for themselves adopt a position that outstrips what the speaker's warrant provides. If the hearer were the speaker with their evidence and reasons, the hearer would not be warranted in adopting the stronger position. Put differently, testimony transmits the speaker's warrant for a particular doxastic attitude. So hearers forming a stronger attitude in response do not have access to warrant for that attitude. What is transmitted to them is just warrant for a weaker doxastic attitude. The wт-generalization results. The doxastic responses that are prima facie warranted are those equal to or less than what's signaled because the р-norm only ensures that what's signaled is what's warranted for speakers.

We need to keep in mind that not all testimony offered by a declarative containing an epistemic term is hedged testimony. We should heed the warning of Wittgenstein (1953, 192) and not "regard a hesitant assertion as an as assertion of hesitancy," or the other way around. Accordingly, a sentence like *I think the corgi splooted* can be used for hedged and unhedged testimony. The kinds of testimony differ in both the content testified and the position that is signaled. When the sentence is used for hedged testimony, what's testified is that the corgi splooted and what is signaled is that the speaker merely thinks this content. When offered as unhedged testimony, the content is that the speaker thinks the corgi splooted and what is signaled is the default that the speaker knows as much. As a result, hedged testimony and unhedged testimony will prima facie warrant different doxastic responses. Unhedged testimony will prima facie warrant belief the hearer can presume is knowledge in an epistemic content. Hedged testimony will prima facie warrant a weaker response to a regular content.

So what does this look like in practice? Warrant depends on what is signaled, and what is signaled is expressed in natural language. So what positions or attitudes are warranted—what is equal to or less than what is signaled—depends on how hearers interpret the epistemic terms. A complete discussion of what is prima facie warranted would take us far afield into the semantics of epistemic terms that can be interpreted as hedges. To discuss what happens in practice while staying on topic, I will focus on three categories that are representative of the variety of epistemic terms. In doing so, I will characterize terms at an altitude that blurs details that are relevant to a compositional semantics.

The first category is non-factive attitude verbs like *believe*, *think*, and *guess*. Whether appearing in matrix or parenthetical positions, these terms are commonly used to hedge (Mackenzie, 1987). What doxastic response the corresponding hedged testimony warrants is the doxastic attitudes these terms denote. The obvious suggestion for *believe* is that it warrants belief. But it might be that they all denote the same attitude.[22] I will not adjudicate this issue here since the norm-

---

[22] Some argue believing is thinking (Hawthorne et al., 2016; Rothschild, 2020), some argue thinking is guessing (Holguín, forthcoming–), and still others argue guessing is believing (Dorst and

based account is indifferent to its resolution. Whatever they denote is what the corresponding testimony warrants.

The second category is epistemic modals like the adverb *possibly* and the auxiliary *must*. The usual way of understanding a modal's semantic contribution due to Kratzer (1977) is quantificationally. Possibility modals behave as existentials. To say that *p* is possible is to say that, given a set of worlds determined by a body of evidence, *p* is true at one of those worlds. Likewise, necessity modals behave as universals. To say that *p* must be the case is to say that *p* is true at all of the relevant worlds. Assuming that hedging uses of modals always exploit a body of evidence that is at least partially the speaker's, testimony hedged with a modal prima facie warrants the hearer in adjusting their body of evidence to support the relevant quantification. For example, *It is possible the corgi splooted* prima facie warrants the hearer in adjusting their body of evidence so that the corgi's splooting is possible. Depending on the body of evidence is constituted, adjusting may require abandoning some evidence or beliefs.

In a range of cases, the hearer's body of evidence will already support the possibility of what is testified. Imagine a hearer who knows the speaker has two dogs—a Pembroke Welsh corgi, a Portuguese water dog—but little else about the dogs or their behavior. It will already be possible for them that the corgi splooted. But this does not mean that the hedged testimony is epistemically useless. The speaker's testimony still provides prima facie warrant. If the hearer was not previously warranted in having their body of evidence support the possibility, they are now. Even if they were already warranted, modals direct attention to which possibilities are live relative to a body of evidence (Ciardelli et al., 2011). So hedged testimony like *It is possible the corgi splooted* can make the hearer newly aware of what epistemic possibilities are warranted for them.

The final category is probabilistic terms like the adverb *probably* or the adjective *likely*. Suppose the semantic contribution of these terms can be understood with quantification alone. Then what the corresponding hedged testimony warrants is akin to what the modal terms warrant. But many opt for richer semantic theories where they contribute to contents that are probabilistic (Yalcin, 2010). To illustrate with a particular approach, suppose the semantics of Moss (2015, 2018). Her account is helpful to work with because she draws a distinction between two interpretations of epistemic terms: "sentences containing epistemic vocabulary are sometimes used to assert thoroughly probabilistic contents, and sometimes used to assert contents about contextually determined bodies of evidence (2018, 36)."[23] Take *The corgi probably splooted*. On one interpretation, the content of this sentence is the same kind of content had by *The corgi splooted*. It is not a thoroughly probabilistic

---

Mandelkern, forthcoming).

[23] Moss does not characterize the probabilistic meaning as a hedging interpretation. Nor should it be. As I argue in van Elswyk (2022), terms can have probabilistic meanings without hedging and terms can hedge while being about a contextually determined body of evidence. But her distinction between two kinds of meaning for epistemic terms is useful to to work when considering how a sentence containing the same term might be used to offer both hedged and unhedged testimony depending on the discourse.

content. It is a content about a contextually determined body of evidence. On the other interpretation, the sentence does have a thoroughly probabilistic content that is unlike the content of the bare declarative.

For Moss, what asserting a probabilistic content does is instruct a hearer on how to adjust their credences. The content can be thought of as indicating a property that a credence should have. For example, asserting *The corgi probably splooted* instructs the hearer to have a credence in the content that the corgi splooted that is greater than .5. I submit that Moss's account of asserting probabilistic content can be understood as characterizing what hedged testimony with probabilistic terms warrants. Unhedged testimony warrants a hearer in having belief they can presume is knowledge in a content about a contextually determined body of evidence. In contrast, hedged testimony warrants a hearer in adjusting her credence in a way matching the probabilistic term.

Other ways of connecting the epistemic terms speaker use to hedge with the doxastic responses warranted for hearers can be worked out. The precise relationship between belief and credence, for example, will be relevant to understanding how different epistemic terms order in strength (§2).[24] But the above discussion highlighted that many kinds of responses can be prima facie warranted due to what the speaker signaled. On the basis of hedged testimony, hearers can come to be prima facie warranted in believing, adjusting their credences, or altering their body of evidence so it supports a possibility.

5 Hedging and the responsibility component of testimony

Having seen how hedging matters to the warrant component of testimony, the essay turns towards understanding how it influences the responsibility component. The responsibility component is where testimony's interpersonal character shines brightest. The normative relationship between speakers and hearers changes when a speaker offers testimony. Other sources of evidence do not similarly bring about a normative difference between two or more individuals. As a result, an account of how hedging influences the responsibility component is key to understanding the interpersonal character of hedging.

On some ways of theorizing testimony, the warrant and responsibility components are treated as disjoint. For example, reductionists cannot explain the responsibility component by only appealing to facts about the hearer and their evidence. With a norm-based approach to warrant, there is no need for explanatory disjointedness. The responsibility speakers incur in offering testimony can be traced back to the same social norms that underwrite testimonial warrant. In this section, I will continue to to develop a norm-based view that has explanatory coverage over both hedged and unhedged testimony.

---

[24] To illustrate, suppose belief is equivalent to credence at or above a context-sensitive threshold (Sturgeon, 2008; Demey, 2013; Leitgeb, 2014; Dorst, 2019). Then hedged testimony warranting an update to credences at or above that threshold would also warrant belief. But hedged testimony warranting weaker updates will not.

Hedging is widely recognized as changing or diminishing what a speaker is responsible for in uttering declaratives (Brown and Levinson, 1987; Fraser, 2010; McCready, 2015; Benton and van Elswyk, 2020). This is why we considered it as a candidate etiological function (§3.2). But the mechanics of how that normative change is caused have gone unexplained. What this section offers is a norm-based explanation of how hedging changes speaker responsibility and how this impacts the interpersonal character of testimony.

The section starts by detailing what responsibility consists in, and how social norms account for it (§5.1). From there, I situate hedging within the norm-based explanation of responsibility (§5.2). The view that emerges is that speakers reduce responsibility by signaling weaker positions because what they are responsible for via the p-norm depends on what was signaled. Weaker signals produce weaker requirements.

5.1 Norms and responsibility

Speakers are responsible for their actions and the outcomes that result. In offering testimony, speakers incur responsibility related to what they testified. Since the proposition testified is not an action, it is a category mistake to say speakers are responsible for the proposition or its truth. Instead, the way to understand the change in responsibility is that speakers take responsibility for signaling that they occupy a particular epistemic position.[25]

Given the κ-hypothesis, the position signaled by a bare declarative is that a speaker knows. From this signal, speakers incur the responsibility to know what they state. We can see evidence of this responsibility in how speakers are challenged. As (14b) and (14c) illustrate, hearers challenge speakers in ways that presume that the speaker knows the content presented (Unger, 1975; Williamson, 2000; Turri, 2010).

(14) (a) The corgi splooted.
(b) How do you know that?
(c) You don't know that.

As a result of the responsibility to know, speakers are blameworthy if they do not. Hearers can properly blame, criticize, or censure speakers when they catch them believing the opposite, or even when catch them sharing information with a bare declarative that they merely suspect.

---

[25] A comparison with Moran (2005, 2018) is instructive. Moran's view is that, in asserting p, one takes responsibility for p's truth. My view is that, in uttering a declarative with content p, speakers signal that they occupy a particular position towards p, and speakers take responsibility for whatever position is signaled. My view is different in three respects. First, testimony does not require the speech act of assertion. Second, speakers are not responsible for p's truth but for the position they signal. Third, responsibility is incurred because of the positional norm. These features of my view make it easier to extend to hedged testimony. Thanks to a referee for encouraging me to make this comparison.

A related normative dimension of testimony is the responsibility acquired by the speaker for a hearer's doxastic response. Suppose *A* testifies *p* on the basis of perceptual experience. Then *A* alone is responsible for addressing challenges to whether *p* is known. But if *A* testifies *p* on the basis of testimony from *B*, then *A* is permitted to defer some challenges back to *B* (Goldberg, 2006; McMyler, 2011; Baker and Clark, 2018). *B*'s responsibility to answer challenges extends to hearers who responded to *A* testimony.

Following Goldberg (2011), both dimensions of a speaker's responsibility find a ready explanation if a speaker's testimony is governed by an epistemic norm. I will illustrate with the κ-norm. Start first with blameworthiness. A speaker who violates a norm is blameworthy. Social norms are brought about by expectations about how people will and should behave, and a violation thwarts those expectations— "It is just not how we do things around here." So a speaker who signals that they know but does not, violates the κ-norm. Blameworthiness results. Exigencies of the situation may excuse the speaker, but their becoming blameworthy is the typical consequence of violating the κ-norm.

Consider next the responsibility the speaker takes for the hearer's doxastic response, or what McMyler (2011) calls the EPISTEMIC RIGHT OF DEFERRAL. Speakers have a choice in what position to signal. Part of taking responsibility for the position signaled is being able to justify or explain why that position as opposed to a different one was signaled. Positions differ in what it takes for them to be warranted. Stronger ones demand more, weaker ones demand less. As a result, responsibility for a position requires being able to justify or explain how the signaled position is warranted. Since hearers without relevant defeaters can presume that speakers are complying with the κ-norm given that it is operative as a social norm, they can presume that speakers are warranted. The right of deferral results. The hearer is permitted to defer challenge because they can presume that the speaker has whatever reasons or "epistemic goods", to use Goldberg's phrase, are needed to warrant the epistemic position signaled.

We can conclude that a norm-based approach can explain some of the relevant normativity that surrounds testimony. In particular, it can explain how a speaker is blameworthy and the hearer's epistemic right of deferral with appeal to the same norm(s) that enable doxastic responses to be warranted. Speakers are responsible for their actions, and position signaling is an act performed in offering testimony. Since signaling is also a norm-governed action, being norm-governed deepens the responsibility.

5.2 Hedging and responsibility

To extend the scope of norm-based explanation of the warrant component to hedged testimony, the ᴘ-norm was introduced in §4.2. That same norm enables us to extend the scope of a norm-based explanation of the responsibility component. The ᴘ-norm requires speakers to occupy whatever position they signaled. Accordingly, the

norm applies to hedged and unhedged testimony alike because the only difference between them is what gets signaled. Norm-based explanations of blameworthiness and the epistemic right of deferral for unhedged testimony straightforwardly extend to hedged testimony too.

To illustrate how the explanations extend, suppose a speaker utters (15b). In doing so, they change what position is signaled. The override the default knowledge signal to indicate that the corgi splooting is merely probable given their evidence.

(15)   (a) Which dog splooted?
       (b) Probably, the corgi splooted.
       (c) # How do you know that?
       (d) Why / how is that probable for you?

By changing what position is signaled, the speaker changes what challenges are felicitous (van Elswyk, 2021). The knowledge-presuming challenge in (15c) becomes infelicitous. A more appropriate challenge is (15d). The change in challenge appropriateness exhibited between (15c) and (15d) broadly reflects the change in responsibility that results from hedging.

The p-norm changes its edicts based on what position is signaled. Applied to (15b), the p-norm requires that the corgi splooting be probable for the speaker. A speaker violates the norm in this instance if they occupy any position on which the splooting is not probable. For example, if it is a mere possibility that the corgi splooted or the speaker knows that the corgi did not sploot, the speaker violates the norm. Since a speaker is blameworthy if they violate a norm, then—barring any excusing exigencies that get the speaker off the hook—blameworthiness results if splooting is not probable for them.

Consider next the epistemic right of deferral. In response to (15b), suppose a hearer adjusts their doxastic state so that the corgi splooting is probable. If they offer *Probably, the corgi splooted* as hedged testimony to a third-party, the third-party may ask (15d) as a challenge. The hearer might then defer back to the speaker from whom they received the hedged testimony. That appropriateness of deferral is explained by the hearer's presumption that the speaker is complying with the p-norm. Since the speaker presumably complied with the norm, they are warranted in having p be probable for them. The hearer can defer back to the speaker who is positioned to meet those challenges.

Importantly, the right of deferral abides by a generalization akin to the wt-generalization (§4.3). What doxastic responses can be deferred back to the speaker depends on what position the speaker initially signaled. Or, stated as a generalization:

DEFERRAL THRESHOLD GENERALIZATION (DT-generalization)
For all acts of testimony $\mathcal{T}$ carrying the signal that the speaker occupies epistemic position $\mathcal{E}$, $\mathcal{T}$ yields the epistemic right of deferral for a doxastic response that is equal to or less than $\mathcal{E}$.

To illustrate, suppose a hearer believes that the corgi splooted in response to (12b) or *It is possible the corgi splooted*, and testifies to a third-party without hedging that the corgi splooted. Then the hearer is not permitted to defer subsequent challenges to whether they know as much back to the speaker. They are fully responsible for justifying or explaining how they know. Consequently, the speaker's responsibility does not extend to any doxastic response to their testimony made by a hearer. It extends only to doxastic responses that result in a position that is equal to or weaker in strength than what the speaker signaled.

Once extended to hedged testimony, the norm-based explanation of the epistemic right of deferral explains the DT-generalization too. The epistemic right of deferral results from presuming that a speaker is compliant with an epistemic norm. Compliant speakers have the epistemic goods to satisfy the norm. So hearers can defer third-party challenges back to the speaker on the presumption that speakers have the goods to meet the challenges. But the P-norm issues different requirements depending on what position was signaled by a speaker. Presumed compliance only entitles the further presumption that the speaker is warranted in occupying the position signaled. As a result, the epistemic right of deferral is only yielded for doxastic responses that are equal to or less than what position the speaker signaled. Stronger doxastic responses are not ones the hearer can presume the speaker has the epistemic goods for. They are not presumably positioned to meet third-party challenges to a stronger position.

We are now in a position to understand the way in which hedging diminishes a speaker's responsibility. The use of a declarative is accompanied by a signal that indicates the speaker's epistemic position towards the content of the declarative. Hedging changes what position is signaled. Given the P-norm, speakers can change what they are responsible for by changing what position is signaled. Since epistemic positions order in strength, speakers can signal weaker positions to incur responsibility that is weaker in that it is less demanding to fulfill.

6 Strategic hedging

The E-function hypothesized earlier is alerting hearers to positional risk (§3). But plainly speakers do not always use hedging for this purpose. They use hedging for a variety of purposes, or what Millikan dubbed speaker functions. Some might regard this discrepancy between what I have proposed and how hedging is practiced as a problem. I see the discrepancy as evidence of strategic uses of hedging. To motivate this perspective, this section focuses on what strategic communication is permitted by the P-norm. Conversations involve strategic use of language in mixed-motive situations between the speaker and hearer (Asher and Lascarides, 2013). Likewise, strategic hedging happens when it is in the speaker's interest to hedge, but it is in the hearer's interest for the speaker to not hedge or to hedge differently than how they did.

To appreciate how strategic hedging can be, revisit what the P-norm prohibits. It

prohibits what I called anti-representation and over-representation. However, the norm permits under-representation. Under-representation is allowed because the p-norm does not mandate what position can be signaled in cases when multiple positions are warranted. It require only that the speaker is warranted in occupying the position signaled. Consider a speaker who knows p. Since knowing p entails thinking p, a speaker is warranted in thinking p if they know p. Were the speaker to offer the hedged testimony *The corgi splooted, I think*, it would comply with the p-norm. They occupy the position signaled.

In §3, alternative proposals for hedging's etiological function were considered. Both alternatives constitute strategic uses of hedging. To illustrate, imagine a speaker who finds themselves in a conversation where most statements are contentiously challenged. In deciding how to contribute, the speaker will be motivated to hedge their statements. One motivation is to diminish their responsibility. By diminishing responsibility, they will be blameworthy for less if they cannot answer the challenges. So speakers can under-represent to diminish the responsibility required by the p-norm. Reputation management is another motivation. If a speaker anticipates a loss to their reputation by not being able to answer the challenges that would accompany unhedged testimony, hedging enables them to contribute in a manner that licenses weaker challenges and challenges that, if not answered, induce a less significant hit to their reputation.

Politeness is another motivation for strategic hedging. Broadly, the point of being polite is conflict avoidance (Kasper, 2000). Speakers who offend hearers risk starting a conflict, and politeness provides a means for harmonious conversation. Speakers can therefore under-represent their position to be polite. Brown and Levinson (1987) identify a number of politeness-oriented motivations for hedging. To illustrate one, imagine a hearer who is bristles when receiving testimony about political matters that are perceived as liberal. And yet, this hearer has just posed a political question to the speaker. To avoid conflict, the speaker might under-represent. They are warranted in occupying a stronger position with respect to the testimony offered, but they signal a weaker position with the hopes that their hesitation will be taken as less offensive.

In all of these examples, the speaker's decision to strategically hedge comes with little cost.[26] They dodge the costs associated with violating the p-norm, and, even if under-representing were prohibited by some other norm, it is very hard to be caught under-representing.[27] In the majority of cases, speakers can credibly defend

---

[26] I am assuming here that strategic hedging is purely strategic, i.e., the hedging is not owed to unjust social forces which incentivize the speaker—in virtue of their social identity—to hedge to avoid discriminatory harm. In these situations, hedging does come with clear costs. The speaker is unjustly barred from fully participating in the conversation as a knower. Such hedging counts as the kind of silencing that Dotson (2011) calls TESTIMONIAL SMOTHERING. Smothering happens when speakers are wrongly coerced to truncate their own testimony.

[27] Though under-representation is permitted by the p-norm, it does not follow that it is not governed by other norms. Hedging plausibly generates a kind of scalar implicature known as a CLAUSAL IMPLICATURE (Gazdar, 1979; Levinson, 1983). Accordingly, a speaker who signals a position $\mathcal{E}$ implicates that they do not occupy any positions stronger than $\mathcal{E}$. Speakers who under-represent

their choice to hedge by appealing to its etiological function. They were hesitant or unsure about what they testified, and they wanted to be transparent about their doubts for the benefit of the hearer.

The same cannot be said for the hearer. When speakers hedge strategically, the hearer is deprived of having a more accurate doxastic attitude warranted. This is a consequence of the norm-based approach to warrant once it is extended to hedged testimony. Since a hearer is warranted only in adopting attitudes equal to or weaker than what was signaled (§4), speakers who signal a weaker position warrant a weaker position for the hearer. They could have warranted a stronger attitude with their testimony but did not. The speaker's strategic gain is the hearer's epistemic loss.

7   Conclusion

This essay made room for hedging in theorizing about testimony by developing and extending a norm-based approach. On a norm-based approach, social norms play a pivotal role: they explain how doxastic responses to testimony are prima facie warranted for a hearer, and how speakers incur responsibility to hearers by offering testimony. Hedged and unhedged testimony were uniformly accounted for as actions governed by the same norm.

Since this essay focused on making room for hedging, attention was not given to whether alternative approaches to testimony can explain hedged testimony as well as a norm-based approach. It remains to be seen whether reductionist or non-reductionist approaches that do not give social norms a starring role are equally as explanatory. Even still, the ease with which a norm-based anti-reductionism can be extended to hedged testimony sets a benchmark for theories of testimony. If it turns out that future attempts to explain hedged testimony with alternative approaches are less straightforward, we may find a new reason to prefer a norm-based approach.

The focus of this essay was also incomplete by only considering the warrant and responsibility components of testimony. The normative pressure for a hearer to doxastically respond to a speaker was not considered, for example. Like the other components of testimony, contemporary discussions have dwelled on the pressure to believe a speaker's unhedged testimony.[28] But the pressure to believe a speaker will not result from every instance of hedged testimony (§4-§5). The pressure will be for a weaker doxastic response. A complete accounting of how hedging makes a difference to the normativity surrounding testimony will need to detail how hedging influences or alters the strength of this pressure.

---

therefore mislead hearers about the strength of their position. In the present case, the speaker misleads by implicating that they do not know. Such misleading is like other instances of misleading with scalar implicatures. To the extent that such misleading is inappropriate, under-representing one's position is.

[28] See Hinchman (2005), Moran (2005), Fricker (2007), Dotson (2011), and Goldberg (2020) for different approaches to this component of testimony.

Despite the limits in scope, this essay clarifies the nature of testimony generally by focusing on hedged testimony in particular. A key takeaway is that what a hearer gets from testimony depends on how a speaker offers them testimony. When covering how hedging influences the warrant and responsibility components of testimony, we encountered two related generalizations: the wt-generalization and the dt-generalization. These generalizations covary. What doxastic response is warranted for a hearer is typically what they have the epistemic right of deferral for. Both components of testimony depend on what epistemic position the speaker chose to signal. Attending only to unhedged testimony has obscured the extent of this dependence on what the speaker signaled. I believe it is hard to overlook once we consider hedging.[29]

References

Jonathan Adler. *Belief's Own Ethics*. MIT Press, 2002.

Scott AnderBois. On the exceptional status of reportative evidentials. *Proceedings of SALT*, 24:234–254, 2004. doi: https://doi.org/10.3765/salt.v24i0.2424.

Nicholas Asher and Alex Lascarides. Strategic conversation. *Semantics and Pragmatics*, 6(2):1–62, 2013. doi: http://dx.doi.org/10.3765/sp.6.2.

Kent Bach. Applying pragmatics to epistemology. *Philosophical Issues*, 18:68–88, 2008. doi: https://doi.org/10.1111/j.1533-6077.2008.00138.x.

Judith Baker and Philip Clark. Epistemic buck-passing and the interpersonal view of testimony. *Canadian Journal of Philosophy*, 48(2):178–199, 2018. doi: https://doi.org/10.1080/00455091.2017.1341781.

Mark Bedau. Can biological teleology be naturalized? *Journal of Philosophy*, 88:647–655, 1991. doi: https://doi.org/10.5840/jphil1991881111.

Mathew Benton and Peter van Elswyk. Hedged assertion. In Sandy Goldberg, editor, *Oxford Handbook of Assertion*. Oxford University Press, 2020.

Matthew Benton. Lying, belief, and knowledge. In Jorg Meibauer, editor, *Oxford Handbook of Lying*. Oxford University Press, 2018.

Matthew Benton. Knowledge is the norm of assertion. In B. Roeber, J. Turri, M. Steup, and E. Sosa, editors, *Contemporary Debates in Epistemology*. Blackwell, 3rd edition, forthcoming.

Cristina Bicchieri. *The Grammar of Society: The Nature and dynamics of social norms*. Cambridge University Press, 2006.

Penelope Brown and Stephen Levinson. *Politeness: Some Universals in Language Usage*. Cambridge University Press, 1987.

David Buller. Etiological theories of function: A geographical survey. *Biology and Philosophy*, 13:505–527, 1998. doi: https://doi.org/10.1023/A:1006559512367.

Tyler Burge. Content preservation. *The Philosophical Review*, 102(4):457–488, 1993. doi: https://doi.org/10.2307/1523046.

Ivano Ciardelli, Jeroen Groenendijk, and Floris Roelofsen. Attention! *Might* in inquisitive semantics. *Proceedings of SALT*, 19:91–108, 2011. doi: https://doi.org/10.3765/salt.v19i0.2520.

C.A.J. Coady. *Testimony: A Philosophical Study*. Oxford University Press, 1992.

Donald Davidson. Deception and division. In J. Elster, editor, *The Multiple Self*, pages 79–92. Cambridge University Press, 1985.

Richard Dawkins and John Richard Krebs. Animal signals: Information or manipulation? In J.R. Krebs and N.B. Davies, editors, *Behavioural ecology: An evolutionary approach*, pages 282–309. Oxford: Basil Blackwell, 1978.

Lorenz Demey. Contemporary epistemic logic and the lockean thesis. *Foundations of Science*, 18(4):599–610, 2013. doi: https://doi.org/10.1007/s10699-012-9292-9.

Kevin Dorst. Lockeans maximize expected accuracy. *Mind*, 128(509):175–211, 2019. doi: https://doi.org/10.1093/mind/fzx028.

Kevin Dorst and Matthew Mandelkern. Good guesses. *Philosophy and Phenomenological Research*, forthcoming. doi: https://doi.org/10.1111/phpr.12831.

Kristie Dotson. Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2):236–257, 2011. doi: https://doi.org/10.1111/j.1527-2001.2011.01177.x.

Don Fallis. Davidson was almost right about lying. *Australasian Journal of Philosophy*, 91(2):337–353, 2013. doi: https://doi.org/10.1080/00048402.2012.688980.

Paul Faulkner. *Knowledge on Trust*. Oxford University Press, 2011.

W. Tecumseh Fitch. Evolving honest communication systems: Kin selection and mother tongues. In D.K. Oller and U. Griebel, editors, *The Evolution of Communication Systems: A Comparative Approach*, pages 275–296. MIT Press, 2004.

Bruce Fraser. Pragmatic competence: The case of hedging. In Gunter Kaltenb ock, Wiltrud Mihatsch, and Stefan Schneider, editors, *New Approaches to Hedging*, pages 15–34. Emerald, 2010.

Miranda Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.

Manuel Garcia-Carpintero. Assertion and the semantics of force-markers. In C. Bianchi, editor, *The semantics/pragmatics distinction*. CSLI, 2004.

Gerald Gazdar. *Pragmatics: Implicature, presupposition, and logical form*. Academic Press, 1979.

Herbert Gintis. *The bounds of reason*. Princeton University Press, 2009.

Sanford Goldberg. Reductionism and the distinctiveness of testimonial knowledge. In Jennifer Lackey and Ernest Sosa, editors, *The epistemology of testimony*. Oxford University Press, 2006.

Sanford Goldberg. Assertion, testimony, and the epistemic significance of speech. *Logos & Episteme*, 1(1):59–65, 2010. doi: https://doi.org/10.5840/logos-episteme20101121.

Sanford Goldberg. Putting the norm of assertion to work: The case of testimony. In Jessica Brown and Herman Cappelen, editors, *Assertion: New Philosophical Essays*, pages 175–196. Oxford University Press, 2011.

Sanford Goldberg. *Conversational pressure*. Oxford University Press, 2020.

Peter Graham. Testimonial entitlement and the function of comprehension. In Alan Millar and Adrian Haddock, editors, *Social Epistemology*. Oxford University Press, 2010.

Peter Graham. Testimony, trust, and social norms. *Abstracta*, 6(3):92–116, 2012.

Peter Graham. Functions, warrant, history. In Abrol Fairweather and Owen Flanagan, editors, *Naturalizing epistemic virtue*, pages 15–35. Cambridge University Press, 2014.

Peter Graham. Epistemic normativity and social norms. In David Henderson and John Greco, editors, *Epistemic evaluations: purposeful epistemology*, pages 243–273. Oxford University Press, 2015.

Peter Graham. The function of assertion and social norms. In Sanford Goldberg, editor, *Oxford Handbook of Assertion*, pages 727–748. Oxford University Press, 2020a.

Peter Graham. Assertions, handicaps, and social norms. *Episteme*, 17(3):349–363, 2020b. doi: https://doi.org/10.1017/epi.2019.53.

Daniel Greco. How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29:179–201, 2015. doi: https://doi.org/10.1111/phpe.12059.

John Greco. What is transmission? *Episteme*, 13(4):481–498, 2016. doi: https://doi.org/10.1017/epi.2016.25.

John Greco. *The transmission of knowledge*. Cambridge: Cambridge University Press, 2020.

W.D. Hamilton. The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7:1–52, 1964. doi: https://doi.org/10.1016/0022-5193(64)90038-4.

John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173:1393–1404, 2016. doi: https://doi.org/10.1007/s11098-015-0553-7.

Edward Hinchman. Telling as inviting to trust. *Philosophy and Phenomenological Research*, 70(3):562–587, 2005. doi: https://doi.org/10.1111/j.1933-1592.2005.tb00415.x.

Edward Hinchman. Assertion and testimony. In Sanford Goldberg, editor, *Oxford Handbook of Assertion*, pages 555—579. Oxford University Press, 2020.

Ben Holguín. Lying and knowing. *Synthese*, 198(6):5351–5371, 2021. doi: https://doi.org/10.1007/s11229-019-02407-2.

Ben Holguín. Thinking, guessing, believing. *Philosopher's Imprint*, forthcoming–. URL https://philpapers.org/archive/HOLTGA-2.pdf.

Sophie Horowitz. Accuracy and educated guesses. *Oxford Studies in Epistemology*, 6:85–113, 2019. doi: https://doi.org/10.1093/oso/9780198833314.001.0001.

Lawrence Hosman. The evaluative consequences of hedges, hesitations, and intensifiers. *Human Communication Research*, 15(3):383–406, 1989. doi: https://doi.org/10.1111/j.1468-2958.1989.tb00190.x.

Elizabeth Jackson. Belief, credence, and evidence. *Synthese*, 197(11):5073–5092, 2020. doi: https://doi.org/10.1007/s11229-018-01965-1.

James M. Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65 (4):575–603, 1998. doi: https://doi.org/10.1086/392661.

James M. Joyce. Accuracy and coherence: prospets for an althetic epistemology of partial belief. In F. Huber and C. Schmidt-Petri, editors, *Degrees of belief*. Springer, 2009.

Daniel Kahneman and Amos Tversky. Variants of uncertainty. *Cognition*, 11(2):143–157, 1982. doi: https://doi.org/10.1016/0010-0277(82)90023-3.

Gabriele Kasper. Linguistic politeness: current research issues. *Journal of Pragmatics*, 14(2):192–218, 2000. doi: https://doi.org/10.1016/0378-2166(90)90080-W.

Christopher Kelp. Assertion: A function first account. *Nous*, 52(2):411–442, 2018. doi: https://doi.org/10.1111/nous.12153.

Angelika Kratzer. What *Must* and *Can* must and can mean. *Linguistics and Philosophy*, 1:337–355, 1977. doi: https://doi.org/10.1007/BF00353453.

Jennifer Lackey. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press, 2008.

Hannes Leitgeb. The stability theory of belief. *The Philosophical Review*, 123(2):131–171, 2014. doi: https://doi.org/10.1215/00318108-2400575.

Stephen Levinson. *Pragmatics*. Cambridge University Press, 1983.

John MacFarlane. Knowledge laundering: testimony and sensitive invariantism. *Analysis*, 65(132–138), 2005. doi: 10.1111/j.1467-8284.2005.00538.x.

Jim Mackenzie. I guess. *Australasian Journal of Philosophy*, 65(3):290–300, 1987. doi: https://doi.org/10.1080/00048408712342951.

Elin McCready. *Reliability in Pragmatics*. Oxford University Press, 2015.

Peter McLaughlin. *What functions explain*. Cambridge University Press, 2012.

Benjamin McMyler. *Testimony, Trust, and Authority*. Oxford University Press, 2011.

Marie-Christine Meyer. *Ignorance and grammar*. PhD thesis, Massachusetts Institute of Technology, 2013. URL https://dspace.mit.edu/handle/1721.1/84420.

Manfred Milinski, Dirk Semmann, and Häns-Jurgen Krambeck. Reputation helps solve the 'tragedy of the commons'. *Nature*, 415:424–426, 2002. doi: https://doi.org/10.1038/415424a.

Ruth Millikan. *Language, Thought, and Other Biological Categories*. MIT Press, 1987.

Ruth Millikan. *Varieties of Meaning*. MIT Press, 2004.

Richard Moran. Getting told and being believed. *Philosophers' Imprint*, 5(5):1–29, 2005. URL http://hdl.handle.net/2027/spo.3521354.0005.005.

Richard Moran. *The exchange of words: Speech, testimony, and intersubjectivity*. Oxford University Press, 2018.

Sarah Moss. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics*, 8(5):1–81, 2015. doi: http://dx.doi.org/10.3765/sp.8.5.

Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, 2018.

Sarah Murray and Will Starr. Force and conversational states. In Daniel Fogal, Daniel Harris, and Matt Moss, editors, *New Work On Speech Acts*. Oxford University Press, 2018.

Martin Nowak and Karl Sigmund. Evolution of indirect reciprocity. *Nature*, 437:1291–1298, 2005. doi: https://doi.org/10.1038/nature04131.

David Owens. Testimony and assertion. *Philosophical Studies*, 130:105–129, 2006. doi: https://doi.org/10.1007/s11098-005-3237-x.

Andrew Peet. Testimony and the epistemic uncertainty of interpretation. *Philosophical Studies*, 173(2):395–416, 2016. doi: https://doi.org/10.1007/s11098-015-0498-x.

Andrew Peet. Etiology, understanding, and testimonial belief. *Synthese*, 195:1547–1567, 2018. doi: https://doi.org/10.1007/s11229-016-1281-z.

Timothy Perrine. In defense of non-reductionism in the epistemology of testimony. *Synthese*, 191(14):3227–3237, 2014. doi: https://doi.org/10.1007/s11229-014-0443-0.

Richard Pettigrew. *Accuracy and the laws of credence*. Oxford: Oxford University Press, 2016.

John Pollock and Joseph Cruz. *Contemporary theoreis of knowledge*. Rowman and Littlefield, 1999.

Joel Predd, Robert Seiringer, Elliott Lieb, Daniel Osherson, H. Vincent Poor, and Sanjeev Kulkarni. Probabilistic coherence and proper scoring rules. *IEEE Transactions on Information Theory*, 55(10):4786–4792, 2009. doi: https://doi.org/10.1109/TIT.2009.2027573.

Ellen F. Prince, Joel Frader, and Charles Bosk. On hedging in physician-physician discourse. In R. D. Pietro, editor, *Linguistics and the professions*, pages 83–97. Ablex: Norwood, NJ, 1982.

Duncan Pritchard. The epistemology of testimony. *Philosophical Issues*, 326–348, 2004. doi: http://doi.org/10.1111/j.1533-6077.2004.00033.x.

Anna Prokofieva and Julia Hirschberg. Hedging and speaker commitment. In *5th International Workshop on Emotion, Social Signals, Sentiment, and Linked Open Data*, pages 10–13, 2014. URL http://www.cs.columbia.edu/~prokofieva/Prokofieva_Hirschberg_Paper.pdf.

David Reby and Karen McComb. Anatomical constraints generate honesty: acoustie cues to age and weight in the roars of red deer stags. *Animal Behaviour*, 65:519–530, 2003. doi: https://doi.org/10.1006/anbe.2003.2078.

Stephen Reynolds. Testimony, knowledge, and epistemic goals. *Philosophical Studies*, 110:139–161, 2002. doi: https://doi.org/10.1023/A:1020254327114.

Daniel Rothschild. What it takes to believe. *Philosophical Studies*, 177:1345–1362, 2020. doi: https://doi.org/10.1007/s11098-019-01256-6.

Paul Rozin, L. Lowery, S. Imada, and Jonathan Haidt. The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality and Social Psychology*, 76:574–586., 1999. doi: http://doi.org/10.1037//0022-3514.76.4.574.

Sunita Sah, Don Moore, and Robert MacCoun. Cheap talk and credibility: The consequences of confidence and accuracy on advisor credibility and persuasiveness. *Organizational Behavior and Human Decision Processes*, 121:246–255, 2013. doi: https://doi.org/10.1016/j.obhdp.2013.02.001.

Jennifer Saul. *Lying, misleading, and what is said: An exploration in philosophy of language and in ethics*. Oxford: Oxford University Press, 2012.

Sebastian Schuster and Judith Degen. I know what you're probably going to say: Listener adaption to variable use of uncertainty expressions. *Cognition*, 203(104285):1–31, 2020. doi: https://doi.org/10.1016/j.cognition.2020.104285.

Thomas C. Scott-Phillips. Defining biological communication. *Journal of Evolutionary Biology*, 21(2):387–395, 2008. doi: https://doi.org/10.1111/j.1420-9101.2007.01497.x.

Thomas C. Scott-Phillips. Evolutionary stable communication and pragmatics. In Anton Benz, Christian Ebert, Gerhard Jager, and Robert van Rooij, editors, *Language, Games, and Evolution*, pages 117–133. Springer, 2010.

Mona Simion. Testimonial contractarianism. *Nous*, 44(4):891–916, 2021. doi: https://doi.org/10.1111/nous.12337.

Mona Simion and Christoph Kelp. How to be an anti-reductionist. *Synthese*, 197(7):2849–2866, 2020. doi: https://doi.org/10.1007/s11229-018-1722-y.

John Maynard Smith. Group selection and kin selection. *Nature*, 201:1145–1147, 1964. doi: https://doi.org/10.1038/2011145a0.

John Maynard Smith and David Harper. *Animal signals*. Oxford University Press, 2003.

Dan Sperber. An evolutionary perspective on testimony and argumentation. *Philosophical Topics*, 29(1/2):401–413, 2001. doi: https://doi.org/10.5840/philtopics2001291/215.

Chandra Sekhar Sripada and Stephen Stich. A framework for the psychology of norms. In P. Carruthers, S. Laurence, and S. Stitch, editors, *The innate mind*, volume 2 of *Culture and Cognition*, pages 237–256. Oxford University Press, 2006.

Andreas Stokke. Lying and asserting. *Journal of Philosophy*, 110(1):33–60, 2013. doi: https://doi.org/10.5840/jphil2013110144.

Andreas Stokke. *Lying and Insincerity*. Oxford University Press, 2018.

Scott Sturgeon. Reason and the grain of belief. *Nous*, 42(1):139–165, 2008. doi: http://doi.org/10.1111/j.1468-0068.2007.00676.x.

Elizabeth Tenney, Robert MacCoun, Barbara Spellman, and Reid Hastie. Calibration trumps confidence as a basis for witness credibility. *Psychological Science*, 18(1):46–50, 2007. doi: https://doi.org/10.1111/j.1467-9280.2007.01847.x.

Elizabeth Tenney, Barbara Spellman, and Robert MacCoun. The benefits of knowing what you know (and what you don't): How calibration affects credibility. *Journal of Experimental Social Psychology*, 44(5):1368–1375, 2008. doi: https://doi.org/10.1016/j.jesp.2008.04.006.

Elizabeth Tenney, Nathan Meikle, David Hunsaker, Don Moore, and Cameron Anderson. Is overconfidence a social liability? the effect of verbal versus nonverbal expressions of confidence. *Journal of Personality and Social Psychology*, 116(3):396–415, 2019. doi: https://doi.org/10.1037/pspi0000150.

Michael Tomasello. *Becoming Human: A theory of ontogeny*. Harvard University Press, 2019.

Stephen Toulmin. *The uses of argument*. Cambridge University Press, 1958.

John Turri. Prompting challenges. *Analysis*, 70:456–462, 2010. doi: https://doi.org/10.1093/analys/anq027.

John Turri. Experimental work on the norms of assertion. *Philosophy Compass*, 12(7):e12425, 2017. doi: https://doi.org/10.1111/phc3.12425.

John Turri. Revisiting norms of assertion. *Cognition*, 177:8–11, 2018. doi: http://doi.org/10.1016/j.cognition.2018.03.023.

John Turri. Truth, fallibility, and justification: New studies in the norms of assertion. *Synthese*, 198:8073–8084, 2021. doi: https://doi.org/10.1007/s11229-020-02558-7.

Peter Unger. *Ignorance: The Case for Skepticism*. Oxford: Clarendon Press, 1975.

Peter van Elswyk. Representing knowledge. *The Philosophical Review*, 130(1):97–143, 2021. doi: https://doi.org/10.1215/00318108-8699695.

Peter van Elswyk. Hedging. Manuscript, 2022.

Peter van Elswyk and Matthew Benton. Assertion remains strong. Manuscript, 2022.

Peter van Elswyk and Yasha Sapir. Hedging and the ignorance norm on inquiry. *Synthese*, 199(3-4):5837–5859, 2021. doi: https://doi.org/10.1007/s11229-021-03048-0.

Colin Vullioud, Fabrice Clément, Thom Scott-Phillips, and Hugo Mercier. Confidence as an expression of commitment: why misplaced expressions of confidence backfire. *Evolution and Human Behavior*, 38:9–17, 2017. doi: https://doi.org/10.1016/j.evolhumbehav.2016.06.002.

Claus Wedekind and Manfred Milinski. Cooperation through image scoring in humans. *Science*, 288(5467):850–852, 2000. doi: https://doi.org/10.1126/science.288.5467.850.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, 2000.

Ludwig Wittgenstein. *Philosophical Investigations*. Oxford: Blackwell, 1953.

Larry Wright. Functions. *The Philosophical Review*, 82(139–168), 1973. doi: https://doi.org/10.2307/2183766.

Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–937, 2010. doi: https://doi.org/10.1111/j.1747-9991.2010.00360.x.

Kevin Zollman. Modeling the social consequences of testimonial norms. *Philosophical Studies*, 172(9):2371–2383, 2015. doi: https://doi.org/10.1007/s11098-014-0416-7.