

The Pragmatics of Hedging: Constructing Uncertainty in English Discourse

Mahnoor Ali*

Master's Scholar, Department of Applied Linguistics, Government College University, Faisalabad, Pakistan Email: mmaahhnoooooorali@gmail.com

Narjis Batool

Master's Scholar, Department of Applied Linguistics, Government College University, Faisalabad, Pakistan Email: narjisbatool078@gmail.com

Syed Irfanullah Shah

Master's Scholar, Department of Applied Linguistics, Government College University, Faisalabad, Pakistan Email: syedirfan4429@gmail.com

ABSTRACT

Uncertainty in English discourse is rarely left to chance; speakers and writers construct it deliberately through a repertoire of linguistic devices known as hedges. This article offers a pragmatic analysis of hedging strategies as they operate across three contemporary discourse domains: academic writing, political communication, and digital or AI-mediated text. Drawing on the foundational taxonomies of Lakoff and Hyland alongside more recent corpus and diachronic scholarship, the analysis distinguishes hedges by their scope of modification, separating devices that qualify propositional content from those that manage the relationship between speaker, claim, and hearer. Illustrative examples drawn from published academic prose, political interview transcripts, and AI-generated text are examined to show how epistemic forms such as modal auxiliaries and evidential shields are routinely redeployed for interpersonal and strategic ends, a phenomenon this article terms functional drift. The analysis further proposes a dynamic, feedback-sensitive framework in which a hearer's inferential uptake of a hedge shapes the speaker's subsequent hedging choices, treating uncertainty construction as an interactional process rather than a fixed property of individual words. The discussion considers how genre conventions, professional risk, and the emergence of large language models are reshaping the pragmatics of hedging in English, and it closes by identifying open questions for corpus-based and experimental follow-up. The article contributes a genre-comparative synthesis that connects classical politeness-based accounts of hedging to its evolving role in twenty-first-century mediated communication.

Keywords: Hedging; Pragmatics; Epistemic Modality; Academic Discourse; Political Discourse; Vague Language; Stance

Introduction

Reduced certainty is a limited commodity in human discourse, and the English language has evolved a complex ensemble of tools for coping with it. A scientist who states that a result 'may imply' a causal relationship, a politician who claims a policy is 'broadly consistent' with previous commitments, and a computer program that begins an explanation with 'I think' are all engaging in the same fundamental pragmatic activity: they are hedging. Hedging is the conscious use of language to mitigate the force, precision, or commitment of an utterance, and comprises an area where the fields of semantics, pragmatics, and social interaction overlap (Hyland, 1998; Lakoff, 1973).

The term is credited to Lakoff (1973), who described sort of and technically as instances of words that alter propositional membership to be "fuzzier or less fuzzy". Lakoff's focus was on the semantic and logical functions of hedges, and how they modify truth-conditional category membership. Following the evolution of pragmatics and sociolinguistics research after the 1980s, this narrow scope has expanded into a resource for managing interpersonal credibility, authority, and epistemic commitment (Brown & Levinson, 1987; Hyland, 1996; Prince et al., 1982).

The paper has asked a comparative question: how does the pragmatic function of hedges vary in the discourse domains (including political discourse) that constitute the major genres of English language use at the moment, academic, political, and more recently digital & AI-mediated? Rather than accept presence in a fixed inventory of hedge forms as the end point of an analysis, it has opted to emphasize variation in scope of modification, drawing a distinction between devices that hedge reference (i.e., modify the informational content of an utterance) and those that hedge the speaker-hearer relationship (following a well-established content-relation distinction in the literature on hedging and vague language, e.g., Channell, 1994; Fraser, 2010).

The article continues in five more parts. In Section 2, a brief theoretical review of hedging studies is offered, from Lakoff's fuzzy-set semantics to Brown and Levinson's politeness framework, to Hyland's functional taxonomy of academic hedges. Section 3 describes the analytical framework in which the sample of exemplary material is selected and interpreted. Section 4 approaches the heart of the pragmatic analysis by describing the devices used, their hierarchical organization into six device types, and the notion of functional drift by which epistemically marked items are recruited into the interpersonal/interventional functions. In Section 5, we compare the three written discourse types, and in Section 6, we draw out the conclusions.

But the implications of this comparison are not merely theoretical. The choices to hedge a forecast, draw conclusions, or make commitments can be consequential in, say, a professional or political context: how a doctor hedges a prognosis will influence how her patient perceives her, while how a scientist hedges her conclusion will influence the trust investors place in her hypothesis to guide their policy. As large language models take on new tasks such as drafting, summarizing, and negotiating, how hedging is invoked and interpreted in an LLM-based context is becoming a substantive issue that pragmatic frameworks, which were tailored for human speakers with authentic epistemic states and face to lose, may not be equipped to address.

Literature Review

From fuzzy sets to pragmatic function

Lakoff's original proposal treated hedges as operators on category membership: sort of a fish makes the truth of X is a fish a matter of degree rather than a binary judgment (Lakoff, 1973). This logical framing explains the semantic behaviour of a narrow set of approximating expressions well, but it does not, by itself, account for why a writer would choose to blur category membership in a given context, nor why such blurring is so heavily concentrated in certain registers, such as scientific reporting, and comparatively rare in others, such as recipes or instructions. The pragmatic turn in hedging research, beginning in earnest with work on politeness and professional discourse, supplies the missing motivational account: speakers hedge not because the world is fuzzy but because managing commitment, face, and accountability is communicatively useful (Brown & Levinson, 1987; Prince et al., 1982).

Politeness, face, and hedging as facework

Brown and Levinson's politeness theory situates hedging within a broader account of how speakers protect their own and their interlocutors' face, the public self-image that every competent language user is assumed to want maintained (Brown & Levinson, 1987). On this account, a hedge such as I could be wrong "softens a potentially face-threatening assertion by lowering the speaker's apparent commitment, thereby reducing the risk to both parties if the assertion turns out to be incorrect or unwelcome. Prince et al. (1982) extended this reasoning to physician-physician discourse, showing that hedges in clinical talk frequently protect professional face and manage epistemic and ethical responsibility for diagnostic claims rather than performing straightforward politeness. Their distinction between hedges that mark the speaker's degree of commitment to propositional truth, which they term approximators, and shields that operate on the relationship between speaker and proposition, anticipates much of the content-versus-relation distinction used later in this article.

Hyland's functional taxonomy of academic hedges

Hyland (1996) and the fuller treatment in Hyland (1998) remain the most widely cited framework for hedging in written academic English. Hyland's corpus-based analysis of research articles across disciplines distinguishes content-oriented hedges, which qualify the truth value or precision of propositional content, such as modal verbs, epistemic adverbs, and approximators, from reader-oriented hedges, which manage the relationship between the writer and the audience by explicitly involving the reader's judgment, for example, it is generally accepted that or as far as we can tell. Hyland further separates hedges by their evidential basis, distinguishing writer-based hedges that signal the author's own uncertainty from attribution shields that displace epistemic responsibility onto other sources, such as reportedly or according to. Crompton (1997) raises a persistent methodological challenge for this taxonomy: many surface forms, particularly modal verbs such as May and could, are pragmatically ambiguous between epistemic and non-epistemic readings, and this ambiguity resists purely lexical or even syntactic disambiguation, requiring case-by-

case contextual judgment. This concern motivates the qualitative, example-driven method adopted in Section 3 of this article rather than a purely automated frequency count.

Vagueness, approximation, and scope of modification

Channell (1994)'s treatment of vagueness complements Hyland's taxonomy by foregrounding a class of hedges, including approximators such as round, approximate, and other categories of vague quantifiers such as and so on, or something, whose pragmatic function is in almost all cases primarily to show that the speaker had a more precise way of saying something but chose not to. Channell's account is important for our purpose because it distinguishes vagueness as a pragmatic device from epistemic uncertainty per se: it is not necessarily uncertain about the price but is not committing to the particular figure, for pragmatic reasons of economy and appropriateness. Fraser (2010), to bring out these two pragmatic distinctions, further advances the scope-and-direction hypothesis that handles all such scopal effects in terms of either modification of propositional content, the illocutionary force, or the "relationship" between sender and message. The content/relationship distinction, and Channell's areal distinction, form the axes for a conceptual map that is used in Section 4 of this article.

Epistemic modality and the diachronic trajectory of hedging

Since there are so many common modals in English, it is noteworthy that Coates (1987) established that epistemic modality in spoken discourse is primarily encoded by modals like may/might/could, which could encode the speaker's probability estimate of a proposition as well as serving as the grammatical underpinning of much hedging. More recent corpus work shows how this pool has evolved over time and – importantly – how it differs across disciplines. In comparing research articles from 1965, 1985 and 2015 from four disciplines, Hyland and Jiang (2016) find a genuine disciplinary divide rather than a universal trend: while hedging greatly decreased in the soft sciences during this interval (by about a third in applied linguistics and a fifth in sociology), it remained stable or increased slightly in the hard sciences (including biochemistry and electrical engineering).

Yao et al. (2023), in analysing a relatively smaller, more recent twenty-five-year time period in just one hard science venue – the journal *Science* – find a significant decrease in the density of hedging over this time period: this is not, as they contend, a signal of prudence, but rather an embedded promotional rhetoric in which diminished hedging is paired with increased boosting for each claim to be true.

Together these two findings show the disparities across disciplines, time frames, and specific publishing venues in what might appear to be a straightforward developmental process for hedging in academic writing. This article applies the same logic to the less-studied spheres of political and digital writing.

Register, gender, and pedagogical dimensions of hedging

Otherwise defined, hedging has been shown to systematically differ across genre, discipline, and even in various language pedagogical contexts: an early study of hedges and boosters in women's and men's speech by Holmes (1990) demonstrated

that hedging cannot be generalized with a single social meaning: I think can function in one domain as an epistemic qualifier and in the other as a positive politeness strategy inviting cooperation, and the two meanings are often conflated in earlier sociolinguistic studies that generalized all forms of tentativeness by treating hedging as a mono-pragmatic marker. This is important for the present article in that it reinstates the distinction between content and relation we argued for in Section 2.4: the social meaning and pragmatic function of a hedge cannot simply be predicted based on the lexical item it consists of, but rather they depend on the extent of scope of the marker and the discourse context in which it is uttered. Skelton (1988) adopts a language-pedagogical perspective in his study of hedging, claiming that English language learners (ELLs) need systematic study of hedging in ELT since the avoidance of, or inconsistent use of, hedges by ELLs may cause native readers to misperceive those learners as overconfident or evasive in contexts where neither impression is intended, a problem related to the nuanced interactional nature of hedging established in Section 2. This connects directly to the cross-linguistic and cross-disciplinary variation evidenced by Vold (2006), who compares epistemic modality markers among disciplines and languages and finds that the systems of markers and frequencies of their occurrences are sufficiently different in various discourse communities that using a hypothetical clause marker translation may unintentionally change the pragmatic force of the hedge, a potentially problematic issue in multilingual science publishing and translation.

Another register distinction is the audience for which a text is written. In their comparison of hedging in specialist scientific writing and a number of popularizations of science (science journalism, etc.), Varttala (2001) finds that popularizations tend to omit many of the hedge-types common in specialists' texts, substituting instead either a much more assertive flat statement or, by contrast, a much more cautious flat statement, either in accordance with a particular editorial convention. Vázquez Orta and Giner (2008) show, between specialist disciplines rather than between specialized and popular language, that the reliance on modal verbs versus suggest or indicate varies sufficiently between disciplines that a taxonomy calibrated on one discipline (notably, the applied linguistics corpora this paper relies on for its research on hedging) will not necessarily be applicable to others. These register effects, together with Markkanen and Schrader (1997)'s reminder that research on hedging crosses the borders of logic, pragmatics, sociolinguistics, and rhetoric, without consistent agreement in terms, scope, or approach, warn us not to treat hedging as a single phenomenon characterized by a single taxonomy. For the purpose of this article, the six-device taxonomy of scope introduced in Section 4 was kept explicitly tied to device function (scope of modification) rather than surface expression, in order to avoid this problem.

Analytical Method

This article employs a qualitative, example-driven pragmatic approach instead of an automated corpus-frequency analysis for the following reasons (anticipated in section 2.3): firstly, many hedging forms are lexically ambiguous, either epistemic or non-epistemic, and secondly, a hedge's interpersonal or strategic function is often only discernible in the discourse environment, thus identifying a hedge based solely on the

form of a potentially-hedging segment may lead to false positives such as counting deontic may, and false negatives, where the hedge is expressed via syntax or intonation rather than one individual lexical item (Crompton, 1997; Vold, 2006).

The illustrative material was thus purposively sampled from three domains where there are different kinds of institutional pressures driving changes in epistemic stance: published prose, both social and natural sciences, where the phenomenon of hedging has been most studied; recorded political discussions, including interview and debate transcripts where hedging interacts with issues of accountability and evasion; and computer-mediated written and spoken language, and computer-mediated (including AI mediated) language, where conventions are still emerging and research on the phenomenon is still in progress (Hu & Cao, 2011; Mo & Crosthwaite, 2025). Within each domain, examples were selected to exemplify the six types of device defined in Table 1, and the phenomenon of functional drift explained in Section 4.6; no claim is made on a quantitative corpus-level measure of the frequency of occurrence of any individual device or device type. Where a quantitative or corpus-based claim is made in this paper, it is located and not recalculated; its contribution to this article is the phenomenon itself, not new frequency data.

There are two practical limitations to this approach which are referred to again in Section 6. Firstly, purposive example selection cannot lead to conclusions concerning the frequency of hedge types across domains but only their qualitative distribution and functions. Secondly, and following from the last, the interpretation of the pragmatic functions of any hedge must, unavoidably, involve a degree of analyst judgment. This is especially the case for borderline modals such as could, which encodes epistemic possibility, social distance and politeness (Coates, 1987).

A Pragmatic Taxonomy of Hedging Devices

Having considered in Section 2 the theoretical distinctions between the types of hedging devices, in this section, the six types are arranged by their function in Table 1, then the dynamic framework in Figure 1, and the conceptual map in Figure 2.

Approximators

Finally, approximators quantify over the amount or category involved without necessarily raising the speaker's degree of speaker-oriented commitment to the fact, reflecting Channell (1994)'s version of the vague language hypothesis. For instance, using approximators like roughly, about, and -ish suffixation, as in "a couple dozen participants" or "the meeting is at three-ish." Approximators are content-oriented in Fraser's sense, affecting the propositional content itself and not the speaker's stance on that content, and thus are found at the broadest end of the scale from everyday conversation to early reporting of experimental data.

Epistemic modal markers

Epistemic modal markers express how confident the speaker is about the truth of a proposition. Usually, they are in the form of modal verbs, e.g., may, might and could, or epistemic adverbs, e.g., possibly and perhaps (Coates, 1987). These examples function in a content-oriented and narrowly epistemic way: directly following the logical tradition going back to Lakoff (1973). In a sentence, the treatment may reduce recurrence rates – the modal auxiliary weakens the truth commitment of the causal

claim, leaving open the two possibilities that the treatment does indeed reduce recurrence rates and that it does not.

Plausibility shields

Plausibility shields (worded differently here to imitate Prince et al. (1982)) mark the uncertainty of the speaker's perceptual or inferential process rather than the proposition and are indicated by it seems, it appears, and apparently. As we have seen, however, unlike a bare epistemic modal, a plausibility shield foregrounds the act of inference, which can have the interpersonal function of eliciting the hearer to join the inference, and perhaps modify it, a property that can be relevant in the feedback loop outlined in Section 4.7.

Attribution shields

Attribution shields place some of the epistemic responsibility for a claim onto another source, such as some argue, or according to the initial reports. For Hyland (1996), this form of hedging constitutes a specific strategy in academic writing because it enables a writer to mention a claim and distance him or herself from it at the same time; this strategy, which he claims, is at least as prevalent in political and journalistic writing, where attribution shields may serve to present personal disagreements while maintaining scepticism.

Mitigators

Mitigators are relation-hedges in Fraser's terms: they do not qualify propositional content but rather speech acts or downplay the interaction between speaker and hearer, as far as in Brown and Levinson (1987) account of face-work. They include parentheticals like I think if I may; syntactic declinings such as tag questions, as in I don't suppose you could look at this again; and dispreferred framing, as in If I may.... Mitigators can combine with a fully specified propositional content (see, for instance, section 4.3.1), illustrating the independence of derelational and propositional content-based hedging in language.

Vague category markers

Imprecise category markers imply enumeration or description as non-exhaustive, e.g. et cetera and various and things like that (Channell, 1994). Such amplifiers are generally more or less content-oriented like approximators, but have to do with the extension of categories as opposed to accuracy of numerical reference, and their associations are with informal registers (they occur less frequently in edited academic and political discourse).

Functional drift: from epistemic marking to interpersonal strategy

Finally, the six device types above are analytically distinct, but actual speech often repurposes epistemically marked forms for interpersonal (or strategic) purposes, a phenomenon I am calling "functional drift" here. In English, a modal auxiliary such as might, whose core semantic function is epistemic, can usually be used to mitigate a request or suggestion (you might want to think again about that), and in which case the degree of tentativeness on the speaker's part is not really relevant; it performs the

face-protective function discussed in Section 4.5. Likewise, the manner of drift away from a plausibility shield into a different kind of indeterminacy can be seen in political speech, where it appears that phrases usually used as an indirect report of inference, implying that the speaker has faced uncertainty in selecting the right conclusion, can be repurposed as a means of graciously making a statement without the gauntlet being thrown to the other person ("Well I guess she was just looking for excitement/Pretty eyes, as usual/I'm sure she tried her best"). Figure 2 illustrates the potential paths of drift for the six devices, from content-based to relation-based effects (vertical axis), and from plausible to strategic scope (horizontal axis). (Dashes indicate the general direction of this drift from one type of device to the other.

Figure 1. A pragmatic framework of hedging as a dynamic, feedback-sensitive process.

Figure 1. A pragmatic framework of hedging as a dynamic, feedback-sensitive process




Figure 1. Six types of devices within a larger interactional cycle. Hedging decisions are rooted in the speaker's own epistemic situation, that is, the actual use of the evidence, and face and institutional risks. This situation is projected via a set of one or more hedging devices which operate within one or more of the scopes of modification, relation or content established in 2.4 and 4.1 through 4.6, respectively. The receiver (hearer) then makes an inferential move to retrieve the maker's comprehension of her own degree of commitment. Such a move does not necessarily converge onto the maker's actual epistemic state, especially given the ambiguity of may (Crompton, 1997: 92). The resulting construct, in the diagram, of uncertainty reflects the fact that the hearer's final interpretation is a construction that results from both speaker evidence and hearer inference, not a straightforward transfer of consciousness. This constructed certainty feeds back to the speaker's subsequent choice of expression of hedging, as the speaker gauges her interlocutor's reception of the proposition with which she presents herself, and she next adapts, repairs, or intensifies her hedge (as indicated by the dashed feedback arrow).

Figure 2. Conceptual map of hedging device types by scope of modification and dominant pragmatic function, with dashed arrows indicating functional drift.

Online ISSN

3007-3197

Print ISSN

3007-3189


<http://amresearchreview.com/index.php/Journal/about>

Figure 2. Conceptual map of hedging device types by scope of modification and dominant pragmatic function, with dashed arrows indicating functional drift

Co-occurrence and hedge clustering

Another of the properties of naturally occurring discourse that is observable here, yet not so much exemplified in the device-by-device treatment presented in Sections 4.1-4.6, is that hedges often come in clusters rather than as single devices. A sentence like it could perhaps be argued that the results might suggest a modest effect stacks an epistemic modal, could, an approximating adverb, perhaps, an attribution-adjacent verb, argued, and a second epistemic modal, might, around a single proposition. This clustering of hedges is not superfluous in the way it might seem; here, for instance, while could soften the illocutionary strength of putting forward the argument in terms of statement power, might independently qualify the evidential import of results-to-effect inference.

This particularly heavy hedge clustering was identified by Hyland (1998) as a feature associated with the most epistemically vulnerable claims made in research articles (e.g. Most of the discussion section), in which the author interprets results in ways unsupported by the data themselves; such a concentration of discussion section claims has also been reported by Salager-Meyer (1994) in medical writing. Clustering thus provides a clausal-level indicator of which claims a speaker or writer perceives to be the most interpersonally or epistemically risky.

Summary taxonomy

Table 1 summarizes the six device types discussed above, their typical realizations, primary scope of modification, and representative sources.

Device type	Typical realizations	Scope of modification	Key source(s)
Approximators	roughly, about, -ish, a couple of	Content	Channell (1994)
Epistemic modal markers	may, might, could, possibly, perhaps	Content	Coates (1987); Lakoff (1973)
Plausibility shields	it seems, it appears, apparently	Content / inferential	Prince et al. (1982)
Attribution shields	some argue, according to, reportedly	Content / evidential	Hyland (1996)
Mitigators	I think, if I may, tag questions	Relation	Brown & Levinson (1987)
Vague category markers	and so on, or something, things like that	Content / category	Channell (1994)

Table 1. A pragmatic taxonomy of hedging devices in English.

Hedging Across Contemporary Discourse Domains

Academic discourse

Academic writing is by now the best documented of all the hedging sub-specialties, and from the corpus-based survey of the past three decades emerges not so much a monotonous uniform density but rather a sensitivity to discipline and genre. As Salager-Meyer (1994) demonstrates in the medical specialized use, the heavy concentration of hedge units in the discussion section of articles contrasts with a systematic discrepancy in subjectivities between English- and Chinese-medium applied linguistics journals in the distribution of hedges vs. boosters in article abstracts (Hu & Cao, 2011), a clear indication that academic genre is not simply a function of a discipline but rather a matter of discursive culture and audience. Likewise, the diachronic reviews collected in Section 2.5 point to an absence of a linear temporal trend: Hyland & Jiang (2016) observe not one but a downward trend in both soft and hard sciences over half a century, a decline over two decades in Science study oriented to a hard-science journal (Yao et al., 2023), and conclude that at least the latter are not to be explained as any one-directional change across all academic English.

Political discourse

Additionally, political argumentation draws upon many of the surface forms already studied in standard and academic writing, but the institutional pressures that influence the use of these strategies may be quite different. Where academic hedging takes on epistemic accountability for a factual claim at least in part because of genuine epistemic uncertainty, political hedging often takes on epistemic accountability

because of institutional pressure to do so but where the speaker's actual epistemic stance may be far more certain than the hedging utterance implies. The functional shifts from provisional acceptability and attribution of that acceptability away from the path of strategic ambiguity introduced in section 4.7 are particularly apparent when making political admissions, denials, and promises, where a hedge allows the speaker to attribute possible validity to some statement without assuming personal responsibility for the claim. The incorporation of hedging into language of strategic ambiguity as a shield to manage face directly mirrors the conceptual discussion in brown and Levinson (1987), since a political hedge is often in service to rather than a denial of true uncertainty about the facts.

Digital and AI-mediated discourse

Digital discourse (social-media text, the output of large language models, et cetera): a domain whose hedging practices are currently being established and studied in real time. Mo and Crosthwaite (2025) study how generative AI systems produce stance and engagement in academic writing tasks. This study is relevant for this article's argument because text produced by AI is hedged in repeating, sometimes formulaic, patterns (e.g. It is important to note, it is worth noting) that do not map simply onto the human epistemic, face, and risk reasons for hedging discussed in Sections 2 and 4. Because AI systems learn from large amounts of human-authored text, their patterns of hedging should be understood as statistical norms of the training data, not as reflecting a system's genuine epistemic stance; thus raising a new kind of pragmatic question of what exactly the point of a hedge is if a system has no epistemic stance of its own to defend. This question, along with the question of how an audience should evaluate trustworthiness in any hedged statement, is related to the wider literature on communicating ignorance or ignorance and certainty in institutional settings, such as, for instance, the finding by Budescu et al. (2009) that people tend to interpret the conjunction of very and likely as indicating an 80% probability, even though many hedging terms are assumed to be calibrated in the literature.

Implications for multilingual and pedagogical contexts

The three domains of comparison above are not in fact hermetically sealed off, and unfortunately strategies for hedging tend to cross linguistic and educational boundaries, albeit imperfectly. As discussed in 2.6, Vold (2006) found, through a cross-linguistic comparison of the markers of epistemic modality, that the inventory of available hedges and the naturalness of different scope-of-modification choices varies sufficiently across languages that when a writer attempts to translate an academic claim that has been hedged, the translation runs the risk of being under- or over-hedged: failing to hedge enough so that a target-language audience perceives the claims as more tentatively indeed intended by the source, or hedging too much so that the translation is flagged as more evidently tentative. This risk will be higher in the area of multilingual academic publishing alone, since, as the comparison of English-medium and Chinese-medium journals by Hu and Cao (2011) established, cultural norms of hedging diverge by publication culture even before translation is introduced. For English medium instruction and academic writing instruction especially, Skelton (1988) has argued for the importance of a systematic approach to hedging even after

nearly forty years, because students who are able to reproduce in writing the hedge-stock of the forms catalogued in Table 1, but who have not previously been exposed to and able to distinguish between the distinctions in scope-of-modification ways of hedge and the move of a soften and in what kinds of interactions according to how and in what ratio seen in Figure 2, they are prone to attacking the effect of diminishing the propositional contents of the claims without the effects of deterritorializing the interactional move of a soften rather than in what circumstances and in what sequence they have encountered forms of hedge, the language learning processes responsible for training students' hedging competence must not only be situated within language learning and teachable within the classroom context, but also located within the description of discipline-specific genres so that learners understand how and where specific forms of hedge function within the genre.

Comparative summary

Table 2 provides a qualitative comparison of the domains (i.e., a summary of those high-level, qualitative judgements derived within Sections 5.1 through 5.3) by highlighting the dominant hedging orientation and institutional motivation that prevailed in each domain rather than the number of occurrences (following the qualitative approach outlined in Section 3).

Discourse domain	Dominant hedging orientation	Primary institutional motivation
Academic writing	Content-oriented; epistemic modals, plausibility and attribution shields concentrated in discussion sections	Managing genuine empirical uncertainty and disciplinary norms of authorial caution
Political discourse	Mixed content- and relation-oriented; frequent functional drift toward strategic vagueness	Protecting institutional face and preserving accountability latitude
Digital / AI-mediated text	Formulaic content-oriented hedging, often decoupled from any underlying epistemic state	Statistical reflection of training-data conventions rather than speaker motivation

Table 2. Qualitative comparison of hedging orientation across three contemporary discourse domains.

Discussion

The analysis of Sections 4 and 5 lends support to three additional, more theoretical, claims concerning language use in modern-day English. Claim 1: the content versus relation distinction Fraser (2010) expects, we expect as an elegant prediction of the approximator-shield contrast Prince et al. (1982) foresaw, offers a simpler means of predicting which mitigation devices will co-occur with high speakers, and which ones will not, namely relation-oriented mitigators. Claim 2: functional drift as a result of

epistemically marked forms being borrowed for a variety of face-work or interpersonal purposes is not effectively localized to the speech act community, and indeed could be quite consequential, in the case of political discourse where the dissonance between a hedge's literal, epistemic import and the actual face-worked function can be wide-ranging. Claim 3: the emergence of text-generating technology as a source of face-worked moderation in intemperate contexts raises a real synchronicity-based 'hard' theoretical conundrum, since the entire framework outlined in Section 2, from Lakoff's (1973) fuzzy set logic to Brown and Levinson's (1987) face-worked policy, revolves around a hedger having beliefs, commitments, and social roles that must be maintained; but the language model under consideration simply does not possess these.

We can also tease out a second prediction from the model in Figure 1, which is that, because hedging is modeled as feedback-sensitive rather than as a global and time-invariant encoding operation, one would expect speaker hedging density to vary systematically over the course of a long interaction, a conversation, a drafting-and-revision cycle, or a political news cycle, as speakers monitor the evolving reception of their hedges and respond to them appropriately. These to-be-tested predictions align well with the diachronic research on academic conventions reviewed in Section 2.5, where one can see how adopted conventions have evolved over decades, at least in roughly predictive ways that could be plausibly related to changes in authors' perceptions of and response to the risk of loss of prestige, at least as evidenced in journal citations, like those of Hyland & Jiang (2016) and Yao et al. (2023), but not (to the present author's knowledge) in a single extended discourse (see also Crossley, 2022).

A final consequence can be derived for the calibration problem mentioned by Budescu et al. (2009, Section 5.3): if in everyday interaction hedges work only if speaker and hearer have the same calibrated expectation that a particular form is likely to communicate a certain message, the fact that the three domains considered in this article are associated with different institutional constraints means that the range of conventions regarding these forms is perhaps more diverse in the latter than in the former cases, which endangers cross-domain calibration. Someone used to reading the kind of hedging considered by Salager-Meyer (1994) in the discussion section-hedge normally expresses remaining empirical uncertainty in a limited domain for specific identifiable assertion-may immediately bring this interpretation to a statement that is hedged in a political speech where, as argued in Section 5.2, the same form most likely functioned strategically as an efficient political face-management device, for instance. As the conclusions in 5.4 prepare one for, the pedagogical solution suggested is the explicit teaching of the function of the forms rather than their lexical form, but the danger of crossdomain miscalibration seems even more crucial in the digital and AI mediated discussion of Section 5.3 in that this is a domain still new enough that the reader may lack a stable interpretive convention in mind about what the machine will communicate through the hedged political claim.

Limitations

The qualitative, example-based methodology employed in this paper, and section 3's conclusion that this is necessary because of the pervasiveness of the form-function

ambiguity noted by Crompton (1997), means the comparative argument in section 5 can only be comparably persuasive. Table 2 compares dominant trends, not actual sample counts, and should not be taken as a quantitative corpus result. A further potential limitation is the digital and AI-mediated domain in which Section 5.3 locates this corpus, which is likely to change at a rapid pace; it remains to be seen whether the detected hedging patterns in current LLM output can be replicated in systems trained differently and with different objectives. Finally, Figure 1's dynamic feedback framework, while fully consistent with the diachronic evidence considered, is a proposed synthesis of hypothetical processes, not a result based on new data, and the core assertion that hearer uptake flexibly influences a speaker's period of hedging in real time remains to be empirically established.

Another limitation arises from the set of languages and varieties included. The focus of the article is on English discourse: Section 2.6 and Section 5.4 reference cross-linguistic studies to help situate the English-specific patterns identified elsewhere in the paper, but do not present original cross-linguistic comparisons. If Vold (2006) and the cross-linguistic work that precedes him are correct to suggest inventory and relative naturalness of hedging strategies vary across languages, the six-device classification defined in Table 1, even if operationalized functionally and not lexically, should be viewed as a prediction of cross-linguistic applicability, not a conclusion. Likewise, the political and digital examples in Section 5.2 and Section 5.3 are based primarily on English media and English-large language model data; conventions of hedging in political communication and in AI systems with strong ean language influences may be different. Extending the comparative framework outlined here to these phenomena remains a task for future research.

Conclusion

To this end, this paper has demonstrated that foregrounding the variability of function along dimensions of scope of modification and discourse domain offers a fresh perspective on the nature of hedging as a pragmatic resource and has provided an overview of the extensive research on the topic in a form compatible with contemporary approaches to interactionally sensitive linguistic pragmatics. Building on classical work by Lakoff (1973), Brown and Levinson (1987) and Hyland (1996), I identified (types of) device, (types of) approximator, (types of) epistemic modal marker, plausibility shield, attribution shield, mitigator, and vague category marker. I demonstrated that the epistemic markers of uncertainty-making are recycled extensively for interpersonal and strategic purposes—a phenomenon of important import given the epistemic then interpersonal domain connotations which dominate accounts of politeness and face for most speech communities—in what I have characterized as functional drift. An interaction-oriented comparison of academic, political, and digital/AI-mediated discourse suggests this drift has a common feedback-sensitive realization in all three domains, in particular in the latter yielding a term frequency and distribution pattern that differs qualitatively from corpora of human discourse. The feedback-sensitive treatment, summarized in Figure 1, treats hedging as a dynamic intra- and intercultural interaction, thus weaving a subtle, empirically grounded fabric that unites classical pragmatic accounts of politeness and epistemic stance with the pending social-technical debate of how to talk about the

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

uncertain in the present technologically thick social environment. This hypothesis remains to be empirically tested through further corpus-based and experimental investigation of digital/AI-mediated interactions.

More generally, the findings hint that hedge research is in a sense about to enter a phase in which its familiar theoretical framework, created nearly entirely around human speakers negotiating face, accountability, and genuine epistemic doubt, must be raised at least one level of generalization to a communicative terrain that involves non-human text generators, globally dispersed multilingual audiences, and political and institutional agents now under a more intense glare of public scrutiny. The six-way taxonomy and the content versus relation distinction presented in Sections 2 and 4, respectively, seem on the basis of the evidence reviewed here to be adequate to this job, because they are functionally rather than lexically defined and so travel quite smoothly across the three domains being compared in Section 5. What remains to be seen—and here the article aimed to be more explicit rather than more conclusive—is how far hearers of all three discourse communities will learn to adjust their interpretation of a hedge when who, why, and perhaps whether the hedge was even uttered must be re-evaluated.

References

- Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage*. Cambridge University Press.
- Budescu, D. V., Broomell, S., & Por, H.-H. (2009). Improving communication of uncertainty in the reports of the Intergovernmental Panel on Climate Change. *Psychological Science*, 20(3), 299–308. <https://doi.org/10.1111/j.1467-9280.2009.02284.x>
- Channell, J. (1994). *Vague language*. Oxford University Press.
- Coates, J. (1987). Epistemic modality and spoken discourse. *Transactions of the Philological Society*, 85(1), 110–131. <https://doi.org/10.1111/j.1467-968X.1987.tb00714.x>
- Crompton, P. (1997). Hedging in academic writing: Some theoretical problems. *English for Specific Purposes*, 16(4), 271–287. [https://doi.org/10.1016/S0889-4906\(97\)00007-0](https://doi.org/10.1016/S0889-4906(97)00007-0)
- Fraser, B. (2010). Pragmatic competence: The case of hedging. In G. Kaltenböck, W. Mihatsch, & S. Schneider (Eds.), *New approaches to hedging* (pp. 15–34). Brill. https://doi.org/10.1163/9789004253247_003
- Holmes, J. (1990). Hedges and boosters in women's and men's speech. *Language & Communication*, 10(3), 185–205. [https://doi.org/10.1016/0271-5309\(90\)90002-S](https://doi.org/10.1016/0271-5309(90)90002-S)
- Hu, G., & Cao, F. (2011). Hedging and boosting in abstracts of applied linguistics articles: A comparative study of English- and Chinese-medium journals. *Journal of Pragmatics*, 43(11), 2795–2809. <https://doi.org/10.1016/j.pragma.2011.04.007>
- Hyland, K. (1996). Writing without conviction? Hedging in science research articles. *Applied Linguistics*, 17(4), 433–454. <https://doi.org/10.1093/applin/17.4.433>
- Hyland, K. (1998). Hedging in scientific research articles. John Benjamins.
- Hyland, K., & Jiang, F. K. (2016). Change of attitude? A diachronic study of stance.

Online ISSN

3007-3197

Print ISSN

3007-3189

<http://amresearchreview.com/index.php/Journal/about>

- Written Communication, 33(3), 251–274.
<https://doi.org/10.1177/0741088316650399>
- Lakoff, G. (1973). Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2(4), 458–508.
<https://doi.org/10.1007/BF00262952>
- Markkanen, R., & Schröder, H. (1997). Hedging: A challenge for pragmatics and discourse analysis. In R. Markkanen & H. Schröder (Eds.), *Hedging and discourse: Approaches to the analysis of a pragmatic phenomenon in academic texts* (pp. 3–18). de Gruyter. <https://doi.org/10.1515/9783110807332.3>
- Mo, Z., & Crosthwaite, P. (2025). Exploring the affordances of generative AI large language models for stance and engagement in academic writing. *Journal of English for Academic Purposes*, 75, 101499.
<https://doi.org/10.1016/j.jeap.2025.101499>
- Prince, E. F., Frader, J., & Bosk, C. (1982). On hedging in physician–physician discourse. In R. J. Di Pietro (Ed.), *Linguistics and the professions* (pp. 83–97). Ablex.
- Salager-Meyer, F. (1994). Hedges and textual communicative function in medical English written discourse. *English for Specific Purposes*, 13(2), 149–170.
[https://doi.org/10.1016/0889-4906\(94\)90013-2](https://doi.org/10.1016/0889-4906(94)90013-2)
- Skelton, J. (1988). The care and maintenance of hedges. *ELT Journal*, 42(1), 37–43.
<https://doi.org/10.1093/elt/42.1.37>
- Varttala, T. (2001). *Hedging in scientifically oriented discourse: Exploring variation according to discipline and intended audience*. [Doctoral dissertation, University of Tampere]. Tampere University Press.
- Vázquez Orta, I., & Giner, D. (2008). Beyond mood and modality: Epistemic modality markers as hedges in research articles. A cross-disciplinary study. *Revista Alicantina de Estudios Ingleses*, 21, 171–190.
<https://doi.org/10.14198/raei.2008.21.10>
- Vold, E. T. (2006). Epistemic modality markers in research articles: A cross-linguistic and cross-disciplinary study. *International Journal of Applied Linguistics*, 16(1), 61–87. <https://doi.org/10.1111/j.1473-4192.2006.00106.x>
- Yao, M., Wei, Y., & Wang, H. (2023). Promoting research by reducing uncertainty in academic writing: A large-scale diachronic case study on hedging in Science research articles across 25 years. *Scientometrics*, 128(8), 4541–4558.
<https://doi.org/10.1007/s11192-023-04759-6>