UNIVERSITY OF CALIFORNIA
SANTA CRUZ

GROUNDS FOR COMMITMENT

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

LINGUISTICS

by

Oliver Northrup

June 2014

The Dissertation of Oliver Northrup
is approved:

Professor Donka Farkas, Chair

Professor Pranav Anand

Professor Adrian Brasoveanu

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Grounds for commitment

by

Oliver Northrup

This dissertation proposes a novel approach to tracking not-at-issue contributions to discourse by bringing together several strands of research concerning evidentiality, illocutionary discourse particles, and speaker bias marking, all of which concern phenomena that indicate the relationships between speakers and the information they express in conversation. I argue that the overall discourse effects in each case should be captured in terms of commitments that are conditioned on various *evidential bases*. These bases derive from the speaker's private beliefs, his interlocutors' discourse commitments, and other contextually-rooted sources. They share the common purpose of publicizing the reliability of the commitments that invoke them.

The body of this dissertation provides arguments for this approach from several empirical domains. Chief among these are the sentence-final discourse particles of Japanese, and biased polar questions in English. For Japanese, I argue that the particles *yo*, *ne*, and their combination *yone* publicize the speaker's beliefs about his *relative authority* to sponsor the content of the particle-marked utterance, compared to that of his interlocutors. These simple conditions, encoded in the evidential base, interact with the default effects of an utterance to derive the total discourse effects of the particle-marked utterances, including the ways that these particles seem to limit possible felicitous responses, and why they are disallowed with certain sentence forms but not others. For English biased questions, I argue that *high negation polar questions* (HNPQs) and two kinds of *tag questions* (TQs) weakly commit the speaker to one or both answers, depending on the question's form. These weak commitments can be based on either a prior or current version of their *default evidential base*, which contains the totality of their private beliefs and any public contextual evidence. The notion of weak commitment is defined in terms of the evidential base's resistance to change in light of future discourse moves.

The results of this investigation are threefold. First, I argue thanks to novel data that each of the above phenomena—as well as others including English polarity particles and

the Singapore Colloquial English (Singlish) particle *lah*—receive superior empirical coverage under the individual analyses advanced here, compared to previous literature. Second, because the general architecture of each solution is shared, the discourse model advanced here allows for a more parsimonious explanation of the formal pragmatic effects of these phenomena. Third, by highlighting the commonalities among these domains and their relationship to evidentiality, the model allows for greater insight into the way that linguistic discourse is organized, the differences between default and non-default utterances, and more broadly, the distinction between formal pragmatics and general reasoning about language use.

Acknowledgments

I'd like to begin by thanking my advisor, Donka Farkas, who has guided this project since it its earliest incarnation as a squib for Semantics B during the winter of 2010. I will miss our weekly meetings, where she provided me with more discussion, advice, and encouragement than I ever could have hoped for; she has been the anchor of my career as a researcher. Thank you also to the other members of my committee, Pranav Anand and Adrian Brasoveanu, who have been both colleagues and role models. Pranav's breadth of knowledge, and the clarity with which he can see the questions at the core of a project, have been professionally crucial and personally inspiring. Similarly, I am constantly impressed by Adrian's ability to balance formal rigor with infectious enthusiasm.

Of course, there are many others who have helped this project along too. Thank you especially to Scott AnderBois, Manfred Krifka, Lauren Winans, Matthijs Westera, Seth Yalcin, and audiences at the UGGS workshop and Stanford's SemPrag group for their thoughts and suggestions. Thanks also to Adam Chong, Maho Morimoto, and Jun Yashima for their help with Singlish and Japanese data.

Finally, I want to acknowledge all the people who have shared this time with me. In particular, the Linguistics Department at UCSC is a close community, and I am grateful to the faculty, staff, and my fellow grad students for making my time here so rewarding. Thank you also to my parents, my grandparents, and especially to Lauren for supporting and encouraging me.

Chapter 1

Introduction

1.1 Overview

Conversation is a game played at several levels. At its most basic, speakers seek to inform one another about the facts of the world (including about each other). At its most abstract, speakers plan their actions and reason about their interlocutors' behavior to extract even more meaning. These are the realms of semantics and pragmatics, respectively, and although they seem quite distinct, we now know that the gulf between them is not as wide as once thought. For one, the underlying goal in both cases can be construed as the same: cooperative speakers seek to identify what the world is like and, ideally, agree on some facets of it to achieve common ground. Further, the inventory of moves that speakers can make is not limited to airing one's beliefs or attitudes. Rather, semantics and pragmatics are intermediated by a *discourse structure* that registers speakers' contributions and allows them to supplement the literal meaning of linguistic expressions with additional information about the context via morphological, syntactic, or intonational markers. The *discourse effects* of the resulting utterances therefore form a bridge between literal (semantic) meaning of linguistic expressions, and the higher-level conceptual strategizing of pragmatics.

This dissertation investigates several kinds of "additional information" utterances can convey using formal markers, in order to arrive at a clearer picture of how discourse is organized. These markers have been formalized independently before, but counter to previous literature, I argue that all of the markers discussed here are deeply similar and that this similarity should be captured at the level of discourse effects. Specifically, I argue that

the markers convey details about the *interloctuors' relationships with the content of their utterances*, in order to facilitate propagation of those utterances' content throughout the discourse structure.

The markers discussed in this dissertation include discourse particles, illocutionary evidentials, and expressions of speaker bias. First, consider the Japanese sentence final particles *yo* and *ne*, as exemplified in (1) below:

(1) *Souta is planning to cook for Hanako.*

a. Waffuru-ga daisuki **da**.
waffles-NOM great.like COP
'She loves waffles!'

b. Waffuru-ga daisuki da **yo**.
waffles-NOM great.like COP YO
'She loves waffles, man.'

c. Waffuru-ga daisuki da **ne**?
waffles-NOM great.like COP NE
'She loves waffles, doesn't he?'

d. Waffuru-ga daisuki da **yone**!
waffles-NOM great.like COP YONE
'She (sure) loves waffles, huh!'

All of these utterances convey the same literal content, but the approximate glosses also suggest that the speaker takes a different stance on that literal content in each case, especially with regard to the his expectations about whether the addressee will agree.

The evidential markers in (2), from Tariana (Arawak; Aikhenvald 2004), on the other hand, do not indicate anything about speaker's relationship with his content relative to his addressee, but rather the speaker's source of evidence:

(2) a. Juse iɾida di-manika-**ka**
José football 3SG-play-KA
'José has played football (we saw it).' *visual direct*

b. Juse iɾida di-manika-**mahka**
José football 3SG-play-MAHKA
'José has played football (we heard it).' *nonvisual direct*

c. Juse iɾida di-manika-**nihka**
José football 3SG-play-NIHKA
'José has played football (we infer it from visual evidence).' *visual indirect*

d. Juse iɾida di-manika-**sika**
José football 3SG-play-SIKA
'José has played football (we assume this on the basis of what we already know).'
inference

e. Juse iɾida di-manika-**pidaka**
José football 3SG-play-PIDAKA
'José has played football (we were told).' *hearsay*

By varying the form of the recent past marker, these utterances again convey identical literal information, but with additional details about how the speaker came to know the information.

Finally, the family of *biased polar questions* in English exhibit what I claim is a similar kind of discourse-level marking:

(3) *Speaker and addressee observe from afar as Jack studies a map.*

a. Does Jack know where Barcelona is? *Positive polar question*

b. Doesn't Jack know where Barcelona is? *High negation polar question*

c. Jack knows where Barcelona is, doesn't he? *Reverse rising tag questions*

Each of the forms above can be employed to ask the same literal question, but the high negation polar question in (3b) and the reverse rising tag question in (3c) indicate that the speaker does not regard both answers as equally likely.

This dissertation argues that in all of the cases mentioned above, the additional information conveyed by the marked utterances concerns the basis that the speaker has for making his contribution to the discourse, i.e., his body of evidence or *evidential base*. I introduce a novel model of discourse that derives the discourse effects of utterances from the evidential bases they introduce in discourse.

The argument proceeds as follows. The remainder of Chapter 1 introduces the basic discourse notions at the core of this dissertation, including the idea that discourse should be modeled in terms of individualized sets of speakers' *discourse commitments*, and that discourse effects can be understood in terms of *initiatives and responses* (Farkas and Roelofsen,

2012). Chapter 2 then expands this discussion by investigating how commitments qualitatively vary from one another, incorporating previous work on sourcehood (Gunlogson, 2003, 2008) and evidentiality (Aikhenvald, 2004; Murray, 2010). The result is a flexible model of discourse where every commitment enters the discourse with an evidential base, which carries certain properties by default but can be refined by the addition of formal markers. This approach allows for a unified treatment of the phenomena introduced above, using formal tools that I argue are independently needed.

Chapter 3 looks more critically at the idea of the discourse-evidential base and questions what kinds of meaning differences the base should be expected to capture. Empirically, it focuses on the Japanese discourse particles in (1), as well as the Singapore Colloquial English particle *lah*. I argue that these particles characterize the *relative authority* that speakers have over their commitments, a notion that necessarily extends the idea of discourse-based sourcehood as discussed in previous literature.

Chapter 4 discusses the idea of commitment strength, and how 'weak' commitment is the appropriate mechanism for capturing speaker bias in expressions like the biased polar questions in (3). The approach provides superior empirical coverage compared to previous analyses of these expressions, and while it initially appears overly-permissive, I show that given reasonable expectations about the space of variation for bias in polar questions, the account in fact predicts exactly the forms that are attested in English.

Chapter 5 concludes with a summary of the findings from Chapters 2–4. It then offers further discussion about the purpose of discourse-evidentiality and formal discourse-evidential marking in terms of general principles of linguistic communication.

Note that Chapters 2, 3, and 4 are presented to be fairly independent; Chapters 3 and 4 both contain summaries of the previous chapters' arguments, and the essential details of the discourse model developed here do not change after Chapter 2. The next section begins the dissertation proper, although the reader who is already comfortable with structuring discourse in terms of speaker commitments may safely skip to the subsection on the initial discourse model offered in §1.2.4 before continuing on to Chapter 2.

1.2 Discourse is a commitment space

This dissertation adopts the view that discourse is best modeled in terms of the commitments that speakers make in discourse, rather than simply a shared representation of their growing knowledge (i.e., their common ground). This section introduces the intuitions behind this approach and discusses how the details of a model built around discourse commitments should be investigated.

1.2.1 The common ground

One of the goals of conversation is to grow the *common ground*, generally thought of as the set of propositions that the interlocutors agree are true of the actual world for the purposes of conversation (Stalnaker, 1978, 2002).[1] Because a proposition describes the set of worlds where that proposition is true, it is sometimes helpful to think of the common ground in terms of those worlds and their similarity to the actual world. By intersecting all the sets of worlds described by the propositions in the common ground, we arrive at the *context set*, the set of worlds under consideration as candidates for the actual world at a given point in the conversation. Growing the common ground over the course of a conversation by adding new propositions is therefore equivalent to shrinking the context set, thereby 'zeroing in' on the actual world.

Note that uniquely identifying the actual world represents a limit case of what it means to pursue the goal of growing the common ground, but reaching this level of specificity is never a goal itself. Rather, speakers focus on more attainable goals by indicating which facts about the world they care about settling in the immediate future. In other words, they prioritize how they would like the common ground to grow, and work together to direct that growth in a "mutually accepted direction" (Grice, 1975).

Managing this process requires a few components. Chief among these is that all of the participants in the discourse must have access to roughly the same information about the shape of their common ground. Stalnaker (1978) argues that the participants in a discourse have their own representations of the common ground (an by extension, their

[1] As Stalnaker (2002) notes, use of the term 'common ground' is due to Grice, although it does not appear as such in his published work. Grice (1989) does, however, discuss that a proposition may have "common ground status."

own context sets). In *non-defective* contexts, these representations are similar enough that their discrepancies do not affect conversation. In *defective* contexts, problematic differences exist among the participants' representations, most often due to differences in what they presuppose about the context. Note that a particular discrepancy between the interlocutors' representations of the common ground can lead to either a defective or non-defective context, depending on whether that discrepancy matters in the present discourse. For example, consider a situation where Mary and Ezra have a friend named Andrew who lives an hour away in the capital, and Ezra erroneously presupposes that Mary knows that Andrew is in town this evening. Ezra and Mary will have representations of the common ground that differ (at least) with respect to the common-ground-status of the proposition that Andrew is in town. As long as Ezra and Mary's discourse does not involve the issue of Andrew's whereabouts, such a context would be non-defective. Consider, however, the discourse below:

(4) *Mary and Ezra are discussing their evening plans. Ezra, but not Mary, knows that Andrew is in town (instead of the capital, where he lives).*

M: What shall we do this evening?

E: We should get dinner with Andrew.

In this discourse, what Ezra attempts to communicate with his utterance—namely that he and Mary should meet Andrew for dinner *in town*—is not the same as what Mary understands, thanks to the discrepancy in their representations of the common ground. Note that many defective contexts can easily self-correct into non-defective ones. For example, if Mary and Ezra never eat dinner out of town, or if the capital where Andrew lives were five hours away, the implausibility of going there for dinner might allow Mary to very easily accommodate that Andrew is in town. That said, this kind of accommodation plays only a small role in the chapters that follow. Therefore, unless otherwise specified, the contexts presented and discussed in this dissertation should be taken to be non-defective. For ease of exposition, I take it for granted that every discourse evolves one unique representation of discourse which all of the interlocutors share, meaning that they agree on which worlds are contained in the context set.[2]

[2]Note that even in defective contexts, at least some of the interlocutors will not know the context is de-

1.2.2 Affecting the common ground

Having access to this shared representation is useless, however, unless the discourse participants are also able to cooperatively change it. They do this by performing *speech acts*, which codify the negotiations that change their common ground (Austin, 1962). Given than the common ground grows according to the goals put forward by the discourse participants, it is useful to think of this process as a kind of structured inquiry. This is clearly reflected in (4), where Mary's utterance indicates what kinds of facts about the world she has an immediate interest in determining, and Ezra's utterance asserts one such proposition.[3]

Exactly what effect Ezra's utterance (an *assertion*) has on the common ground is a point of debate in the literature. The debate hinges on the question of whether a single speaker should be able to directly manipulate what has been claimed to be a shared representation, or if the speaker is limited to making proposals that must be ratified before the primary effect of the utterance can affect the common ground. The distinction depends on whether rejecting the content of an assertion is the same as the assertion itself failing. Stalnaker (1978), for example, argues that the essential effects of an assertion with content p is to update the common ground with p, as long as "there are no objections" from the speaker's interlocutors. This means that if the asserted content is rejected (as Ain's is in (5)), the essential effect of the assertion fails.

(5) *Ain and Burr are arguing about the events of the day.*

A: Sonny got kicked out of class for fighting.

B: No he didn't!

Under this view, a speech act is a 'proposal' only in the sense that any other interlocutor may veto the move. If speakers can unilaterally change the common ground, speech acts are quite powerful, and their ability to inflict changes must be reined in by objections from the other participants.[4]

---

fective (e.g., Ezra in (4)). Implicit in the assumption of a shared representation of the common ground is the simplifying idea that speakers act as though the context is non-defective at all times, as well.

[3]Roberts (1996), among others, extends this view, arguing that all discourse is structured in terms of questions, whether overtly introduced or not. See below.

[4]It is important to keep in mind that assertions carry secondary effects, such as the mutual recognition that the assertion occurred (Stalnaker, 1978). This distinction evokes Sperber and Wilson's (1986) *informative*

A related strategy is proposed by Murray (2010), who argues that asserting p imposes an ordering on the context set, such that p-worlds are ranked over $\neg p$-worlds. If the addressee accepts, the less-preferred worlds are removed, otherwise the ordering is removed. In (5), this would mean that Ain alters the common ground by ranking worlds where Sonny got kicked out of class for fighting over those where this was not the case, but Burr rejects this ordering. In this way, the assertion does change the common ground, but in an easily reversible way.[5]

Alternatively, it is possible to take seriously the idea that the common ground can only be changed through negotiation, and that the essential effect of an assertion is therefore to propose a change to the common ground that can later be ratified, without altering the relevant part of the common ground itself (Gunlogson, 2003). This approach requires a more nuanced relationship between the portion of the discourse structure where proposals are made, and the shared body of information that the interlocutors are building. Both of these components are arguably 'common ground', but there is reason to believe that the negotiation of discourse moves deserves special status within the discourse structure.

In particular, it is clear that speakers track carefully the differences among their states. For example, if a speaker makes a proposal and then that proposal is questioned, the original speaker is still responsible for his original contribution. We can see this by considering the oddness of the continuation of (5) below:

(5′) *Ain and Burr are arguing about the events of the day.*

Ain: Sonny got kicked out of class for fighting.

Burr: No he didn't!

Ain: # Yes/oh, you're right.

Under a Stalnaker-style analysis, when Burr denies Ain's claim, Ain's utterance 'fails', in the sense that it does not alter the common ground in the way she (presumably) expected it to when she chose to make an assertion. In fact, though, her utterance has significant lasting effects; it is odd for Ain to go on to agree with Burr, because she is still 'on the hook' for

versus *communicative intentions*, where the former is the intent to inform, and the latter is the intent to inform of the intent to inform.

[5]The proposal also formalizes the idea of separate channels for 'at-issue' and 'not-at-issue' content. See Chapter 2.

her previous utterance. We might expect that in the name of growing the common ground, she should be able to agree with Burr and move on, but in fact the later move is instead self-contradictory. In a similar vein, Ain and Burr might simply agree to disagree, at which point they are both still responsible for their original stances (Farkas and Bruce, 2010).

Because of the importance speakers place on the kind of negotiated information detailed above, many researchers have been lead to restructure the representation of discourse in terms of not only what interlocutors share, but in terms of their individual *commitments* in discourse.

1.2.3 Discourse commitments

A commitment is a self-binding resolution to act as though a given proposition is true. This definition has its roots in Austin (1962), who argues that the effect of speech acts involves "committing" the speaker to "certain future conduct." The idea that commitments are propositional in nature is due to (Hamblin, 1970), which allows for immediate parallels with the idea of the common ground. Formally, commitments can be captured by articulating the common ground into a number of discrete sets for propositions that the individual interlocutors are committed to. These sets are known by various names, include *commitment-stores* (Hamblin, 1970), *commitment-slates* (Hamblin, 1971), and *speaker beliefs* (Gunlogson, 2003), although I will employ Farkas and Bruce's (2010) *discourse commitments*, abbreviated as *DC*s. To arrive at the classical common ground, one need simply intersect the interlocutors' commitments; whatever lies within this set is shared.

Whether this intersection is the entire common ground, or if there are separate repositories for, e.g., presupposed content, depends on how commitment is conceptualized; there is some variation in the literature in terms of what 'counts' as a commitment. For example, after uttering (6) below, we might ask whether the speaker should be taken to be committed to anything beyond the literal content of his utterance.

(6) Marle told me that Lucca is missing.

This utterance carries a number of entailments; but is the speaker actually committed to these, in the same way he is to the proposition that Marle told him that Lucca is missing? Stretching further, we might imagine that the context of utterance is such that everyone

knows that Marle is extremely trustworthy, so that (6) can be taken to reasonably convey that Lucca is in fact missing. In such a context, is the speaker therefore committed to the content of the embedded clause? For Hamblin (1971), at least, "all consequences of commitments are commitments," meaning that a speaker can be taken to be committed to a wider set of commitments than those he has overtly made. Walton (1993) suggests instead that a speaker's commitments should be partitioned into a "light side" of overt commitments and a "dark side" of those "only partially apparent or plausibly surmised." I remain agnostic on this point, although the model of discourse presented here only formally captures the "light side" of overt commitment. The common ground of a discourse is therefore taken to be the intersection of the interlocutors' (overt) commitments, plus a large body of presupposed information.

This stance is partially due to the fact that, as suggested by the original definition above, commitment is a public act, rather than an expression of knowledge or belief (Hamblin, 1970; Gunlogson, 2003; Beyssade and Marandin, 2006, among many others). The fact that commitment is independent of a speaker's mental state is clear when we consider that two known liars can converse without incident. They might leave their conversation knowing no more than when they started (beyond the secondary effects of having had the conversation), but the fact that they are manipulating personal promises of conversational consistency in the name of cooperation insulates the discourse from the effects of their deceit.

Of course, in most cases, we take it that our interlocutors' commitments do in fact mirror their thoughts and beliefs. This is because even though the immediate function of trading commitments is to promote consistency in order to grow the common ground, this consistency is most easily achieved when speakers are generally truthful. The closer a speaker's discourse commitments reflect his mental state, the less effort is required to maintain the differences between these two bodies of information.

As a result, I take it that making an utterance with content p does not mean that the speaker believes p, but that he presents himself as believing p. De Brabanter and Dendale (2008) explain this as the difference between understanding commitment as a "mental state" and as "the *expression* (sincere or not) of a mental state [emphasis theirs]". From this, his interlocutors can trust that the speaker will continue to act as though he believes that p is accurate. In most cases the interlocutors can then also conclude that the speaker

probably believes p, even though that is not what is being registered in discourse.

1.2.4 An initial model of discourse

What the discussion above lacks so far is a way to capture how exactly commitments are to be negotiated. To this end, this dissertation adopts and augments the "Tabletop" discourse model developed by Gunlogson (2003), Farkas and Bruce (2010), Farkas (2010), and most recently Farkas and Roelofsen (2012). The model follows a long tradition back at least to Hamblin (1958, 1970, 1971), Stalnaker (1978), Carlson (1983), and Roberts (1996). The structure of discourse under this model consists of two main components, the *Table*, a conversational scoreboard that registers a stack of issues under discussion; and sets of discourse commitments (DCs), one for each participant in the discourse. As participants engage in conversation, they raise issues by putting one or more alternative resolutions to the issue *on the Table*. Note that Farkas and Roelofsen's (2012) notion of alternatives draws from the framework of Inquisitive Semantics (Groenendijk and Roelofsen, 2009; Ciardelli and Roelofsen, 2011, among others), which I do not require here. I assume instead that the alternatives on the Table (and, for now, the discourse commitments populating the DCs) are classical propositions.[6]

The Table serves a more circumscribed role than, for instance, the Question Under Discussion (QUD) of Roberts (1996). As an example, consider the situation below:

(7) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a pineapple.*

R: Ahmed likes dragon fruit.

Raj's utterance raises a number of questions, many of which center around the fact that the purpose of this utterance is to direct Pepper to select a different fruit. In terms of *issues*, however, the semantics of the utterance supply just one proposition, that Ahmed likes dragon fruit. The Table only registers this single proposition; determining the extended pragmatic effects of Raj's utterance is left for other—presumably more general—reasoning processes.[7]

[6]The reader is invited to take the Inquisitive Turn for herself; there are a number of benefits with regard to homogenizing the effects of informative and inquisitive content which are beyond the scope of this dissertation.

[7]This story is more complicated if Raj's utterance contains any kind of focus marking, for instance on *dragon*

What the above discussion reveals is that whether an utterance involves a single proposed resolution (as in (7)) or several alternatives depends entirely on the semantics of the utterance: When an interlocutor employs a declarative sentence like *Ahmed likes dragon fruit*, it places the proposition that Ahmed like dragon fruit on the Table, while employing an interrogative sentence like *What does Ahmed like?* will place on the Table its full set of possible answers.

The fullest explication of the Table is due to Farkas and Bruce (2010), who argue that the Table actually holds not just the utterance's semantic content, but its syntax and phonology as well. In short, the Table stacks utterances, of which issues are just one component. This strategy is useful for, e.g., resolution of anaphora, but it is not relevant for this issues raised here, so I follow the example of later work by ignoring the Table's broader uses here.

The immediate goal of conversation (i.e., the way that interlocutors act when being maximally cooperative) is to *clear the Table* by settling the issues on it. Issues can be resolved in any way the interlocutors see fit, but the canonical resolution in most cases is for one of the alternatives to propagate throughout the various interlocutors' DCs. This is because the ultimate goal of discourse—again following Stalnaker (1978)—is to grow the common ground. Recall that any propositions in the intersection of the interlocutors' DCs are common ground, so the more that the interlocutors agree in their commitments, the greater the overlap will be.[8]

1.2.4.1 Default discourse effects

The major advantage of tying the issue-raising capacity of an utterance to its form is that it minimizes the need for illocutionary operators in order to bridge between the semantics of the utterance's content and the utterance's total discourse effects (Farkas and Roelofsen, 2013). Accordingly, the discussion here assumes that the default mapping between these domains is straightforward: When a speaker makes an utterance, he raises the content of

fruit; under some analyses such as Rooth (1992), focus marking invokes a set of alternative resultions much like the wh-question in the paragraph below. In (7), I assume Raj utters his assertion without focus.

[8]Committing to a proposition seems to project a clear path toward growing the common ground even without the Table, so it is tempting to try to do without the Table at all. I do not take this step here, because the Table affords a parsimonious way to make distinctions between *at-issue* and *not-at-issue* content, as discussed in Chapter 2. See also AnderBois et al. 2010 for further examples of not-at-issue content arising from appositives.

the utterance as an issue, and commits to acting as though the actual world lies somewhere within the space delimited by the proposition or propositions that comprise the issue. More schematically, default assertions and questions can be captured by the following illocutionary operation:[9]

(8) DEFAULT DISCOURSE EFFECTS

When a speaker X performs a utterance with content φ:

i. φ is placed on the Table.

ii. a. If φ is a proposition, φ is added to DC_X.

b. If φ is a set of propositions, $\bigcup\varphi$ is added to DC_X.

For utterances involving a normal declarative sentence like (7) above, the result is that by virtue of the utterance's form, we derive the effects of a *default assertion*—commitment to the propositional content, coupled with a call for that proposition to propagate throughout the discourse structure to the other interlocutors' DCs. For an utterance employing an interrogative sentence like (9) below, we derive the effect of a *default question*, which is to commit the speaker to the notion that some answer or other is true, and call for the issue to be settled.

(9) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed.*

R: Does Ahmed like pineapples?

Here, Raj's commitment is to the notion that Ahmed either does or does not like pineapples. While the overall effect of this commitment is trivial, note that it does make Raj responsible for any presuppositions the question carries. See Chapters 2 and 4 for more discussion on this point.

It is important to note that the default discourse move discussed above is not the only way an interlocutor can become committed to act as though a proposition is true. For instance, if Pepper wishes to signal that she accept the content of Raj's assertion in (7) so that it can removed from the Table and become common ground, she has many options:

[9]Note that in Inquisitive Semantics, a declarative sentence denotes a singleton set of propositions, rather than just a proposition. This makes the disjunction in (8b) unnecessary and further reduces the burden of the default illocutionary component.

(10) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a pineapple.*

R: Ahmed likes dragon fruit.

P: a. Yes, I know that.

b. Do you see one?

c. They're out of season.

In (10a), Pepper accepts p overtly with the particle *yes*, thereby adding p to DC_{Pepper}. In (10b–c), we see that no such overt acceptance is needed; Pepper does not challenge Raj, which can be taken as implicit acceptance of p. Either way, p is now a member of both DC_{Raj} and DC_{Pepper}, and so is common ground between them.

Crucially, the kind of implicit acceptance shown in (10b–c) is possible only because Raj's raised issue includes just a single alternative, p. This option goes out the window if multiple alternative resolutions are put on the Table, as when Raj asks a question:

(11) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed.*

R: Does Ahmed like pineapples?

P: a. Yes, (he does). *Commits Pepper to positive answer*

b. Do you see one? *Fails to commit Pepper to positive answer*

c. They're out of season. *Fails to commits Pepper to positive answer*

When Raj asked the polar question above, he raises the issue of whether Ahmed likes pineapples by placing $\{p, \neg p\}$ on the Table, rather than just p, where p is the proposition that Ahmed likes pineapples. Because there are multiple alternatives under consideration, it is no longer possible for Pepper to implicitly accept one, as shown by the infelicity of (11b–c), assuming her goal is to signal acceptance of p. The *yes* response in (11a), on the other hand, picks out p from the alternatives and signals Pepper's commitment to it.[10]

1.2.4.2 Investigating discourse structure

Following Farkas and Roelofsen (2012), I take it that the default discourse effects introduced above also act the formal starting points for defining various *non-default* discourse

[10]Why *yes* signals commitment to p instead of $\neg p$ is not addressed here, although it follows from Farkas and Roelofsen's (2012) full account.

moves, which can vary in terms of contextual requirements for their use, the kinds of responses they permit, and especially the commitments they impose on their speakers. This dissertation investigates a number of these non-default moves, so before moving to the body of the dissertation, this final section details how the 'non-defaultness' of an utterance should be understood in terms of the present discourse model.

Non-default discourse moves are *marked*, relative to default discourse moves, in a few senses. First, non-default moves tend to be formally marked, by means of intonation, morphology, or lexical choice; consider the intuitive difference between the effects of the default assertion and the utterances employing final particles below:

(12) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a pineapple.*

R: a. Ahmed likes dragon fruit. *Default declarative*

b. # Ahmed likes dragon fruit, huh. *huh-marked declarative*

c. # Ahmed likes dragon fruit? *Rising declarative*

(12b–c) are intuitively similar to (12a) in that they apparently all involve a declarative sentence form, but (12b–c) feature additional components beyond what is canonically expected for declaratives; (12b) includes a final particle, while (12c) uses final rising intonation instead of falling intonation.[11] The forms of these utterances are therefore formally marked compared to the default case.

The non-default utterances are also marked in terms of the conditions under which they become available for use, as shown by the fact that neither non-default option shown above is felicitous in the context provided. Changing the context so that Pepper picks up a dragon fruit, rather than a pineapple, in fact leads to rather different felicity judgements:

(12′) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a* ***dragon fruit****.*

R: a. Ahmed likes dragon fruit. *Default decl*

b. Ahmed likes dragon fruit, huh. *Right-marked decl*

c. Ahmed likes dragon fruit? *Rising decl*

[11] I ignore here the 'uptalk' use of rising intonation. See Gunlogson (2003, 2008) and Chapter 2 for significant discussion of these sorts of examples.

It seems that these non-default utterances are available when the context provides Raj with some grounds for his utterance.

Finally, non-default utterances are marked with regard to how they affect the discourse themselves, especially in terms of the commitments they impose. All of the utterances in (12′) seem to involve raising as an issue the proposition that Ahmed likes dragon fruit, but we understand that Raj's relationship with this proposition is quite different in each case. (12′a), for example, seems to offer approval or support of Pepper's choice. (12′b) instead shows that Raj has arrived at this conclusion on the basis of Pepper's actions. Lastly, (12′c) is a rather complicated move where Raj seems to seek verification of an inkling of his. This difference can be drawn out by testing whether Raj can deny foreknowledge of the proposition:

(12″) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a* ***dragon fruit****.*

R: a. Ahmed likes dragon fruit. # I had no idea. *Default decl*

b. Ahmed likes dragon fruit, huh. I had no idea. *Right-marked decl*

c. Ahmed likes dragon fruit? I had no idea. *Rising decl*

Only the default declarative seems contradictory with a denial of foreknowledge, suggesting that the discourse effect of these utterances—specifically, what Raj commits to when he performs them—is different in the later cases. How exactly the effects of these utterances differ will require further manipulations of the discourse.

In short, the discourse effects of non-default utterances can be investigated by careful manipulation of the preceding and following context, in order to determine how an utterance is limited by or itself limits the evolution of discourse. Throughout this dissertation, the investigation of non-default discourse moves, especially those than affect the quality of commitments, proceeds along these lines.

Chapter 2

An evidential model of discourse

2.1 Introduction

The previous chapter argued that modeling discourse in terms of the commitments that utterances impose on their speakers is a natural way to capture the dynamics of conversation and the evolving state of the interlocutors' common ground. This was achieved using a set of individualized commitment slates which propositions can be added to or removed from as the discourse progresses. This chapter introduces an additional complication to this picture by investigating how the status of a speaker's commitments can vary, both in comparison to other commitments, and over time. It concludes that not all commitments are the same, and introduces a novel discourse model to parsimoniously capture those differences that are linguistically encoded.

Conversational agents possess imperfect knowledge of the world, which means that sometimes they will need to change their commitments in light of new information. Some commitments, however, are easier to change than others. Consider the contrast below, when A tries to retract his commitment to having seen Tal leave on a train (13) versus merely having heard about it (14). Assume in both cases that A's utterance successfully implicates the embedded proposition.

(13) A: I saw Tal get on a train to Budapest this morning.
B: No, he's flying there.
A: #Oh, my mistake.

(14) A: I heard that Tal got on a train to Budapest this morning.

B: No, he's flying there.

A: Oh, my mistake.

In both examples, A reports some details of Tal's movements, which B contradicts. Assuming that B has superior knowledge of Tal's plans and so utters her contradiction appropriately, the result is a crisis of commitment, which A hopes to resolve in B's favor by retracting his report. The retraction is not equally successful in both cases, however, because A's direct, visual evidence in (13) is more difficult to write off than the hearsay in (14).

Note that the issue at hand crucially should not be confused with the question of whether A can change his mind. Recall from Chapter 1 that commitments are public declarations of consistency in *behavior*, not belief; when a speaker commits to a proposition p, he puts himself under the obligation to behave as though p is true for the purposes of the conversation, and nothing more.[1] Given the assumption of cooperativity as discussed by Grice and a non-trivial amount of (non-linguistic) reasoning about a speaker's motivations for making discourse commitments, a commitment to p can be taken to convey belief in p. While (13) might accurately reflect the changes in A's beliefs about Tal's travels, A's utterances are insufficient to publicly reflect these changes.

This is not to say that A's original commitment in (13) is unretractable. A can of course defer to B, but to do so he needs to take further public action to explain away the aberrant visual evidence. This suggests that the pressure of consistency, which drives interlocutors to maintain their commitments, is affected by the properties of the information that leads them to make commitments. Put another way, the 'staying power' of a commitment is a function of the reliability of the information that underwrites it.

Chapter 1 showed that commitments are an integral component of how speakers organize discourse, and argued that building a formal discourse structure around them is a prudent approach. This required breaking up Stalnaker's (1978) notion of the common ground into sets of propositions ascribed to different discourse participants. This chapter follows a similar rhetorical arc, by showing on the basis of examples like (13–14) that not all commitments that a speaker makes are equal either. An ideal discourse model should

[1] In fact, the requirements for belief are much more stringent than those of commitment; see McCready (2008a).

be able to reflect those differences in commitment that affect the felicity of future discourse moves. This requires explicit reference to both the content of a commitment, and the motivations for it.

This chapter establishes a model of this sort by drawing together documented examples of variation in commitment from the previously-cited work of Gunlogson (2008), Farkas and Roelofsen (2012; F&R), and others in §2 with data from the literature on (illocutionary) evidentiality in §3. It argues that these data can and should be captured in a uniform way, whereby commitments are understood to carry with them enough information about their strength and provenance to successfully capture how commitments vary. Drawing formal connections between these phenomena is the core innovation of this dissertation. The outline of this approach, and its formal skeleton, are presented in §4. (A number of refinements appear in Chapters 3–4.) As a digression, §6 considers the question of how a speaker determines whether a given commitment is an appropriate move, given his private knowledge. Answering this question fully is beyond the scope of this dissertation. §7 concludes.

2.2 Variation in commitment

This section introduces a few well-motivated distinctions among commitments as discussed by Gunlogson (2003, 2008), Farkas and Bruce (2010), Farkas (2010), and F&R. These distinctions revolve around *sourcehood*, a notion of conversational sponsorship whereby every commitment in discourse has at least one participant to undertake the burden of vouching for the truth of the commitment's content.

2.2.1 Sourcehood

Gunlogson (2003, 2008) shows that different utterances with superficially similar goals can yield different effects on the future of the discourse, by virtue of their form. Consider (15) below, where every possible response from B, at its core, is an attempt to agree with speaker A.

(15) A: The server's down.

B: a. Yes, I know. / Yes, that's right.

b. #Yes, I didn't know that. / #Yes, is it?

c. Oh (I didn't know that.) / Oh, is it?

As shown by the contrast between B's responses in (15b) and (15c), *yes* and *oh* correlate with different information states regarding the server's status. When B responds '*Oh*', she seems to merely accept the truth of the server being down on A's word. The fact that the particle can be followed with an echo question lends further support to this view, given that questions by default stake no claim about the asker's commitments.[2] On the other hand, if B replies '*Yes*', she cannot felicitously go on to explain that she *didn't* in fact know that the server was down, meaning that a '*Yes*' response requires that B knew the server was down independently of A's utterance.

Gunlogson (2008) captures this difference by introducing the notion of conversational *source*, the discourse participant who vouches for the truth of a commitment's content. She further proposes that every discourse commitment has at lease one source, meaning its content is marked as having been sponsored by a particular discourse participant who is assumed to have evidence of its truth. By default, a speaker is the source of his own commitments. This follows from the fact that, in normal discourse, we expect participants to have a preponderance of evidence for the commitments they make public and attempt to convince others of—a theme that will recur throughout this chapter. Note that in the same way that committing to p is different from believing that p, acting as source for p does not entail having unimpeachable proof of p. In (15), B can take it for granted that A has a preponderance of evidence for his claim, but A's assertion is silent on the details.

A participant who is not the source of a proposition can nonetheless accept it and become *dependent* on the proposition's source. This is exactly the distinction shown between *yes* and *oh* in (15) above; *yes* indicate that B is an independent source for the server's being down, while *oh* indicates acceptance of the idea, dependent on A's say-so. Here too, it is important to note that if B accepts A's proposition as a dependent, it does not mean that B must completely lack knowledge about the status of the server; if B works for IT, for example, she may be acutely aware of the server's troubled operational history, which would make her positively disposed toward accepting that it is down. By accepting as dependent, B simply signals that A's commitment is the crucial component that underwrites her own

[2] See Chapter 4 for an analysis of some illuminating exceptions.

commitment for the purposes of the present discourse.

F&R model sourcehood by subdividing a participant's discourse commitments into two separate 'bins', DC_X^s and DC_X^d. When X commits to p as source, p is added to DC_X^s. When X is dependent on some other interlocutor for p, he adds p to DC_X^d. When a participant places propositional content on the Table, by default he also commits to it as source. It need not be specified that the commitment that accompanies default assertions, for example, be performed as source; this simply follows from the fact that every proposition in the discourse structure requires a source. Committing as dependent, while not uncommon in discourse, is intuitively a more marked move, because of the more 'precarious' way in which it stacks a commitment on top of earlier ones; a dependent commitment places a special condition on the discourse, that some else must be committed to the same proposition. In terms of the structure above, this means that new commitments are added to DC^s, unless a move specifically marks that the commitment rests on previous ones, in which case its content is added to DC^d.

Based on the above, it is tempting to conclude that commitment as source is an integral component of what it means to assert under a discourse model that includes sourcehood. This is, however, not borne out.[3] Consider the version of the IT conversation below, where B follows her dependent *oh* response with what appears to be a simple assertion:

(16) A: The server is down.
B: Oh. Well, the server is down then.

In this example, B commits to the server's being down as dependent, and then asserts that the server is down as a sort of echoing move. This does not seem to require that B has suddenly become a source for the server's status, however, which means commitment as source is not tied to the notion of assertion. Rather, we again see that commitment as source is simply a default; When a speaker performs an assertion in absence of an appropriate source, he becomes the source. When a source is already available, like in (16), the assertion can make use of the existing source to indicate dependence.

Note that the sourcehood distinction as introduced here bears only tangentially on the introductory examples in (13–14). In both cases, A makes a normal assertion using a

[3]Special thanks to Donka Farkas for noticing this fact, which I missed repeatedly in previous revisions of this chapter.

declarative sentence, thereby committing to having seen (13) or having heard (14) that Tal got on a train to Budapest. Both commitments implicate that A believes that Tal did in fact get on the train, but nowhere is it represented in the discourse model that A's commitment in (13) is more difficult to divest. Sourcehood is a binary distinction: a commitment is either made as source or dependent, regardless of any nuances in the speaker's (epistemic) certainty. The binary character of sourcehood is similarly displayed in (13), this time across speakers. When A commits to the server being down and B responds with *yes*, there is no indication of whether A's or B's evidence for the server's status is better. In Chapter 3, I will argue that Japanese grammaticalizes this very distinction.

2.2.2 Conditional commitment

One point that is constant through the above examples is that when a speaker uses a normal declarative sentence to perform an assertion, he commits as source (by default) or dependent (under specific circumstances), and any more nuanced information he communicates about how strongly he trusts his information must result from other means. There are, however, bona fide cases where the form of an utterance affects the strength of a commitment. Consider the exchange below, drawn from Gunlogson (2008):

(17) *Laura has just entered the room, where Max sees her for the first time that day. Max says:*

  a. Did you get a haircut?

  b. You got a haircut?

  c. You got a haircut.

Moving down through these examples, we understand Max to be increasingly certain of Laura's timely trim. The question in (17a) does not commit Max to anything about Laura's hair, while (17c) clearly does. Less obvious is how the so-called *rising declarative* in (17b) should be treated. Gunlogson (2008) terms the more tentative discourse move encoded here a *contingent commitment*, though I will adopt F&R's term *conditional*. This kind of commitment is 'conditional' in the sense that it will only become an actual commitment if certain conditions obtain. Specifically, a conditional commitment is contingent on the *addressee's* commitment to the proposition in question. This means that if Laura responds

to Max's question positively as in (18) below, Max's conditional commitment strengthens, and he is not able to felicitously deny having thought she had gotten her hair cut.[4]

(18) *Laura has just entered the room, where Max sees her for the first time that day.*

M: You got a haircut?

L: I sure did!

M: I thought so. / # Oh, I had no idea. / # Really? It doesn't look like it.

Note also that Max cannot simply reverse his position, as shown by the infelicity of the final response above, *Really? It doesn't look like it.* The rising declarative therefore conveys Max's *bias* toward receiving a positive answer—a concept fleshed out more fully in Chapter 4.

If Laura disconfirms the bias of Max's rising declarative, the condition required to transform Max's conditional commitment into an actual one is not met. As expected, this means that he cannot go on to make other discourse moves that presuppose the positive answer, as shown in (19a).

(19) *Same as (18)*

M: You got a haircut?

L: No, not yet.

M: a. # Who cut it?

b. # (Yeah,) I didn't think so. / # Yes, I know.

c. Really? It looks like you did.

Similar to the confirming case, a disconfirming answer to a conditional commitment does not let Max off the hook, as shown by the alternative continuations in (19b). Even though Max's conditional commitment does not transform into an actual commitment, he still presents himself as having been biased, and so cannot commit as a source for the disconfirming answer. This is not to say that Max is precluded from objecting, however. If he is not ready to give up his bias, he can press further, as in (19c). Because commitment

[4]Note that in sufficiently rich contexts, 'rising declaratives' do not require rising intonation to achieve their intended effect. This fact does not inveigh strongly against the claim that rising declaratives have a distinct, formal pragmatic effect (Poschmann 2008, p.255). See §2.5.2.

is a public stance and not a perfect reflection of belief, changes in commitment never entirely cut off a speaker's ability to doubt the appropriateness of a discourse move. In other words, a conditional commitment does not become an actual one automatically. Rather, it publicizes the conditions under which the speaker could choose to make his commitment actual.

F&R model the conditionality of a commitment as a separate dimension of variation on par with the sourcehood distinction introduced above. When combined with sourcehood this way, each participant's discourse commitments splinter into four; an actual commitment can be made as source or dependent as above, but so can a conditional commitment. These possibilities are diagramed in Fig. 1.

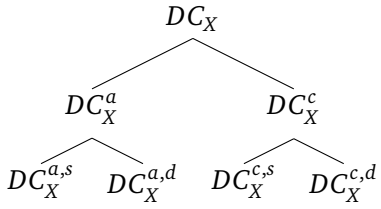

Figure 2.1: The structure of discourse commitment lists (F&R, p.50)

Note that regardless of whether a conditional commitment is made as source or dependent, the condition remains the same: Only after the addressee commits as source can the speaker's commitment become actual. Conditional commitments as source and dependent present the speaker as having different relationships with the proposition independent of the addressee, and so represent different classes of discourse moves. If a speaker makes a contingent commitment as source, he seeks a 'confirming' response to corroborate the evidence he already holds independently. If a speaker made a contingent commitment as dependent, his commitment is doubly tentative: He indicates both that he will not commit unless his addressee does, and that he *could not* cooperatively do so. F&R motivate this distinction by comparing rising declaratives and tag questions. In this dissertation I pursue a different analysis of these constructions, so see Chapter 4 for a review of their arguments. For present purposes, I instead will focus on the idea of conditional commitment as a whole, arguing that it is already dichotomous, even without the subdivision in Figure 1. Because of this, choosing to divide a speaker's DC in this way obscures deeper explanations about the ways that commitments vary, even if it is sufficient to capture the rising declarative data.

2.2.2.1 Implicit addressee authority

The conditionality of a commitment is a complex notion. First, whether a commitment to some proposition p is made as source or dependent, expressing conditional commitment to p reveals that the speaker expects his addressee to be able to definitively determine whether the commitment is an appropriate one. In Gunlogson's (2008) terms, this means the addressee must be *implicitly authoritative* about p, or able to act as a source for either p or its negation. A conditional commitment used appropriately therefore rules out the possibility of the addressee responding in a dependent fashion, as demonstrated below.

(20) *Max points at something at fruit stand where Laura often shops.*

M: That's a persimmon?

L: Yes, (it is)./#Oh, (I had no idea).

This example highlights that the behavior demonstrated here is not a result of Max's assumption that Laura knows what kind of fruit he's pointing to. What is interesting about the infelicitous response above is that it is *Laura* who seems to make an inappropriate move when she responds with *oh*; we could imagine a situation where by answering dependently, Laura signals that she does not in fact know, e.g., whether the fruit in question is a persimmon and expose Max's mistaken assumption about her knowledge, but this is not possible. It seems that the first component of conditional commitment is therefore a forward-looking sourcehood requirement imposed on the addressee's expected commitment.[5]

The reader might notice that this description, which involves the speaker reaching across the discourse structure, was argued against in Chapter 1. I stand by that point. As

[5]In one sense, this makes conditional commitment very similar to placing issues on the Table, which is perhaps the most forward-looking component of the discourse structure, given that its function is to register propositions to be considered for addition to the common ground. Indeed, some approaches complicate the Table, rather than the DCs, to account for 'in between' cases like the rising declarative in (17b). See Ettinger and Malamud (2013) for such a view. Briefly, they split the table into two parts, one which can register multiple alternatives, and one which can only support a single alternative (and so reveals stronger commitment). Lexical items and grammaticalized strategies can choose where an issue is raised, or can 'advance' an alternative from the multi-alternative partition to the single-alternative one. Given that (17b–c) are both declarative sentences, however, such views must divorce the issue-raising capacity of an utterance from its strict semantics. Under the present view, the declarative sentence form has a uniform effect on issue-raising, and only the character of commitments can change.

will be shown, overreach is avoided in the model developed here because the speaker (A) introduces the addressee's evidence as underlying his *own* commitment. The infelicity of a dependent response from the addressee (B) arises from the circularity of B making a commitment based on A's commitment that is in turn based on B's (now implicitly denied) evidence.

2.2.2.2 Weaker speaker authority

The second component of conditional commitment is that regardless of the sourcing, the speaker's certainty of that commitment must be weaker than the commitment she assumes her addressee will provide. Consider:

(21) *Laura has just entered the room, where Max sees her for the first time that day.* ***Laura*** *says:*

L: # I got a haircut?

M: # Yes, (you did)./#Oh.

Ignoring the possible uptalk interpretation of Laura's rising declarative, this utterance is pragmatically odd, because it seems to require that Max know more authoritatively what's happened to Laura's hair recently than she does. The situation is not improved even if Max is assumed to be implicitly authoritative on the issue (perhaps because he is renowned for his detailed knowledge of his friends' appearances). In other words, it is not simply enough to say that a conditional commitment anticipates an authoritative answer from one's addressee. Instead, that expectation is coupled with the knowledge that the addressee's response will trump whatever the speaker has committed to. The conditional commitment is therefore *weaker* than the response it anticipates.

Given two commitments to the same proposition, what does it mean for one commitment to be weaker than the other? As noted by Harnish (2005), commitment is not a monolithic notion in two regards. First, commitments are made for "reasons," meaning simply that every proposition a participant commits to is underwritten by larger bodies of information. This is clearly illustrated by the sourcehood distinction introduced above; some commitments are backed by a speaker's body of private knowledge, and some are backed (merely) by his addressee's say-so. Second, commitments can be made at various "degrees,"

based on the strength or reliability of the reasons underwriting them. Gunlogson (2008) leverages both of these dimensions of variation in characterizing conditional commitment. What a conditional commitment *lacks* at utterance time is backing by the addressee's presumed answer. And while the backing it does have might be based on independent reasons if the speaker is a source (like Max's knowledge of how Laura's hair looked yesterday versus today), these reasons do not support the commitment reliably enough for the speaker to commit while the addressee's answer is lacking. What is conditionalized in a conditional commitment is the speaker's credence in the commitment's content. One commitment is therefore weaker than another when the evidence that supports it leads less reliably to the conclusion that the commitment is correct.[6]

2.2.2.3 Independence of the two dimensions

Why go to the effort of reanalyzing conditional commitment into the combination of the two components above? Positive proof for this move comes from that fact that each component can be invoked independently. The first component, implicit addressee authority, already bears a striking resemblance to Gunlogson's (2008) sourcehood distinction, as discussed in §2.2.1. Chapter 4 will provide even stronger support from Japanese, where the discourse particle *ne* appears to encode exactly this notion. The second component of conditional commitment, weaker speaker authority, also appears as an independent phenomenon, as shown by the discourse particle *wohl* in German (Zimmermann, 2004, p. 543):

(22) Hein is auf See.
Hein is at sea
'Hein is at sea.'

(23) Hein is *wohl* auf See.
Hein is *wohl* at sea
= Speaker assumes that Hein is at sea

As indicated by the glosses above, *wohl* contributes the notion that the speaker's assertion is based on reasoning, rather than any kind of direct experience. The particle does not modalize the proposition, nor does it introduce a separate layer of not-at-issue meaning:

[6]This characterization leaves open the question of how variation in commitment relates to variation in epistemic modality. This point is best addressed after the formal model is introduced, and returns in §2.5.

> Crucially, a sentence containing *wohl* does not say that the state of affairs described by the sentence without *wohl* holds. Rather, the presence of *wohl* has the effect that the state of affairs described by the sentence is still unresolved. (Zimmermann, 2004, p. 554)

Poschmann (2008) argues that *wohl*'s effect is best captured as an indicator of reduced speaker commitment, in line with the second component of conditional commitment as discussed above. Further examples of this sort, drawn from the evidentiality literature, appear in the next section.

2.2.3 A fledgling model of discourse

At this stage, the amount that the model of discourse being developed here has changed from the end of Chapter 1 is quite modest. Eschewing Gunlogson's (2008) and F&R's implementations for capturing the sourcehood and (bipartite) conditional commitment data in this section, it seems that the empirical space yet to be mapped can be summarized by (24iib) below, where each commitment is associated with additional information about the evidence underwriting the commitment.

(24) KEY COMPONENTS OF DISCOURSE (SUMMARY)

i. The Table, to register issues

ii. A DC_X for each participant X containing:

   a. The propositions that X is committed to

   b. (Some) information about the **strength of and reasons for** each proposition

But how should this additional information about commitments be formalized? As we've seen, existing accounts offered by Gunlogson (2008) and F&R essentially 'tag' commitments with sourcehood and conditionality information, and tie configurations of these tags to the forms of utterances. I am wary of this approach for two reasons. First, as discussed above, conditional commitment is a complex notion that involves a particular combination of commitment sourcing and commitment weakness. Ideally, we'd like an explanation for why the two factors come together in this way. Second, there are likely other distinctions that are made among commitments by virtue of their form, and simply subdividing the commitment sets further to account for these as they are discovered will rapidly lose explanatory power.

I argue instead that by jointly tracking the components of conditional commitment, reasons and strength, we can arrive at a more flexible model that still captures these existing variations in a natural way. Fortunately, we needn't look far to find a prototype for such an approach. The following section introduces evidentiality, a suspiciously-similar phenomenon where languages mark the medium through which information was learned and, by extension, the reliability of that information.

2.3 Evidentiality as grounds for commitment

From a theoretical perspective, the central goal of the present chapter is to expand on (24iib) above; what information should a commitment convey, beyond its propositional content, and how should it be formalized? There are a number of semantic and pragmatics notions that could fall under this expanding umbrella, including evidentiality and epistemic modality. This section begins with the former, evidentiality, and argues that it shares a family resemblance with the commitment variations of the previous section. This is because at some level, every successful approach to evidentiality requires paring the main information that the evidential utterance conveys with information about the body of evidence that underwrites that information—the very relationship I have argued is reflected by the phenomena in the previous section.

2.3.1 Evidentiality primer

Evidentiality "[indicates] a particular relationship between a person and a proposition," specifically a relationship that specifies how it is that the speaker came to know or suspect that the proposition is true (Murray 2010, p. 46–7). Note already how this definition seems to paraphrase the facts about commitment from the previous section. In the case of evidentiality, this "particular relationship" is marked by introducing the source of evidence for a proposition, such as *hearsay* in the example below.

(25) A: Where's Milgrim?

B: **Word is/I hear/They say** he skipped town.

All of the options in the bolded portion of B's response above accomplish the same function; they indicate that hearsay evidence is what allows B to posit that Milgrim skipped town.

Note that the notion of an evidential source is distinct from sourcehood as discussed in the previous section; B is in fact the 'source' participant here, in virtue of the fact that he has answered A's questions with an assertion and that there is no prior commitment that would license use of the dependent base.

2.3.1.1 Evidentials versus evidentiality

Aikhenvald (2004) notes that while all languages can express evidentiality, about a quarter of the world's languages use specific lexical items to do so. The words employed in this lexical strategy are *evidentials*, whose primary meaning is to express evidentiality. Languages tend to reliably grammaticalize portions of the same small set of distinctions in evidence source, for example, direct evidence—visual, auditory, etc.—might be contrasted with various kinds of indirect evidence such as inferences or conjectures (Willett, 1988; Speas, 2004). Among indirect evidence, languages often distinguish between inference on the one hand, and hearsay on the other. Tariana (Arawak), for example features a five-way contrast in evidentials (Aikhenvald, 2004, p. 2–3):

(26) a. Juse iɾida di-manika-**ka**
José football 3sg-play-**rec.p.vis**
'José has played football (we saw it).' *visual direct*

b. Juse iɾida di-manika-**mahka**
José football 3sg-play-**rec.p.nonvis**
'José has played football (we heard it).' *nonvisual direct*

c. Juse iɾida di-manika-**nihka**
José football 3sg-play-**rec.p.vis**
'José has played football (we infer it from visual evidence).' *visual indirect*

d. Juse iɾida di-manika-**sika**
José football 3sg-play-**rec.p.assum**
'José has played football (we assume this on the basis of what we already know).'
inference

e. Juse iɾida di-manika-**pidaka**
José football 3sg-play-**rec.p.rep**
'José has played football (we were told).' *hearsay*

These evidential markers are rigidly grammaticalized, as evidenced by the fact that they fuse with tense markers in this language.

Evidentiality can also be expressed periphrastically, often using verbs with primary meanings related to the particular kind of evidence involved; expressing evidentiality is a secondary use. Consider, for example periphrastic English translation of the Japanese visual evidential *mitai*, below.[7]

(27) Hanako-ga moo yonda mitai.
Hanako-NOM already read VISUAL.EVIDENTIAL
'It looks like Hanako already read (it).'

Further examples of periphrastic evidentiality are sprinkled through the previous sections; any attributive parenthetical or sensory verb construction (e.g., "I heard that Tal got on a train to Budapest this morning" in (14)) falls into this category.

2.3.1.2 Illocutionary and epistemic evidentials

Broadly, there are two families of evidentials that differ in whether the evidential content they introduce is integrated into the main content or set apart from it (Matthewson et al. 2007; Murray 2010). The evidential content supplied by *illocutionary evidentials*, such as those of Cuzco Quechua (Faller 2006), is "parenthetical," meaning that it does not interact with the main content of the utterance. The meaning of these evidentials is best paraphrased in English using actual parentheticals, like *That's the postman, I take it* (Murray 2010). As their name suggests, these evidentials are often formally captured as speech act operators.

As discussed by Faller (2002) and Matthewson et al. (2007), one important property of illocutionary evidentiality is that the evidential component is not directly negotiable—although it can of course be negotiated indirectly. This means that polarity particles cannot target the evidential content, suggesting it is "not-at-issue," literal content contributed by an utterance which nevertheless does not affect the future direction of the discourse.

(28) A: That's the postman, I gather.

[7] *Mitai* has a non-evidential use as well, often seen in noun phrases such as *ookami mitai na inu* ('(a) dog that looks wolf-like').

B: # No, you saw it happen. / What makes you think so?

This property is shared with the variations in commitment from the previous section:

(29) A: That's the postman?

B: (# No,) I'm not in a position to settle that any more than you are.

Here, A makes a conditional commitment concerning the identity of the postman. B can object to A's assumption that she is able to settle this issue, but cannot use *no* to reference the commitment-related content.

The other class of evidentials, *epistemic evidentials*, does not share this property. The evidential content of these expressions seems to be integrated into the main utterance content, much like the English modal in *That must be the postman*:

(30) A: That must be the postman

B: a. That's not true. That's the UPS guy.

b. That's not true. There are other plausible answers. That might be the UPS guy.

(adapted from a similar example of Matthewson et al.'s (2007))

In (30a), B denies the proposition that the subject is the postman, while in (30b), he instead denies the modality of A's proposition. Both are licit, which indicates that the modal contribution of A's utterance goes beyond a modification of A's illocutionary force. Epistemic evidentials behave identically and are usually analyzed as a special class of modals (see Izvorski 1997; Matthewson et al. 2007). As such, they fall outside the scope of the main discussion here.

2.3.2 The not-at-issue effects of evidentiality

The discussion above sneaks one further distinction into our discourse model, briefly mentioned in Chapter 1. This is between *at-issue* content, roughly the 'main point' of an utterance, and *not-at-issue* content, which includes presupposed content but also new, backgrounded content, like what is introduced by the parentheticals above (Farkas and Roelofsen, 2012). Within the model proposed so far, the at-issue/not-at-issue distinction is derived

from whether a given proposition has appeared on the Table. Because the Table is the locus of negotiations about changing commitments, anything that is placed on it is by definition at-issue. If a speaker performs an utterance with not-at-issue content, he commits to it as normal, but fails to register it as an issue in the discourse. This does not mean the content 'sneaks by' the addressee in any way—simply that the speaker assumes it will be uncontested.

2.3.2.1 Finding the 'main point'

I assume that identifying the at-issue and not-at-issue components than an utterance contributes can be read from the denotation of the sentence the utterance employs, such that the literal meaning of the sentence is the main point, and the non-literal meaning is not. It is worth noting, however, that this is likely a simplification. Papafragou (2006) and Simons (2007) argue convincingly that context plays a major role in determining which contribution of an utterance is the main point. For example, in the context of A's question in (31), all of B's responses are naturally construed as answers, even though in (31b–c), the answer is presupposed or implicated.

(31) A: Who was Louise with last night?

B: a. She was with Bill.

b. Henry thinks/I think that she was with Bill.

c. Henry is/I'm convinced that she was with Bill.

(Simons, 2007)

This suggests that the main point of all three utterances is the literal meaning of (31a), and that the matrix clauses in (31b–c) contribute an attributive—even evidential—meaning that conveys B's "limited certainty" in the answer.

The view that the these utterances warrant an evidential interpretation is further strengthened by the fact that matrix clauses which do not serve to bolster B's answer are inappropriate in this context:

(32) A: Who was Louise with last night?

B: a. (?) Henry hopes/I hope that she was with Bill.

b. ? Henry wishes that she was with Bill.

c. ? Henry dreamt that she was with Bill.

(Simons, 2007)

The notable exception is that if Henry "has a reputation as a seer," (32c) is licit. In other words, if the context allows Henry's dreams to count as appropriate evidence, the utterance can be interpreted as expressing evidentiality.

It seems clear that the literal content expressed by an utterance is, by default, its main point. Recall that under the 'simple' view of main point selection, this was the only option. Based on the examples above, it is tempting to conclude that in contexts that support it, non-literal content including presuppositions and implicatures can be promoted to main-point status. Under this 'choice' view of main point selection, an utterance supplies a number of propositions which from which the main point is chosen. Even this, however, turns out to be too strict an understanding. Consider B's replies to the question below, which Simons (2007) calls *exclusionary* answers:

(33) A: Which course did Louise fail?

B: a. Henry, the idiot, thinks she failed calculus.

b. Henry, entirely wrongly, is convinced that she failed calculus.

c. Henry, falsely, said that she failed calculus.

(Simons, 2007)

These utterances all convey that calculus is not the answer to A's question. If this exclusionary response is taken as the main point of these utterances, it is problematic for both the 'simple' and 'choice' views of main point selection, because the main point of these utterances only comes about via an interaction between the content of the matrix clause, the subordinate clause, and a parenthetical phrase. For example, only by combining the proposition that Henry thinks Louise failed calculus with B's parenthetical evaluation of Henry's mental acumen does (33) express that the calculus is an incorrect answer to A's question.

It therefore appears that definitively settling how the main point of an utterance is calculated is discourse is a tricky notion. The differences between these conceptions will not play a pivotal role in this dissertation, so the reader is referred to AnderBois et al. 2010 for a fuller exploration of the issue through the lens of appositives. What is common to all

of these views, and crucial for the present enterprise, is that expressions of illocutionary evidentiality fall squarely outside the main point of evidentially-marked utterances. Illocutionary evidentiality is not-at-issue, so what we require now is a clearer idea of how not-at-issue content interacts with the discourse structure.

2.3.2.2 At-issue and not-at-issue as paths

To say that some content conveyed by an utterance is not-at-issue is to say that the speaker contributing that content does not seek to have it ratified by his addressee. There are a number of reasons that this might be the case. For example, a speaker might wish to register his attitude toward a topic of conversation, without diverting the conversation's trajectory:

(34) *Lucca returns to the campsite to find Janus missing. She sits down with Glenn, who looks cross.*

L: Where did Janus go?

G: The jerk ran off again without saying!

Here, Glenn's immediate goal is to address Lucca's question. He also wishes to voice his negative impression of Janus, so he does it in a with a not-at-issue using the expressive *the jerk* (Potts, 2007). Part of the reason expressive content like this is appropriate to be treated as not-at-issue is because there is very little Lucca can do but accept it; in normal contexts, Glenn is the ultimate judge of his own emotions, so there is no point in raising his expressive content as an issue.

Murray (2010, 2011, 2012) presents a parallel arguments for illocutionary evidentiality, and conceives of at-issue and not-at-issue content as followed two different 'paths' for entering the discourse. Consider a evidential utterance like (35) from Cheyenne (Algonquian), which contains a direct evidential marker (realized here as a null morpheme):

(35) É-hó'tȧhéva-∅ Floyd.
3-win-DIR Floyd
'Floyd won, I'm sure.' (Murray, 2010)

Under Murray's (2010) account, this utterance contributes two different updates to the common ground. Note that Murray's analysis eschews individualized *DC*s in favor of a

unified common ground. The first update is non-negotiable, and concerns the expression of direct evidentiality. The fact that speaker has personal experience that supports the proposition that Floyd won automatically becomes common ground, thereby reducing the worlds under consideration (i.e., the context set).[8] The second update concerns the at-issue meaning of the sentence, that Floyd won. Rather than eliminating worlds, this update orders the context set so that worlds where Floyd won are preferred to those where he did not. This leaves the common ground in a state where the addressee can then choose to eliminate the dispreferred worlds. In other words, the not-at-issue content imposes an update on the common ground, while the at-issue content proposes one.

Murray's final analysis is significantly more nuanced that the sketch above, but even at this informal level, it is clear that the presence of an evidential source is exactly the kind of information suited for the not-at-issue path into the discourse structure. After all, no interlocutor could hope to contest existence of the speaker's direct evidence in (35), or especially the inference contributed by *I gather* in (28), repeated below.

(28) A: That's the postman, I gather.

B: # No, you saw it happen. / What makes you think so?

An identical argument can be made for sourcehood configurations in §2; if a speaker performs an assertion, he raises the issue of whether the propositional content of his assertion should be accepted by all of his conversational participants. He does not invite challenges to his personal suitability as the assertor—this is taken for granted.

What is lacking under the above view, however, is a satisfactory explanation of how exactly the relationship between the not-at-issue evidential source and the at-issue content should be captured. After all, the evidential source isn't merely secondary information, the way that Glenn's use of *the jerk* to refer to Janus is in (34). Rather, we understand the at-issue content of evidentially-marked utterances to rely on the evidentially-marked content.

[8]The idea of uptake is avoided here. See Ch 1 for arguments against the idea of imposing *anything* on the common ground.

2.3.2.3 Modulating strength

Finally, another parallel between the variations in commitment discussed in the previous section and illocutionary evidentiality as described here is that under certain conditions both modulate the strength of commitments. Consider the examples below, from Central Alaskan Yup'ik (Eskimo-Aleut; Krawczyk 2009, p. 5–6):

(36) a. Aya-llru-uq.
leave-PAST-3RDSG
'She left.' *Unmarked*

b. Aya-llru-llini-uq.
leave-PAST-INFER-3RDSG
'Evidently, she left.' *Inferential evidential*

c. Ayag-ni-llru-a.
leave-said-PAST-3RDSG
'He said that she left.' *Verb 'say'*

d. Aya-llru-uq-gguq.
leave-PAST-3RDSG-HEARSAY
'It is said she left.' *Reportative evidential*

Krawczyk (2009) reports that while (b) commits the speaker to the subject having left, (36d) in particular does not:

(37) a. #Aya-llru-llini-uq . . . Aya-ksaite-llru-yuka-a
leave-PAST-INFER-3RDSG leave-past-think$_t$hat-3RDSG
#'Evidently, she left. . . I don't think that she left.' *Inferential*

b. Aya-llru-uq-gguq. . . . Aya-ksaite-llru-yuka-a
leave-PAST-3RDSG-HEARSAY leave-past-think$_t$hat-3RDSG
'It is said she left. . . I don't think that she left' *Reportative*

In these examples, the evidentially-marked utterance is followed by a contradictory bare utterance. Similar to the English glosses, the result is infelicitous when the speaker employs the inferential evidential, but felicitous with the reportative. This suggests that the inferentially-marked utterance commits the speaker to the matrix content, while the reportative-marked utterance does not.

This section has briefly presented some facts about evidentiality in general and evidentials in particular. These grammatical markers encode relationships between speakers than the propositions they commit to in discourse, akin to sourcehood and conditionality from the previous section. The next section crafts a discourse model to help make these similarities more explicit.

2.4 A revised model for commitments

This section introduces a novel model of discourse, based on Farkas and Roelofsen (2012; F&R) but incorporating a more flexible view of the way that commitments can vary, and especially the way that certain not-at-issue content can be used not only to augment at-issue content, but underpin it. In addition to bringing formal symmetry to the data in §2–3, the resulting model will be applied to capture the discourse effects of biased questions in English (Chapter 4), as well as some cross-linguistic generalizations about discourse particles (Chapters 3 and 5).

The core of the system is an expanded notion of commitment; building from the example of evidentiality, I argue that all commitments—whether overtly encoding evidentiality or not—can be made at varying degrees of confidence, based on a particular set of background information.

2.4.1 Model architecture

The basic architecture adopted at the end of Chapter 1 and expanded slightly in §2.4 is repeated below. Discourse is modeled with two major components, the conversational stack or Table, and a collection of Discourse Commitments for each participant in the discourse:

(24) KEY COMPONENTS OF DISCOURSE (SUMMARY)

i. The Table, to register issues

ii. A DC_X for each participant X containing:

a. The propositions that X is committed to

b. (Some) information about the **strength of and reasons for** each proposition

The notion of the Table as an issue-raising component of discourse is unchanged from earlier work. What differs under the present model is the structure of a speaker's discourse commitments. Unlike F&R, DC_X is not partitioned on the basis of sourcehood and conditionality. Instead, individual commitments bundle within themselves the capacity to reflect a broader set of distinctions. In addition to its propositional content p, a commitment carries this additional information in the form of an *evidential base* E. The pair of these components $\langle p, E \rangle$ comprises my reformulation of a commitment.

2.4.1.1 The evidential base

The evidential base E is the set of propositions that, when taken together, serve to underwrite a commitment. This set is directly analogous to the modal base of Kratzer (1981, 1987)—a connection discussed more fully in §2.5.3. When a speaker makes an evidentially-marked utterance, it introduces an appropriate E, which is immediately used as a component of any commitments that the utterance spawns. For example, in (36d) from Central Alaskan Yup'ik, repeated below, the hearsay evidential *-ggug* introduces the existence of a set of propositions that derive from having heard that the speaker's friend left the house:

(36d) Aya-llru-uq-gguq.
leave-PAST-3RDSG-HEARSAY
'It is said she left.'

This evidential base underwrites the proposition that she left, which remains the (only) at-issue component of meaning offered by this utterance, as shown previously.

At this point, it is crucial to make a stark distinction between the *content* of an evidential base on one hand, and its *character* on the other. The content of an evidential base is an enumerable set of propositions, known to the speaker alone, and not available to the other participants of the discourse, unless they can infer them. The character of the base, on the other hand, is what is publicly disclosed by an evidentially-marked utterance. For the hearsay evidence in (36d), this means E might be rather opaquely captured as a set $\{p : p$ contributes hearsay evidence that the speaker's friend has left her home $\}$, notated E_{HEARSAY}. In other words, what is entered into the discourse structure is a description of the evidential base, rather than its exact propositional content (c.f. Kratzer 1987).

2.4.1.2 The default base

If every commitment is a pair of a proposition and a description of the evidential base that underwrites it, then what happens when an utterance fails to introduce any overt evidential content? Any commitments that such an utterance spawns clearly still rest on a base of evidence, but that base must be rather generic, considering the variety in the evidence that can underwrite commitments from non-evidential utterances. For example, consider (38a–c) below, which constitute a single discourse where Fatima makes three assertions about the geography of Western Africa:

(38) *Oumar and Fatima are filling out a crossword puzzle about the geography of Western Africa.*

a. O: What is the capital of Mali?

F: (*Recalling a conversation she overheard at a coffee shop:*)
The capital of Mali is Bamako.

b. O: Thanks. What about the Gambia?

F: (*Thinking back to her school days:*)
The capital of the Gambia is Banjul.

c. O: Great, last one for now: Togo? I'm totally stumped.

F: (*Giving a sigh of frustration and consulting a map:*)
The capital of Togo is Lomé.

In each case, Fatima commits to the proposition she has asserted, but the specific evidence that allows her to make these commitments varies from her perspective. Fatima's evidence could be described as hearsay (38a), past experience and beliefs (38b), and her interpretation of a contextually salient map (38c), respectively. Nonetheless, the details of the differences among these bodies of information is not publicized by her utterances, which means that the character of the evidential base employed in each case must be identical. This is especially clear for (38a–b); Oumar has no way of knowing that Fatima is recalling an overheard conversation in one case and her studies in the other, so their shared representation of the discourse will not reflect this distinction. Similar considerations apply to (38c), as well. While it's reasonable for Oumar to conclude that Fatima's evidence in this case come from the map, Fatima's utterance does not convey this information. This is

apparent from the fact that it is felicitous for Fatima to deny Oumar's conclusion in (38c′) as long as she previously used a non-evidential utterance:

(38) c′. O: Great, last one for now: Togo? I'm totally stumped.

F: (*Giving a sigh of frustration and consulting a map:*)
The capital of Togo is Lomé, (#apparently).

O: That's cheating!

F: (*Holding up the map, which is clearly of France:*)
No, I knew the answer. I'm bored, that's all.

In this example, Oumar's accusation of cheating at the crossword puzzle reveals that he thinks Fatima's evidence comes from the map. Fatima can easily deny this charge, unless she indicated her indirect evidence by means of an evidential marker like *apparently*. Without the evidential marker, Fatima's actions—sighing and looking at the map—can be construed as misleading, but her words cannot. This means that the evidence her utterance invokes must be compatible with having contextual information, but must not require it, on par with (38a–b). If Fatima answers Oumar using an evidential marker, on the other hand, both her words and actions are misleading. Note that the utterance with *apparently* clearly does not pick out the map itself as the source of evidence any more than the non-evidential utterance did. Rather, by characterizing the base, Fatima takes on responsibility for Oumar's potentially faulty inferences about the content of the base, which he must make either way in order to reach arrive at his objection. An evidentially-marked utterance therefore invites the addressee to puzzle out the evidential base's content in a way that a non-evidential utterance does not. This provides further evidence that in the absence of overt indicators of evidentiality, commitments are introduced with an extremely general base.

When an utterance lacks overt evidential content, any commitments that result from the utterance are conditioned on a generic E, the *default evidential base*. Based on the variety in (38), it is reasonable to conclude that this base is comprised of all of the information available to the speaker: his private beliefs especially, as well as any public contextual information. This assumption follows naturally from the expectation that speakers will use the totality of the information available to them to determine the appropriateness of their

discourse moves, unless they explicitly mark their intention to do otherwise.[9]

2.4.1.3 The dependent base

The default base represents a generic body of evidence that can be winnowed down by marked discourse moves. Traditional evidential bases are one way the base can be narrowed, but discourse-specific factors like the source/dependent distinction can be captured analogously. In the dialogue in (39a–c), Algot and Dalia both make commitments about who loves and hates ice cream (bolded), but Dalia's commitments are made as source, while Algot's are as dependent.

(39) *Algot the intern has been tasked with going on a dessert run. He approaches his supervisor Dalia and asks:*

a. A: How about ice cream?

   D: **I love ice cream**.

b. A: Would Fiske be okay with ice cream?

   D: No, **Fiske hates ice cream**. He chucked a pint out a window once.

c. A: Ok, **you love ice cream** and **Fiske hates ice cream**. I'm just going to buy pie.

Recall that committing as source is the default when a speaker makes a commitment. If our notion of the default base already includes the speaker's private beliefs, default commitment as source does not require any special characterization; even though the specific evidence that Dahlia relies on in (39a) derives from introspection and her evidence in (39b) from Fiske's behavior, the default base is broad enough to cover both situations while indicating that Dahlia's commitments are independent of her interlocutor's past or future discourse moves. In other words, the difference in Dahlia's ultimate evidence is not pragmatically or morphologically marked, so the default base suffices. Algot's bolded utterances in (39c), however, are understood as dependent commitments. To verify this, note that replacing *ok* with an overt marker of independent sourcehood such as *that's right* is pragmatically odd, given that one of the commitments involves Dahlia's personal preferences.

[9] Details about the nature of the default base are discussed in depth in Chapter 4.

Dependent commitment to a proposition p is captured in the current model by characterizing the evidential base as crucially containing the *addressee's* commitment, either to p itself or to some closely related proposition.[10] Note that because the existence of a commitment is by definition known to all of the interlocutors in the discourse, the addressee commitment that the dependent's commitment makes salient is part of the contextual information that is covered by the default base. What is unique about the dependent evidential base is that the addressee's commitment is given unusual prominence:

(40) DEPENDENT COMMITMENT

Let q stand for the fact of speaker B's independent commitment to p (or a closely related proposition). When a speaker A commits to p as dependent on speaker B, he conditions his commitment on the base E_{DEP} s.t.:

i. $q \in E_{\text{DEP}}$ and

ii. the speaker could not felicitously commit to p given a base $E_{\text{DEP}'}$ where $E_{\text{DEP}'} = E_{\text{DEP}} - q$.

As indicated by (40ii), the defining feature of the dependent base is that the addressee's commitment is the crucial fact, without which the base would not be viable. In this way, the speaker publicizes that his commitment is contingent on his addressee maintaining that commitment.

2.4.1.4 The weak base

§2.2.2 discussed how the commitment that results from using a rising declaratives like (18) below bundles two distinct requirements.

(18) *Laura has just entered the room, where Max sees her for the first time that day.*

M: You got a haircut?

L: I sure did!

M: I thought so. / # Oh, I had no idea. / # Really? It doesn't look like it.

[10] Alternatively, the base could be construed as crucially containing p itself, toward similar ends. I instead pursue the addressee commitment path because it better highlights that the dependent commitment is tied to the addressee's discourse moves, not merely the information the addressee has conveyed.

First, the speaker presents himself as believing that his addressee is implicitly authoritative about his utterance's content p, meaning that she is expected to be able to settle whether p is true (Gunlogson, 2003). Second, the speaker presents himself as believing that the fact that the addressee is implicitly authoritative matters greatly for the speaker's own ability to remain committed to p. In other words, his commitment is *weak*, because he is less implicitly authoritative about p than his addressee.

When a speaker presents a commitment as being weak, it means he thinks that the relationship between p and the body of evidence that he is able to bring to bear to support p is more tenuous than the bodies of evidence that others might be able to offer. This is clearly Max's intent in (18); he has some evidence that suggests that Laura got a haircut, but he recognizes that his ability to maintain that commitment is beyond his control, because the evidence is insufficient. The general unreliability of the weak base is its defining characteristic:

(41) WEAK COMMITMENT

Let q stand for some proposition that serves as evidence for p. When a speaker weakly commits to p, he conditions his commitment on the base E_{WEAK} s.t.:

i. $q \in E_{\text{WEAK}}$ but

ii. the speaker could not felicitously commit to q (given any base) and

iii. the speaker could not felicitously commit to p given a base $E_{\text{WEAK}'}$ where $E_{\text{WEAK}'} = E_{\text{WEAK}} - q$.

The definition of E_{WEAK} amounts to an admission of insufficient evidence; the base underwrites p, but its viability depends on content that the speaker cannot commit to. The continued viability of the commitment therefore depends on the speaker's interlocutors. Note that the characterization of the base needn't specify that the addressee is authoritative; unless they make contributions that bolster the base, the commitment cannot be trusted. In such cases, we can say that the commitment is *fragile* because the base is *mutable* in ways outside of the speaker's control. This is what makes the weak commitment in (21) so odd:

(21) *Laura has just entered the room, where Max sees her for the first time that day.* ***Laura*** *says:*

L: # I got a haircut?

M: # Yes, (you did)./#Oh.

In this context, there is no reason why Laura's commitment about the state of her own coiffure should be fragile, even in the face of discourse moves that don't confirm it. It would be odd indeed if Max's input constituted a crucial factor that affected whether Laura could safely commit.

The definition of E_{WEAK} is quite similar to that of E_{DEP} in (40); both indicate that the viability of the base rests crucially on a proposition (q) that is not under the speaker's control. The subtle difference between E_{WEAK} and E_{DEP} emerges when we contrast what properties are ascribed to q, and what expectations about the future of the discourse invoking each kind of base projects. The dependent base crucially rests the addressee's commitment to p (or a related proposition). In doing so, it marks the commitment that invokes it as reliant on the addressee's continued commitment to remain solvent. It therefore projects that as long as the addressee's commitment stands, so too will the speaker's. The weak base, on the other hand, rests on a promissory note; the crucial evidence is something the speaker cannot commit to, which serves to acknowledge that the continued viability of the commitment that invokes it depends on this evidence being offered. If no such evidence is found, the commitment will fail.

In a sense, both weak commitment and dependent commitments are 'weak', in that their continued viability rests on the actions of other discourse participants. Weak commitment, however, highlights a much more general kind of weakness because it doesn't specify what kind of information is needed to bolster the base. This leaves room for the bolstering q in (40) to consist of, for example, the addressee clarifying a confusing or conflicting discourse move:

(42) Let p = Todd is allergic to pineapple.

A: I'm going to buy this pineapple for Todd.

a. B: Isn't he allergic? *weak commitment to p*

b. A: Oh, it's for an art project. *evidence supporting p*

c. B: Oh, an art project. Ok. I knew he was allergic. *normal commitment to p*

In this example, B employs a *high negation polar question* to suggest that A's purchase conflicts with her belief that Todd is allergic to pineapples. When A clarifies this point, B

is able to 'rescue' her commitment, but the commitment is crucially *not* dependent on A. Rather than any notion of dependency, what B's question conveys is *speaker bias*, which may or may not result in dependent commitment. Chapter 4 argues at length that various non-default question forms like the high negation polar question in (42) convey different forms of speaker bias for their answers, and that weak commitment is the appropriate way to capture the phenomenon.

It is possible, however, for a commitment to be conditioned on a base that is both dependent and weak. This is actually what we see in the case of rising declaratives; they cannot last alone (in general), and must be bolstered by the addressee's ratifying commitment (in particular). A fuller example of this appears in the subsection below. Chapter 3 presents another case where an utterance employs a base that collects multiple, compatible characterizations by stacking discourse particles.

Before moving on, however, it is worth noting that the expanded definition of commitment sketched above is not by any means radical. For F&R, for instance, truly understanding the nature of a given commitment involves knowing both its propositional content and its location within the partitioned structure of DC_X. The novel formalism here replaces this geographic information with a more flexible, albeit more abstract representation, which I motivate further in the coming sections.

The information encoded here in the evidential base could arguably be presented in a format more recognizable from the perspective of $F\&R$ by 'tagging' propositions with E values so that DC remains at its core a set of propositions, rather than moving to a conception of DC as a set of more complex objects (i.e., as p^E, rather than $\langle p, E \rangle$). This formulation would have the advantage that deriving a speaker's *commitment set* $\bigcap DC_X$, which describes the set of worlds the speaker entertains, would remain unchanged from earlier work. Weak commitment makes this proposition-centric formulation problematic, however, because it is possible to hold multiple weak commitments that conflict, so long as the bases underwriting each commitment are distinct. This is in fact what appears to be the case in (42); the speaker is weakly committed to Todd's pineapple allergy given what she knows, and weakly committed to the opposite given what A has said and done.[11] Culling from DC_X the full set of propositions that a speaker has committed to for any reason and

[11] See Chapter 4 for arguments to this effect.

at any strength is therefore not a useful conception; if a speaker holds conflicting weak commitments, the resulting commitment set would be empty. Because commitments are relevant for future discourse calculations only in virtue the evidential bases that underwrite them, I refrain from giving the propositional content formal primacy as a reminder of this fact.

Finally, recall the thoughts of (Harnish, 2005, p. 38) on what it means to express commitment:

> Expressing might come in degrees in virtue of the fact that its analysis is in terms of providing ***reasons***, and reasons can come in degrees, or at least **levels** e.g. good reasons, weak reasons, (and maybe neutral reasons?). [Bolding added -ON]

The model presented here captures exactly this insight: Commitments are bundled with an evidential base that captures the "reasons" that underwrite them, and whether that those reasons are "weak" or not.[12]

2.4.2 Examples

There is still much to discuss, but with the basic architecture in place, this section pauses to demonstrate how the discourse model captures the basic data from Chapter 1 and earlier in Chapter 2.

2.4.2.1 Default assertion

Consider the simple declarative sentence below, used to perform an assertion:

(43) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a pineapple.*

R: Ahmed likes dragon fruit.

[12]Harnish also mentions "neutral" reasons, which can be understood two ways. First, he might intend that there is an intermediate level between good and weak reasons on a scale. More intriguingly, neutral reasons might be those that the speaker refrains from evaluating, as a way of making a commitment of negligible strength. Expressions of "pure conjecture" in Yucatec Maya (AnderBois, p.c.) and Cuzco Quechua (Faller 2002) suggest this might be useful kind of evidential base to have, although I will not pursue the matter here.

When Raj utters this sentence, he immediately imposes two changes on the discourse structure:

(44) DISCOURSE EFFECTS OF (43)

	Before	→	*After*
Table:	-		p
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEFAULT}} \rangle\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$

First, he places an issue on the Table, containing the propositional content of the utterance. In this case, this is the proposition that Ahmed likes dragon fruit (henceforth p), in virtue of the declarative form Raj employs. Concurrently, Raj commits to p, which requires an evidential base (E). The previous subsection introduced a number of ways that the evidential base can be modified: a base can invoke a particular evidence source, it can signal dependence, and it can be highly mutable, or 'weak'. Further characterizations will appear in Chapters 3–5. A simple default assertions like (43), however, does not impose any of these special restrictions on the evidence that underwrites the commitment, which means that the default base—Raj's private beliefs and any contextually salient evidence—is invoked. The base is neither weak, nor is it automatically dependent.

Note again that if Pepper had previously committed to p, as in (45), Raj's commitment can be either dependent (45a) or not (45b):

(45) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a dragon fruit.*

P: Ahmed likes these.

R: a. Oh, Ahmed likes dragon fruit. Interesting.

b. Yes, Ahmed likes dragon fruit. Interesting.

The choice of response particle signals this difference. In particular, *oh*, as the dependent marker, indicates that the base underwriting Raj's commitment to p crucially references Pepper's previous commitment (indicated by the dagger):

(46) DISCOURSE EFFECTS OF (45a)

	After P's utt.	→	*After R's utt.*
Table:	p		-
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEP}^\dagger}\rangle\}$
DC_P:	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle\}$		$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle^\dagger\}$

The *yes* response in (45b), on the other hand, imposes no such requirement:

(47) DISCOURSE EFFECTS OF (45b)

	After P's utt.	→	*After R's utt.*
Table:	p		-
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle\}$
DC_P:	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle\}$		$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle\}$

There are two further points illustrated above that bare mentioning. First, note that Pepper raised p as an issue (just as Raj had in (43)). Raj's utterance should do the same, but because every member of the discourse is committed to p by the end of his discourse move, the issue is already resolved. Second, (47) uses E_{DEFAULT} for both participants' commitments. This does not mean that the content of the bases is the same, merely that they use the same characterization. Raj's and Pepper's default bases do overlap in terms of contextual information, but their private beliefs likely vary.

2.4.2.2 Rising declarative

To see how a non-default assertion changes the commitment involved, we need only look at an utterance that employs a rising declarative:

(48) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a dragon fruit.*

R: Ahmed likes dragon fruit?

When Raj utters this sentence, its declarative form contributes the same components to its overall discourse effect as it did in the previous example: Raj raises the issue of p and commits to it, given some E. The utterance includes no markers of the source of evidentiality, so again the default base is employed. By virtue of its rising intonation, however, the default base is further characterized as both dependent and weak.

(49) DISCOURSE EFFECTS OF (48)

	Before	→	*After*
Table:	-		p
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEP+WEAK}} \rangle\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$

The notation E_{C1+C2} indicates that the base is characterized as both a E_{C1} base and a E_{C2} base. In this case, the $E_{\text{DEP+WEAK}}$ means that Raj's commitment to p crucially requires that Pepper is committed to p or a closely related proposition, and that unless further evidence can be brought to bear to enhance the base, Raj's commitment will not survive. These requirements are brought together by the fact that unlike (45a), Pepper has yet to make a commitment to p. Raj's utterance therefore anticipates that Pepper will make this commitment in the process of settling the issue Raj has raised, at which point he will be dependently committed:

(50) DISCOURSE EFFECTS OF (48), FOLLOWED BY A POSITIVE REPLY

	Start	→	*After R's utt.*	→	*After P's positive reply*
Table:	-		p		-
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEP+WEAK}} \rangle\}$		$\{\ldots, \langle p, E_{\text{DEP}^\dagger\text{+WEAK}} \rangle\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEFAULT}} \rangle^\dagger\}$

Note, however, that P's reply does not transform Raj's commitment, as shown previously by the fact that he can still object, even after the conditions imposed by his evidential base have been met:

(51) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a dragon fruit.*

R: Ahmed likes dragon fruit?

P: Yes

R: Are you sure? Do you have proof?

If Raj wishes to accept Pepper's response, he must do so overly, thereby making a new commitment that replaces his original, weak one. He is constrained, however, in that the new commitment must not invalidate his old one, meaning than a non-dependent answer like (51b) is infelicitous:

(52) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a dragon fruit.*

R: Ahmed likes dragon fruit?

P: Yes

R: a. Oh, ok.

b. # Yes, that's what I thought.

In terms of the effects these answers have on the discourse structure, what this means is that once Raj has performed an utterance than commits him to being dependent in the future, he must follow through with that dependence:

(53) DISCOURSE EFFECTS OF (52a), FOLLOWED BY A POSITIVE REPLY AND RAJ'S ACCEPTANCE

	After R's utt. →	*After P's positive reply* →	*After R's dep. acceptance*
Table:	p	-	-
DC_R:	$\{\ldots, \langle p, E_{\text{DEP}+\text{WEAK}}\rangle\}$	$\{\ldots, \langle p, E_{\text{DEP}^\dagger+\text{WEAK}}\rangle\}$	$\{\ldots, \langle p, E_{\text{DEP}^\dagger}\rangle\}$
DC_P:	$\{\ldots\}$	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle^\dagger\}$	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle^\dagger\}$

If Raj instead tries to commit independently as in (52b), he ignores the evidence conditions he invoked with the rising declarative and reveals that he misrepresented himself.

(54) DISCOURSE EFFECTS OF (52b), FOLLOWED BY A POSITIVE REPLY AND RAJ'S FAILED ACCEPTANCE (AS SOURCE)

	After R's utt. →	*After P's positive reply* →	*After R's acceptance*
Table:	p	-	-
DC_R:	$\{\ldots, \langle p, E_{\text{DEP}+\text{WEAK}}\rangle\}$	$\{\ldots, \langle p, E_{\text{DEP}^\dagger+\text{WEAK}}\rangle\}$	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle\}$
DC_P:	$\{\ldots\}$	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle^\dagger\}$	$\{\ldots, \langle p, E_{\text{DEFAULT}}\rangle^{\dagger??}\}$

Finally, recall that when a speaker utters a rising declarative, it actively constrains the *addressee's* licit responses as well:

(20) *Max points at something at fruit stand where Laura often shops.*

M: That's a persimmon?

L: Yes, (it is)./#Oh.

Here, Laura's infelicitous *oh* response attempts to adopt the proposition that the fruit is a persimon based on Max's commitment. In terms of this model, this faulty dialogue can be represented in this way:

(55) DISCOURSE EFFECTS OF (20)

	Start	→	*After M's utt.*	→	*After L #acceptance*
Table:	-		p		?
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEP}+\text{WEAK}} \rangle\}$		$\{\ldots, \langle p, E_{\text{DEP}^\dagger+\text{WEAK}} \rangle^\ddagger\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{DEP}^\ddagger} \rangle^\dagger\}$

The final state of the discourse structure reveals a paradox: Max's commitment depends on Laura's committment, and Laura's depends on Max's. In failing to commit using independent evidence, Laura causes Max's evidence to become unviable. It is therefore simply not possible for Laura to commit on Max's say-so—even as a means of indirectly indicating that Max's expectations about Laura are wrong.

2.4.2.3 Evidentially-marked assertion

Returning to Raj and Pepper at the market, let us say that instead of any of the utterances above, Raj utters an evidentially-marked version of the assertion in (43) again employing a falling declarative form:

(56) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Pepper picks up a pineapple.*

R: I hear Ahmed likes dragon fruit.

This utterance is identical to the original default assertion, with the exception that the evidential marker *I hear* indicates that Raj's commitment will be based on indirect or hearsay evidence, rather than his default base:

(57) DISCOURSE EFFECTS OF (56)

	Before	→	*After*
Table:	-		p
DC_R:	$\{\ldots\}$		$\{\ldots, \langle p, E_{\text{HEARSAY}} \rangle\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$

Note that it follows from this proposal that it will not be possible to make a dependent commitment based on one's own hearsay. Such a base, $E_{\text{HEARSAY+DEP}}$ would have to characterized as both a body of hearsay evidence, and crucially reliant on a commitment of the addressee's. Commitments in the current discourse are not hearsay, so no such base is possible.

2.4.2.4 Default polar question

The effects of default and non-default polar questions form the bulk of the data in Chapter 4, but it is worthwhile to briefly introduce how a default, informaiton-seeking question differs from the assertions above. See also Chapter 3 for a discussion of rhetorical questions.

(58) *Pepper and Raj are at the market, shopping for a (single) snack for Ahmed. Raj spots a nice looking dragon fruit.*
Raj: Does Ahmed like dragon fruit?

When Raj utters this question, he again places an issue on the Table containing the propositional content of his utterance. In this case, that issue is a set containing both the proposition that Ahmed likes dragon fruit ($= p$) and its negation ($= \neg p$), as specified by the interrogative sentence's form. In placing an issue with more than one proposition on the Table, Raj invites Pepper not to simply accept the set of answers, but to help settle the issue in favor of one resolution or the other. At the same time, Raj makes a commitment, although it is one that is easy to overlook: If his question is honest, he must commit to the possibility of either answer. In other words, Raj makes a trivial commitment to $(p \vee \neg p)$. Like a default assertion, default question introduces no special characterization of the evidential base. Therefore, this trivial commitment underwritten by Raj's default base:

(59) DISCOURSE EFFECTS OF (58)

	Before	→	*After*
Table:	-		$\{p, \neg p\}$
DC_R:	$\{\ldots\}$		$\{\ldots, \langle [p \vee \neg p], E_{\text{DEFAULT}} \rangle\}$
DC_P:	$\{\ldots\}$		$\{\ldots\}$

After this point, Pepper can commit to either answer, which sets Raj up to make a dependent commitment to settle the issue, or object if he sees fit.

2.4.3 Summary

The discourse model introduces in this section organizes discourse into a shared Table raising and settling issues, and a set of discourse commitments DC_X for each participant X in the discourse. Unlike the F&R model on which the present one is based, DC_X is not a set of propositions, and it is not partitioned in terms of Gunlogson's (2008) sourcehood distinction or the conditionality of commitments. Instead, commitments themselves are pairs of a proposition and the evidential base that supports it, which captures the data that F&R capture while also drawing formal parallels to illocutionary evidentiality. Using the system for simple utterances requires assuming that there is a default evidential base that is employed when an utterance does not specify its characterization. Multiple characterizations can be imposed on a single base as long as they are compatible.

2.5 Further evidence and issues

With the model established, this section provides further support for particular view of evidentiality presented here, as well as the idea that connecting commitment variables like sourcehood with evidentiality is a natural move.

The strongest argument in favor drawing evidentiality and discourse distinctions like sourcehood together is the simple fact that evidentiality already covers a homegenous group of factors, and yet some logical possibilities are unattested cross-linguistically. The model put forward in this chapter does not stretch the umbrella of evidentiality much further, and even predicts its edges: I argue that the utility of illocutionary evidentiality lies in its ability to facilitate the management of commitments. It does this by encoding only those differences in the provenance of information that speakers can reliably use to evaluate the quality of the marked commitments.

To fully explore this characterization, this section then shows how illocutionary evidentiality is related to, but distinct from, issues in information management including hedging and epistemic modality.

2.5.1 Attested and unattested bases

One of the great insights of Faller (2002) is that evidential expressions do not fall neatly into a single hierarchy. Instead, there are at least two distinct criteria involved:

(60) 1. **Amount of inference:**
visual > auditory > other sensory > inference from results reasoning
2. **Number of intervening speakers:**
direct > secondhand > thirdhand > hearsay/folklore

Evidential markers differ in the amount of inference that is required on the part of the speaker, as well as the length of the grapevine between the origin of the information and the speaker. These two lines are distinct, meaning that inferential and reportative evidentials are generally not ordered with regard to one another, even in languages that otherwise exhibit a strict evidential hierarchy. The expansions likened to evidentiality in this chapter, dependent marking and weak-commitment marking, can also be construed in terms of inference strength and the number of intervening speakers:

(61) 1. **Amount of inference:**
strong evidence (default) > weak evidence
2. **Number of intervening speakers:**
independent source (default) > dependent

Making a commitment based on weak evidence requires working from a less reliable (and likely smaller) body of information than making a commitment based on strong evidence does, in the same way that reasoning from visual evidence requires less inference than working from pure reason. Similarly, making a dependent commitment quite literally means that the addressee forms the link that connects the speaker and the proposition he commits to.

Common to all of the above scales is that each scale runs from a highly reliable source of information, to significantly less reliable ones, but the scales do not extend to any completely unreliable sources. This is because all known evidentials carry a doxastic flavor; they encode information about how a speaker comes to believe what he does. If evidentiality were simply a matter of introducing a source of evidence with no question of how that evidence affects reliability, we might also expect to see counterfactual evidentials–expressions

that indicate a particular proposition is impossible or unlikely, based on some false knowledge of them:

(62) UNATTESTED: Percy wrote a book [false-rumors-ev]

This admittedly odd notion is nonetheless coherent, unless we grant that evidentiality serves a more specific purpose than merely marking the evidential source of information. Instead, evidentiality, like sourcehood, serves to modify commitments in predictable ways so that interlocutors can keep track of the provenance of the information they share.

This view is further supported by the fact that cross-linguistically, grammaticalized evidential systems exhibit only coarse distinctions like those on the scales above. We can entertain a simpler notion of the purpose of evidentiality, where all that an evidential expression contributes is a backgrounded assumption about the existence of evidence of a particular sort. Given such a system, we would in principal expect a language to allow with the following evidentials:

(63) a. ATTESTED: Percy wrote a book [hearsay-ev]

b. UNATTESTED: Percy wrote a book [hearsay-from-family-ev]

Many languages have lexical items that correspond to a hearsay evidential, but none have a hearsay evidential specifically for evidence that came from the speaker's family. In general terms, it seems that attested evidential bases have a "unitary character," which dictates that they are general and that they capture distinctions in evidence big enough to reliably affect the strength rating of the commitments they underwrite.

Of course, this property is not particular to evidential systems. Rather, it hints at a larger truth about the purpose of grammaticalization, namely that languages grammaticalize concepts that are relatively frequent and general enough to be useful for communication.[13] Of course, languages have many strategies to get around this lexical discrimination. The simplest strategy is to overtly mention the base:

(64) a. [Word is] Percy wrote a book.

[13]Donka Farkas (p.c.) suggests that morphological tense marking exhibits the same bias toward common, useful distinctions. Languages have no tense specifically for 'yesterday' or 'the day before yesterday', because the combination of a grammaticalized past tense and some lexical strategies is more useful.

b. [According to my aunt], Percy wrote a book.

Alternatively, it is possible to use a less-specific base, and rely on others to accommodate its content:

(65) I just saw my aunt. [Word is] Percy wrote a book.

This level of granularity needn't be reflected in the grammar. By tying the expression of evidentiality to the act of commitment, the space of reasonable evidential markers can be automatically constrained.

2.5.2 Commitment modification versus 'hedging'

Employing the weak evidential base is not the only way to indicate a lack of certainty in one's contributions to a conversation. One extremely common alternative is *hedging*, an example of which appeared at the very start of this chapter:

(14)′ A: I heard that Tal got on a train to Budapest this morning.

B: No, he's flying there.

A: Oh, my mistake.

Unlike illocutionary hearsay markers such as the German *wohl* (see §2.2.2.3), *I heard* is part of the propositional content of the utterance. This is apparent from the fact that the speaker can easily respond to either the embedded clause (above) or the matrix clause:

(66) A: I heard that Tal got on a train to Budapest this morning.

B: No, you read that in an email. / No, you saw him board, remember?

Hedging is a term which has been applied to a number of different constructions and strategies whereby a speaker distances himself from the information he is contributing to the discourse. Hedging is therefore a mostly-descriptive cover term for a number of loosely-related processes, and is useful only at a higher level of description than the model being constructed here seeks to capture.

Itani (1995) argues that hedging has occurred whenever a speaker has communicated "limited commitment" to a proposition, and is a purely pragmatic notion. Consider the examples below, where B's *according to* clause functions as a hedge:

(67) *Travelers A and B are discussing where to visit next.*

A: Is Kyoto a beautiful town?

B: According to what people say, it is. (Itani 1995, p. 23)

(68) I know, from this grey sky, it will rain today, and according to the weather forecast, it will. (Itani 1995, p. 24)

In (67), B's use of *according to* seems both to justify her answer and communicate that she wishes to reduce her "responsibility" (i.e., commitment, possibly) for the proposition that Kyoto is beautiful. One could imagine an analysis for the construction in these terms, where the proposition is given reduced status within B's discourse commitments, but (68) rules out this possibility. Here, the *according to* clause does contribute a justification, but the result (if anything) is a kind of bolstering. Itani (1995) rightly concludes that *according to* therefore contributes exactly what it seems to: simple attribution.

The fact that in some cases attribution seems to weaken responsibility and in other cases bolster it means that hedging is a *post hoc* label we can apply to cases where attribution has a distancing effect, but there is no reliable effect on "commitment" in the sense employed in this dissertation, because the hedged utterance does not cause the effect *by virtue of its form*. Nevertheless, it is common to see hedging referred to as an act of illocutionary modification in the literature. Holmes (1984) famously talks of hedging as one of many ways to "attenuate" or "boost" illocutionary "degrees of strength." This description evokes my discussion of the weak evidential base, but it applies to the intensions behind speakers' utterances, rather than their immediate discourse effects.

The key point of this discussion is to stress that just because an utterance *can* be employed toward particular ends does not mean those ends are encoded in the form that the utterance employs. The next subsection explains this difference in more detail through the lens of epistemic modality. See also Beyssade and Marandin (2006) for a related discussion about *direct* versus *indirect speech acts* and modification of illocutionary force.

2.5.3 Commitment, evidentiality, and epistemic modality

Epistemic modality involves a modulation of the speaker's belief in a proposition that in many ways mirrors the commitment calculations discussed above. The evidential base

in particular evokes the *conversational backgrounds* of Kratzer (1981, 1987). Under this analysis, a proposition *p* embedded under a modal operator introduces a *modal base*, a set of propositions that projects the set of the worlds consistent with the speaker's evidence (usually the propositions that the speaker takes to be true). An *ordering source* is then employed to order the worlds in the modal base in a way that reflects the viability of the worlds given the ordering source's assumptions.

Some attempts have been made to collapse evidentiality and epistemic modality into a single phenomenon. Faller (2011), for example, provides an alternative analysis of the evidential system of Cuzco Quechua that captures evidence source as an evidential base and inferences as an ordering source, and ultimately concluding that evidentials as a whole cannot be subsumed into epistemic modality, although aspects of certain evidentials can.

In terms of the relationship between epistemic modality and commitment, I wish to maintain a similar distance. One strong argument for maintaining this distinction is that there are cases where a single utterance might employ *both* epistemic modality and illocutionary evidentiality:

(69) a. Pilar yachay wasi-pi ka-sha-n-**man**
Pilar know house-LOC be-PROG-3-**man**
p = 'Pilar might be at school.'

b. Pilar yachay wasi-pi-**n**/-**s**/-**chá** ka-sha-n-**man**
Pilar know house-LOC-**n**/-**s**/-**chá** be-PROG-3-**man**
p = 'Pilar might be at school.'
ev = -**n**: Sp has best possible grounds for saying that Pilar might be at school (= inference).
-**si**: Sp was told that Pilar might be at school.
-**chá**: Sp conjectures that Pilar might be at school. Cuzco Quechua; (Faller, 2002, p. 85)

In these examples, the proposition the speaker commits to is modalized. In (69b), three different evidentials are appended, altering the evidential base that underwrites the speaker's commitment, but *not* altering the fact that this commitment is to a modalized proposition. Even though *-n* (an allomorph of *-mi*) is generally associated with strong (default) commitment, commitment to a proposition with a possibility modal is licit.

What this requires is that the evidential base that underwrites the commitment and the modal base that drives the modalized proposition are distinct. In other words, the first example in (69b) exhibits strong commitment given a base B_1 to [the possibility of p, given B_2], where B_1 and B_2 are distinct. Crucially, there is no way to collapse B_2 into B_1 and still retain the intended reading; the speaker's commitment is not weak, even though what he is committed to is. In other words, (69b) is an instance of (70a), and crucially not (70b):

(70) THE SCOPE OF UNCERTAINTY

a. The speaker makes a strong commitment to *might-p*

b. The speaker makes a weak commitment to p

This is the precise distinction discussed in the hedging examples above; even through (70a) and (70b) accomplish the same goal in most cases, they are distinct speech acts that impose different expectations on how the speaker will act with regard to p. Because these 'levels' of conditionalization must be distinguished to capture the data for Cuzco Quechua, they should in principle always be distinct.

2.5.4 Evidence source, or reliability?

It is important to acknowledge that the literature on evidentiality is extremely varied, meaning that the position adopted here is not without its detractors. Under the view advanced here, the purpose of marking illocutionary evidentiality is to indicate the basis for a commitment. Evidential expressions introduce not-at-issue sources of evidence, which, while generally expressed in terms of a commitment to the existence of appropriate evidence, here instead introduces a veiled body of evidence itself called an evidential base (c.f. Aikhenvald 2004, Faller 2002, Murray 2010 and others). This base primarily characterizes the evidence source underwriting the evidential commitment, but in principle there is nothing preventing a language from using a single marker to characterize both an evidence source and a non-default level of mutability that the evidence imparts to the commitments it is used in. The simplest case, making use of base characterizations already introduced, would be an evidential marker that introduces a base that encodes both evidence source but is also weak. It is also reasonable to propose a base that only characterizes the mutability of commitments that invoke it; this is exactly what the weak base does.

Remarkably, some researchers reject the idea that evidentiality has anything to do with evidence source at all, arguing instead that the core meaning of an evidential encodes only a particular strength of evidence (Krawczyk 2009). This runs counter to both the source-only camp (Aikhenvald 2004, Faller 2002) and the combination source/reliability camp (Davis et al. 2007, McCready and Ogata 2007), and in my terms amounts to the claim that the weak base (and potentially other bases like it) are the only ones needed to capture the effects of evidentiality. Upon inspection, however, it seems that what has been uncovered here is not an error, but yet more variation in the ways that evidentiality can work cross-linguistically.

The strongest evidence for the reliability-only view comes from the phenomenon of "evidential promotion", where a speaker can choose to use a different evidential just in case it conveys the level of confidence that he wants to convey. Krawczyk (2009) cites the following example from Central Alaskan Yup'ik (Eskimo-Aleut), cited previously to compare the inferential and reportative markers:

> "Say you and I have a friend, and we are driving to her house to see if she is home. As we approach her house, we cannot tell if she is home or not. I ring her doorbell, and her husband answers and tells me that she has left. Given that he is her husband, he lives with her, and I don't necessarily think he is lying to me, I can use any of the following constructions below when I tell you that she has left." (Krawczyk 2009, p. 6)

(36) a. Aya-llru-uq.
leave-PAST-3RDSG
'She left.' *Unmarked*

b. Aya-llru-llini-uq.
leave-PAST-INFER-3RDSG
'Evidently, she left.' *Inferential evidential*

c. Ayag-ni-llru-a.
leave-said-PAST-3RDSG
'He said that she left.' *Verb 'say'*

d. Aya-llru-uq-gguq.
leave-PAST-3RDSG-HEARSAY
'It is said she left.' *Reportative evidential*

Test for commitment of the usual sort reveal that (a–d) are ordering in decreasing "strength,"

albeit in Holmes's (1984) sense; in particular, the inferential evidential in (36b) presents the speaker as much more certain of the utterance's literal content than the reportative in (36d). In terms of evidential source, (36d) is the appropriate choice given the scenario above. Interestingly though, any of these are possible as long as the speaker trusts his hearsay evidence. The argument therefore is that the so-called 'inferential' evidential is simply a higher-strength option than the 'reportative' one, and the fact that the evidence in this example is hearsay is inconsequential.

On closer inspection, however, this data is not at all problematic for the other views. Hearsay evidence exists at the weak end of Faller's (2002) intervening-speaker scale, and so is canonically associated with a low level of speaker credence in Central Alaskan Yup'ik. In other words, hearsay in this language is inherently untrustworthy. The fact that the speaker trusts his evidence source in the example above is therefore additional information that he can use to infer higher confidence in the proposition that the friend in fact left. Further, presuming that the speaker is aware of the uncertainty that simply using a hearsay evidential marker would convey, the speaker not only can, but *should* condition his commitment on the more canonically reliable inferential base instead.[14]

A more notable case comes from Cuzco Quechua (Amarind), as investigated by Faller (2002), which famously offers its speakers use of the *mi* evidential to mark that the speaker has "best possible grounds" for his utterance, repeated below:

(69′) Pilar yachay wasi-pi-**n** ka-sha-n-**man**
Pilar know house-LOC-**n** be-PROG-3-**man**
p = 'Pilar might be at school.'
ev = Sp has best possible grounds for saying that Pilar might be at school (= inference). Cuzco Quechua; (Faller, 2002, p. 85)

This clearly does seem to be a case where the evidence source itself is less important than its effect on mutability. Under the model advanced here, this is expected; again, there is no prohibition against evidentials that under-specify one or the other components of evidentiality.[15] The immediate question becomes: What other characterizations of commitment

[14]One prediction that Krawczyk (2009) makes and the above solution does not is that "evidential demotion" would be available as well; given a case of pure inference, if a speaker can use the hearsay evidential to show he doesn't trust his reasoning, it would disprove my analysis.

[15]That said, (Faller, 2012, p. 300) mentions a situation similar to the Yup'ik example and seems to indicate

mutability are attested, and how should they be captured? For the Cuzco Quechua *-mi* marker, for example, we can easily imagine a converse of the weak base which characterizes the evidence underwriting a commitment as basically immutable; rather than relying on corroboration from the addressee, the base would be resistant against any evidence the addressee could offer. The analysis of Japanese *yo* in Chapter 3 offers one possible route toward such an analysis.

2.6 Conclusions and next steps

This chapter has presented the basic features of the evidentiality-centric discourse model that is the project of this dissertation. The discourse-sourcing and evidence-sourcing notions introduced in §2–3 are subsumed into a general evidential base, which has certain default properties that can be modified or overwritten by specific expressions, with reliable effects on the discourse structure. The model is general enough that it does not impose and particular views of the nature of the relationship between the evidential base that underwrites a commitment and the strength of that commitment, although this remains an issue worth of a separate investigation. The model also exhibits appealing parallels with the study of epistemic modality, though I refrain from forcing any particular interpretation of those parallels.

The model introduced here still requires a great deal of nuancing, which can only come from putting it to task explaining patterns of natural linguistic data. Chapter 3 beings this process with an exploration of Japanese sentence final discourse particles, which in addition to being evidential-like in their distribution, perform a function similar to familiar sourcehood marking. Chapter 4 then turns to the issue of biased questions in English in order to explore the nature of the default evidential base, as well as the idea that the weak base is used to express speaker bias.

---

that evidential promotion is *not* possible when it is understood that direct, visual evidence could in principle be available (counter to Krawczyk (2009)):

> [I]f I was told by a trustworthy source that it is raining in Lima and I am completely convinced that this is true, I can still not use the Direct [*-mi* form.]

Chapter 3

Japanese discourse particles as relative discourse evidentials

3.1 Introduction

Discourse particles in Japanese form a closed class of short words that can occur at the far right edge of an utterance. They are notable because rather than affecting the propositional content of an utterance, these particles encode information about how the speaker's relationship with a proposition compares to his interlocutors' relationships with it. These facts are analyzed here in within the framework established in Chapter 2, as part of a further exploration into how evidential bases can vary. The result is a new pair of evidential bases, whose simple characterizations as markers of a speaker's relative authority successfully capture the particles' complex discourse effects better than previous approaches.

By far the most widely used discourse particles in Japanese are *yo* and *ne*, which are generally believed to mark information as "hearer-new" and "hearer-old," respectively, as shown below (from McCready 2008b).

(71) Ame-ga futteiru (yo/ne).
rain-NOM falling YO/NE
'It's raining, (man/huh).'[1]

[1]Most translations are my own, and attempt to mirror to the greatest extent possible the pragmatics of each utterance. When choosing between a translation more faithful to the literal Japanese and one appropriate for an identical context in English, I will prefer the latter. Note however that many of the example sentences

In truth, these particles are a good deal more complicated than the pithy quotations above suggest; there is a sizable body of literature—most of it in Japanese—about the use and meanings of *yo* and *ne*, but very little of it is formal beyond recent work by McCready (2008b), Davis (2009, 2011), and the authors they cite. Broadly, this line of research suggests that *yo* and *ne* are best understood as encoding information about the relevance of the speaker's utterance to the addressee's goals, and that the apparent hearer-new vs. hearer-old distinction is a derived effect. While relevance is clearly a key component of the meanings of these particles, I argue here based on novel data that relevance is not enough to fully explain their behavior.

With this in mind, the immediate goal of this chapter is twofold. First, I offer an overview of *yo* and *ne*'s use, with particular focus on a few data points that are not discussed by McCready (2008b) or Davis (2009, 2011). I do this by contrasting *yo*, *ne*, and their perennially understudied combination, *yone*, whose effect I argue is compositional. My second goal is to offer an analysis of this dataset within the framework developed in the previous chapter.

In the process of accomplishing the second goal, this chapter will expand and nuance the model of discourse that is the central project of this dissertation. While Chapter 2 argued that previously-observed phenomena like illocutionary evidentiality, conditional commitment, and Gunlogson's (2008) sourcehood distinction can all be captured as a family of commitment modifiers, expanding the model to account for the Japanese particle data allows for a more thorough exploration of what kinds of content *should* be formalized this way. This chapter therefore lends further support to the main points of Chapter 2. It also takes further steps toward constraining the sorts of distinctions that can be reflected in the evidential base that underwrites a commitment.

The theoretical contribution of this investigation is the introduction of a novel discourse notion I term *relative authority*. A discourse participant's relative authority over a proposition tracks how qualified that participant is to vouch for the truth of the proposition, compared to his interlocutors. Expressions that mark relative authority make explicit the speaker's assumptions about how his epistemic or deontic authority to sponsor a particular addition to the context compares to that of his interlocutors, constraining future discourse

themselves come from McCready (2008b) and/or Davis (2009, 2011), and the use of *man* in the translation of examples involving *yo* is directly due to McCready (2008b).

moves much like the phenomena in Chapter 2. In terms of the model developed here, it will be argued that the essential function of the particles *yo* and *ne* is to introduce special discourse-evidential bases whose characters reference relative authority. The contributions of these particles overlap with their more-constrained combination *yone*, driving a pragmatic competition similar to the sort found in languages with robust evidential systems (Faller, 2012).

The chapter proceeds as follows: §3.2 discussed the distribution of *yo*, *ne*, and *yone*, including previous authors' observations and my own novel ones. §3.3 presents a brief summary of an existing formal treatment of these facts, drawing from McCready (2008b) and Davis (2011). §3.4 summarizes the model of discourse as developed in Chapter 2, with an eye to the components that remain 'fuzzy' and that the present data can sharpen. It also offers a few extensions, including an expansion to account for imperatives, based on Portner (2007) and Farkas and Bruce (2010).

§3.5 then presents some preliminary observations about *yo* and *ne* within this framework in terms of sourcehood. §3.6 formally define *relative* versus *absolute authority*, and introduces my novel analysis of these particles. §3.7 argues for the retention of Davis's (2011) analysis of the intonational morphemes that co-occur with *yo*. Before concluding, §3.8 offers an extension to the analysis by way of comparison with the Singlish particle *lah*, which overlaps in important ways with *ne*. Finally, §3.9 offers conclusions and raises questions to be addressed in the next chapter.

3.2 Japanese discourse particles

This section provides data on felicitous and infelicitous use of the final particles *yo* and *ne* with declaratives, imperatives, and interrogatives (and by extension, their associated canonical speech acts).[2] It is worth noting that there is a great deal of variation in how these

[2]There are many particles which will receive no attention here, including *zo(o)*, *zee*, 'feminine' *wa*, and various regional particles such as Kansai-dialect *wa*. Each of these is unique and complex, and must be taken into account when developing a comprehensive analysis of discourse particles in Japanese. This is especially important under the approach developed here, given that evidentials appear to compete. (See §3.6.4.1.) Because the focus of this chapter is on developing my model of discourse, I choose to focus in on just *yo* and *ne* and leave the comprehensive analysis of Japanese for future work.

particles are used, even between dialects, and so these characterizations are not intended to be exhaustive. The most accessible and complete characterization to date appears in Saigo (2011).

3.2.1 A simplified picture of *yo*

Sentence-final *yo* occurs with two different intonational patterns, as described by Davis (2009, 2011). Before considering the differences between these patterns, though, it is useful to first sketch the generalizations common to both. This characterization is due to McCready (2008b).

Regardless of its intonation, *yo* is generally taken to mark an utterance as containing hearer-new information. Also, this information must be somehow relevant to the addressee's goals:

(72) *Souta sees Ayaka hasn't noticed that her train has arrived.*

S: Densya kita #(yo).
train came #(YO)
'Your train is here!'

(73) *Souta sees Ayaka at a busy train station, but he knows nothing of her plans.*

S: Densya kita (?yo).
train came (?YO)
'A train has arrived.'

(72) is therefore felicitous only in situations where the speaker, Souta, believes[3] that his addressee, Ayaka, doesn't know that her train has arrived, but where it matters to Ayaka whether it has.[4] (73), by contrast, is peculiar with *yo*, because the information isn't relevant. Generalizing, we could propose the following as a good first approximation of *yo*'s behavior:

[3]Throughout, I make the simplifying rhetorical assumption that discourse involves speaker belief, but it would be more accurate to talk about what the speaker *presents himself as believing*; as discussed in Chapter 1, it does not matter what the speaker's beliefs truly are, only what he presents to his interlocutors for the purposes of the current discourse.

[4]For ease of exposition, most examples in this chapter feature two participants, Souta (male) and Ayaka (female). Their friends Taro (male) and Hanako (female) make guest appearances.

(74) INITIAL EMPIRICAL GENERALIZATION FOR *yo*

An utterance expressing a proposition *p* is marked with *yo* iff the speaker believes *p* is **informative** and **relevant**:

i. *p* is informative iff *p* is new information for his addressee.

ii. *p* is relevant iff the addressee cares whether *p* is true.

And indeed, this characterization is the starting point for many analyses of the particle. Before simply accepting it, note that this initial empirical generalization predicts that *yo* will be infelicitous whenever *p* is shared knowledge; no matter how relevant the proposition, if the addressee already knows it, the condition in (74i) will not be met. Examples like (75) below, however, reveal that this prohibition is false, and that the relevant notion is not informativity, but evidence. In particular, *yo* is felicitous in situations where the speaker and addressee share a belief in the truth of a particular proposition, as long as the speaker's evidence is stronger:

(75) *Souta and Ayaka both know that Hanako arrived in town today. Souta knows because he saw her a few minutes ago. Ayaka knows because someone told her so.*

A: Hanako-ga tuita tte.
Hanako-NOM came QUOTE
'I hear Hanako has arrived.'

S: (N.) Hontou da yo.
(Yeah.) true COP YO
'Yeah, that's right.'

In this example, both Souta and Ayaka know that Hanako is in town, but Ayaka has indicated she knows this only indirectly. Because Souta has direct evidence of Hanako's arrival, he may use *yo*. If *yo* were prohibited with shared knowledge, (i.e., if *yo* required true hearer-novelty), this example would be infelicitous. Because of this, no account that relies on an absolute hearer-newness requirement for *yo* will be able to account for (75). This observation, which I will abbreviate as the *weak informativity* characteristic of *yo*, drives much of the analysis presented in the latter half of this chapter.

Note that the use of *yo* is infelicitous if the context in (75) is reversed, so that Souta has (weaker) indirect evidence and Ayaka has direct evidence. This confirms that it is not simply inequality of evidence, but inequality *in the speaker's favor* that licenses *yo*.

In (75), the speaker's and addressee's evidence were disjoint, but this is not a requirement for *yo* either. Consider:

(76) *The 'check engine' light in Ayaka's car is on, leading her to suspect engine trouble. Souta is a mechanic who both sees the light and inspects the engine, and determines there is trouble.*

S: (Yappari) mondai atta yo.
(as-suspected) problem was YO
'Yup, there's a problem.'

Here, Souta and Ayaka have both seen the 'check engine' light, but in addition, Souta has looked at the engine himself. His evidence that the car is having engine trouble is therefore a proper superset of Ayaka's, and *yo* is licit. Based on this discussion, we can replace the generalization in (74) with:

(77) REVISED EMPIRICAL GENERALIZATION FOR *yo*

An utterance expressing a proposition p is marked with *yo* iff the speaker believes p is **weakly informative** and **relevant**:

i. p is weakly informative iff the speaker has stronger evidence for p than the addressee.

ii. p is relevant iff the addressee cares whether p is true.

This generalization captures both weak informativity cases like (75) and hearer-novelty cases like (72).

Native speakers often consider a *yo*-utterance to be more "forceful" than its *yo*-less counterpart, and sometimes manipulate conversations to avoid using it. That said, there are many cases where *yo* becomes required, in particular when the speaker is actively trying to convince the hearer of something, as in (78):

(78) *Ayaka and Souta have conflicting information about where Hanako is.*

A: Hanako, kaeranakatta no ni...
Hanako didn't go home although
'Even though Hanako didn't go home...'

S: Kaetta #(yo)!
went home YO

'She DID go home!'

Here, Souta is issuing a correction, which makes *yo* required. Leaving it off makes it sound as though Souta's utterance has no bearing on Ayaka's claim, even though pragmatically it seems as though the connection should be obvious.[5] Without *yo*, Ayaka might reasonably infer from this discourse that Souta is not paying attention to what she said.

As shown above, *yo* is required in response to statements the speaker disagrees with (i.e., in corrections and similar moves that risk conversational crisis). Another place *yo* shows up in responses is to questions, as shown below.

(79) *Ayaka wants to know whether Hanako went home.*

A: Hanako, kaetta no?
Hanako went home cop
'So then, did Hanako go home?'

S: N, kaetta #(yo).
(yes) went home YO
'Yeah, she went home.'

yo is generally required in responses to questions. This is unsurprising, given that responding to an information-seeking question is probably the most clear-cut place one would want to provide relevant, novel information to the asker.

So far, all the examples discussed have dealt with declarative sentences, but it is important to note that *yo* occurs with imperatives as well.[6] An example follows in (80), where *yo*'s main contribution is forcefulness (to the point of pleading), as indicated by the translation, '*come on*'.

(80) *Ayaka and Souta are shopping. Ayaka has all the money.*

S: Atarasii geemu kat-te (yo).
new game buy-IMP YO
'(Come on,) buy me a new game.'

[5]This point, that the availability of these particles blocks pragmatic inference, returns in §3.7.1.

[6]As noted by Schmerling (1982); Farkas and Bruce (2010); Condoravdi and Lauer (2012) and others, imperative sentences are used for a wide array of speech acts. In this chapter I will make the simplifying assumption that all imperatives will be associated with what Farkas and Bruce (2010) terms the *core imperative force* speech acts, which includes commands, requests, permission, and advice.

More about imperative use follows in the finer-grained treatment of *yo* in the next section.

So far, we have seen that weak informativity and relevance are the key requirements for felicitous use of *yo*. Further, when these conditions are met, *yo* is not optional; failing to use it when issuing a correction or answering a question often results in infelicity. An analysis's ability to capture this generalization will be a key test of its viability. Further, a successful analysis of *yo* should be able to explain why the particle sounds intuitively 'forceful' in assertions, and 'pleading' in imperative speech acts. Before moving on, however, it is important to note that intonation plays a key role in the felicitous use of *yo*.

3.2.2 *yo* and intonation

Davis (2011) notes that *yo*'s use is complicated by the fact that it occurs with two different intonational patterns, a final rise (notated ⇑) and a final fall (notated ⇓). These two patterns have subtly different properties with regard to what information the speaker believes his addressee is committed to prior to his utterance. In short, rising *yo* is used to inform, while falling *yo* is used to correct. In both cases, the weak informativity requirement in (77i) holds, as shown above. The difference in intonation merely determines what it means for the utterance to be relevant (i.e., how the utterance meets (77ii)). In the analysis to come, I will follow Davis (2011) in treating *yo* and its intonational contours as three distinct morphemes.

Note that this differs from the treatment of rising vs falling intonation as it applies to assertions versus questions (in both Japanese and English). This is because the intonational markers on *yo* are specific to this particle, and distinct from the normal tunes found elsewhere in the language. The problem of intonation's multifaceted role in deriving discourse effects is muted here.

3.2.2.1 The distribution of rising *yo*

With declaratives, an utterance of the form *yo*⇑(φ) (i.e., a rising *yo*-utterance expressing a proposition φ) indicates φ is relevant and 'hearer-new', as discussed above. It is generally required in such cases. For example, in the situation in (81), failure to use *yo*⇑ would make it sound as though the speaker were simply remarking on the presence of the train, rather than attempting to spur the addressee to action.

(81) *Souta sees Ayaka hasn't noticed her train has arrived.*[7]

S: Densya kita #(yo⇑).
train came #(YO⇑)
'The train is here.'

Again, this is notable because even though pragmatically it is very clear what the speaker is attempting to do in either case, the inference is blocked when *yo*⇑ is absent.

With imperatives, *yo*⇑(φ) indicates φ is not obvious or inferable to the addressee, as shown in (82).

(82) *Souta has an earache and Ayaka is about to listen to some classic rock.*

S: Sizuka-ni site kure (yo↑).
Quiet-DAT do.PROG give.IMP (YO↑)
'Stay quiet for me, would ya?'

Also, the command must reflect the speaker's wishes in order to be felicitous with *yo*. This is met in the example above, because the speaker's wish to avoid pain prompts the imperative in the first place. While in (82) Ayaka may not know what exactly prompted Souta's imperative, she knows it has to do with the speaker's wishes (e.g., an earache, a desire for concentration, a dislike of rock music, etc.) and not, for example, an angry landlord. Without *yo*, such speaker-external reasons remain active options, and may actually be preferred.[8]

As evidence for the speaker-centric nature of *yo*-marked imperatives, note that it is impossible to grant permission with *yo*, as shown by the contrast between (83) and (84).

(83) *Souta is trying to settle a dispute among his employees.*

S: Mata nanika attara, soudan ni kite kudasai ?(yo⇑).
again something be-COND consultation for come please ?(YO⇑)
'If anything else happens, please come talk to me again.'

[7] The issue of whether this example shows an effect of relevance or hearer-newness is potentially at play here. One could imagine a slightly different context where Ayaka sees her train has arrived, but fails to recognize that this means it's time for her to board. In such a situation, the *yo*-marked utterance in (81) is dispreferred to an option like *dettyau yo*⇑, 'It's going to leave on you!' where the train's arrival itself is not discussed. This tendency to avoid marking the obvious with *yo*, regardless of how relevant it may be, falls out from the authority component of the final analysis.

[8] See §3.7.1.

(84) *Ayaka knocks on Souta's door.*

S: Haite kudasai (#yo⇑).
come.in please (#YO⇑)
'Come in please.'

In (83), the command to come for another consultation originates with Souta's wishes for harmony in his office, so adding *yo*↑ is preferred, while in (84), Ayaka indicates a desire to enter the room by knocking, making it infelicitous for Souta to tag the imperative with *yo*; to do so would be to imply that it was some internal state of Souta's, rather than Ayaka's knocking, that prompted the imperative.

With questions, *yo*⇑ is infelicitous:

(85) Mada, ame, futteru ka (#yo⇑)?
still rain falling Q (#YO⇑)
'Is it still raining?'

This is not the case with falling *yo*, as will be shown in the next section.

3.2.2.2 The distribution of falling *yo*

yo⇓(φ), on the other hand, is similar to *yo*⇑(φ), but replaces the relevance requirement with the requirement that the addressee is supposed to be somehow misinformed or incorrectly taking something for granted. For this reason, it is only falling *yo* that occurs in corrections:

(86) *Ayaka and Souta have conflicting information about whether the prime minister has died, and Ayaka has just said, "The prime minister died."*

S: Sinde-nai yo⇓/#yo⇑!
died-not YO⇓/#YO⇑
'He did NOT die!'

As in (78), failing to use *yo* at all in (86) would be odd because it would sound as though Souta's utterance had no connection to Ayaka's. Rising *yo*, however, is even worse than no *yo* at all, and simply doesn't make sense for many speakers. Falling *yo* must be used whenever the addressee is incorrect (from the speaker's perspective).

Because the misinformation in both (78) and (86) relates to the polarity of some proposition, it is tempting to construe *yo*⇓ as a *reverse polarity* marker, in terms of Farkas (2011) and others, but as shown in the examples below, there is no requirement that the addressee's misinformation be so directly connected to the speaker's utterance:

(87) *Souta and Ayako know someone has died, and Ayaka has just said, "It was the prime minister who died."*

S: Sinda no-wa Hanako da yo⇓/#yo⇑.
Died one-TOP Hanako COP YO⇓/#YO⇑
'It's Hanako who died.'

(88) *Ayaka has just asked, "Did you eat already?"*

S: Tabeta (yo⇓/yo⇑).
ate YO⇓/YO⇑
'(Yeah), I ate.'

The context for (87) is almost identical to that of (86), except that in the original context, Souta is refuting Ayaka's claim about the prime minister directly, while in (87), he instead gives an alternative that does not simply differ in polarity from Ayaka's utterance. The case in (88) is more complicated because the corrective component of Souta's response is not required by the immediate utterance context, the way it was in the previous example. Because of this, rising *yo* is perfectly acceptable, and using falling *yo* indicates that the speaker objects to the addressee's question. Ayaka might infer from this utterance that she's missed something and it should be obvious whether Souta has eaten, or she might simply infer that Souta doesn't want to be grilled about his eating habits. Either way, falling *yo* indicates misinformation on the part of the addressee, and often displeasure on the part of the speaker.[9]

This characterization also applies to imperatives. An imperative of the form *yo*⇓(φ) is very similar to an imperative with rising *yo*, except that it also indicates that the hearer is already failing the command or request expressed by φ. McCready (2008b) provides the example below:

[9]Falling *yo* rarely occurs when issuing corrections in teaching contexts, most likely because it is so easy to interpret it as an indicator of being disgruntled.

(89) *The speaker is repeating a request that Godai (the addressee) denied.* (from *Maison Ikkoku*)

Sonna... Godai-san tuiteite oage-nasai yo⇓/#yo⇑.
that.kind.of Godai-san go.with HON-please YO⇓/#YO⇑

'Hey now—Godai, go with her.'

Here, only falling *yo* is acceptable, because Godai has already refused once, and hence is not acting in accordance with the speaker's wishes. Returning to the classic rock example in (82), falling *yo* would be the only option if Ayaka had already turned on the radio before Souta's *yo*-marked request for quiet.

Unlike rising *yo*, falling *yo* may occur with questions, as long as they are rhetorical, as shown below by the contrast between the regular question in (90) and the rhetorical one in (91).[10]

(90) Konna hon, kau (ka) (*yo⇓)?
this.kind.of book buy Q (*YO⇓)
'Are you going to buy a book like this?'

(91) Konna hon, dare-ga kau ka (yo⇓).
this.kind.of book who-NOM buy Q (YO⇓)
'Who the hell would buy a book like this!?'

yo⇓(φ) sounds snide and sarcastic compared to its *yo*-less counterpart. Also, mirroring the generalizations about imperatives, rhetorical questions with *yo* are more speaker-oriented than those without; in many contexts, (91) might be better glossed as 'Who the hell would buy a book like this!? *I* sure wouldn't'.

(91) seems to imply a negative response (e.g., 'No one would'). It is difficult to find felicitous examples of *yo* occurring with rhetorical questions implying positive answers (e.g. 'Is the Pope Catholic?', 'Don't you know who I am?'), this may have more to do with the relative difficulty of using rhetorical questions for these cases even without *yo*; Japanese

[10]The term *rhetorical question* can refer to a couple different phenomena, which differ as to who answers the question (if anyone). All of the examples discussed here are of the obvious-answer type, where the answer is assumed to be known to all the discourse participants. Self-answered rhetorical questions ('Who's in charge here? I am.') do not occur with *yo*. See Chapter 4 for a fuller discussion of rhetorical questions.

speakers tend to prefer other, arguably more directly cooperative expressions such as *mot-iron* ('of course').

That said, it is clear that *yo* is fine with rhetorical questions that are non-negative:

(92) Ore-wa dare da to omotteiru n da (yo⇓).
1st.macho-TOP who COP C thinking N COP (YO⇓)
'Who the hell do you think I am!?'

In this example, the implied answer is something like 'no one', or 'someone unimportant'; in other words, while the rhetorical question isn't negative in polarity, it still seems to imply something negative (in the sense of 'bad') about how the addressee is treating the speaker. As a final example, consider (93) below, (based on an example from Han (2002)):

(93) Dare-ga unda to omotteiru n da (yo⇓)?
who-NOM gave-birth C thinking N COP (YO⇓)
'Who the hell do you think give birth to you?'
(c.f. 'Who do you think's taken care of you all your life?)

This question is again 'negative' in the broad sense, but unlike the previous example, that negativity is confined to the question-asker's attitude. The implied answer to the question itself is simply 'you', although the question suggests that the addressee has been acting as though the answer is 'not you'. In sum, *yo* is licit with at least emotionally negative rhetorical questions, but the use of rhetorical questions in Japanese is generally restricted to these emotionally negative cases, so the negativity does not signal anything particular about *yo*.

3.2.2.3 Interim conclusions about *yo*⇑ / ⇓

So far, this section has presented an overall characterization of *yo*, and then moved to a more nuanced view based on that presented by Davis (2009, 2011) where *yo* has two variants that differ in intonation. The empirical generalizations that an analysis of these facts must satisfy appears below, broken down by sentence form.

(94) FINAL EMPIRICAL GENERALIZATION FOR *yo* WITH DECLARATIVES

a. An utterance expressing a proposition *p* is marked with *yo*⇑ iff the speaker believes *p* is **weakly informative** and **relevant**:

i. *p* is weakly informative iff the speaker has stronger evidence for *p* than his addressee.

ii. *p* is relevant iff the addressee cares whether *p* is true.

b. An utterance expressing a proposition *p* is marked with *yo*⇓ iff the speaker believes *p* is **weakly relevant** and **corrective**:

i. (As above) *p* is weakly informative iff the speaker has stronger evidence for *p* than his addressee.

ii. *p* is corrective iff the addressee is mistaken about *p* or facts surrounding it.

c. *yo* (with either intonation) is **obligatory** when the above conditions are met. Failure to use it indicates **inattention** to the conversation.

With declaratives, it was shown that a speaker can must *yo*⇑ felicitously in cases where he knows he has better evidence than his addressee and his utterance is somehow relevant to his listener. *yo*⇓ carries a similar evidence requirement, but trades rising *yo*'s relevance requirement for a 'corrective' one, whereby the speaker must believe that his addressee is actively wrong about something in the discourse. Further, note that *yo* with either intonation is not optional, and failing to use it signals a lack of engagement in the conversation. Weak informativity (94a/bi) and obligatoriness (94c) are therefore generalizations that can be traced directly to the contribution of *yo*, while the relevance (94aii) and correction requirements (94bii) reflect intonation.

With utterances employing imperatives, we see mainly the effects of intonation:

(95) FINAL EMPIRICAL GENERALIZATION FOR *yo* WITH IMPERATIVES

a. An utterance expressing *p*! is marked with *yo*⇑ iff *p*! **originates with the speaker**.

i. *p*! originates with the speaker iff fulfilling *p*! derives from the speaker's own deontic authority.

b. An utterance expressing *p*! is marked with *yo*⇓ iff the speaker believes the addressee is **failing to heed** *p*!.

i. The addressee is failing to heed *p*! iff he acts in a way that does not help bring *p*! about.

c. Failure to use *yo* (with either intonation) in an appropriate context is infelicitous.

With imperatives, we see significant symmetry, especially in the case of falling *yo*, where the addressee's failure to comply with the command or request (95b) is very similar to the addressee being misinformed (94bii). It is less clear how the generalizations for rising *yo* can be related. Also, *yo* remains obligatory with imperatives (95c), although the result of failing to use it does not convey inattention. Note that *yo*'s most salient effect, weak informativity (94a/bi), has not seen specific support with imperatives, although at this stage it is unclear what the parallel of informativity would be in the domain of the core imperative speech acts.

Finally, *yo* displayed very limited felicity with interrogatives. Specifically, rising *yo* was never compatible with them, while falling *yo* could appear with rhetorical questions:

(96) FINAL EMPIRICAL GENERALIZATION FOR *yo* WITH INTERROGATIVES

a. An utterance questioning p? is never marked with *yo*⇑.

b. An utterance questioning p? may be marked with *yo*⇓ iff the question is *rhetorical*. Using it sounds **snide**.

The ideal account of *yo* should be able to account for all of these generalizations. In particular, weak informativity remains the key goal; the most crucial claim of the next section is that counter to McCready (2008b) (and tacitly counter to (Davis, 2011)), the hearer-newness requirement on *yo* is not absolute. Many of the other generalizations about *yo* appear to be contingent on intonation.

3.2.3 The distribution of *ne*

Before launching in *yo*, I now turn to *ne*, which in many ways is *yo*'s opposite. When used with a declarative sentence expressing a proposition φ, *ne*(φ) is taken to presuppose that φ is hearer-old.[11] In the example below, *ne* is felicitous when the speaker believes that he

[11]Note that *ne* alternates with a few other forms, including the longer *nee* and more colloquial *naa*, which tend to mark commiseration. There is reason to believe these are more than simply stylistic variants of *ne*. I leave an exploration of these facts to other work.

and his addressee both think φ is true. Omitting *ne* implies that the addressee might not have been aware of φ prior to the speaker's utterance.

(97) *The meteor approaches. / #Souta feels sick from eating too much yakisoba.*

S: Owari da ne.
end COP NE
'(This is) the end, isn't it.'

In the licit context, it is not necessary that Ayaka has commented overtly on the probable outcome of the meteor's impact. All that is required is that Souta judges Ayaka competent enough to draw her own conclusions about what will happen. This explains the contrast with the alternative context above, where Souta feels as though he will die from overeating. Ayaka is in no position to reach the conclusion that Souta has, so it is illicit for Souta to presume that she agrees with him on this point.

As suggested by the translation, *ne*-marked utterances require a *yes* or *no* response. This issue will reappear at the end of §3.5.1, but for now, it is sufficient to point out the similarity between *ne* and the English tag question; in addition to both phenomena constructions requiring a response, both may occur with either falling declarative intonation, as above, or question intonation, as below:

(98) *Souta is considering cooking breakfast for Ayaka's friend Taro, and so is seeking advice from her.*

S: Waffuru-ga daisuki da ne?
waffles-NOM great.like COP NE
'He loves waffles, doesn't he?'

With imperatives, *ne*(φ) indicates a polite request.[12] Default intonation and a final question rise are both acceptable, as shown below.

(99) *Ayaka is tutoring Souta on Inquisitive Semantics.*

[12]The grammaticalized politeness system of Japanese interacts with the imperative to create a number of different forms that encode the core imperative speech acts. The example here is polite, but *ne* is equally good with less polite imperatives, such as *Ikinasai ne* "Go (there), ok?" Less good with *ne* are direct imperatives such as *Ike!* "Go!," likely because they are too abrupt to 'soften' with *ne*.

A: Tyanto benkyoo site kudasai ne./?
properly study do please NE
'Please study properly, ok?'

Native speakers note that use of *ne* "softens" a request, and makes it easier for the speaker to refuse and cancel the imperative. This makes it useful both with imperatives that ultimately benefit the addressee like (99) above, as well as those that constitute polite requests:

(100) *Ayaka is sick with the flu.*

A: Mizu, motte kite ne./?
water bring come NE
'Bring me some water, would you please?'

Finally, note that *ne* is felicitous with interrogatives:

(101) Sonna hon, kau (ka) ne?
that.kind.of book buy (Q) NE
'I wonder if (he'll) buy that kind of book.'

As suggested by the gloss, however, the result is not a regular question. Rather, the speaker seems to be musing about the question. The addressee can still answer the question if he wants to, but unlike with normal questions, there is no obligation to do so:

(102) *In the bookshop, Souta and Ayaka are watching Taro from afar as he examines an English reference grammar with great intensity.*

S: Sonna hon, kau ka ne?
that.kind.of book buy (Q) NE
'I wonder if (he'll) buy that kind of book.'

A: a. Kau yo⇑.
buy YO⇑
'He's going to!'

b. Soo ne.
that-way NE
'Good question.' (lit. 'It *is* that way, huh.')

Note that Japanese is not alone in allowing these kinds of musing questions. Farkas (2010) reports that the special question particle *oare* in Romanian has a similar effect. Farkas

(2010) further notes that the *oare* morpheme in involved in the formation of free choice items, such as *oarecare* ('any old'). She uses this to argue that *oare*-marked questions license a broader array of canonical answers than normal questions, meaning that these utterances remain questions at their core. In the final analysis of *ne* presented here, I instead argue that apparent questions like (102) are instead metalinguistic; the question is mentioned, but actually asked. Arguably, these two strategies single goal: to allow a speaker to present a question without requiring an answer.

Compared to the complications in characterizing *yo*, the generalizations about *ne* are fairly straightforward, thanks to the lack of intonational complications:

(103) EMPIRICAL GENERALIZATIONS FOR *ne*

a. An utterance expressing *p* is marked with *ne* iff the speaker believes *p* is **header-old**. It requires **confirmation**.

b. An utterance expressing *p*! can be marked with *ne* to indicate the speaker's **weaker deontic authority**.

c. An utterance questioning *p*? is marked with *ne* iff it is a **musing** question, which relieves the addressee of the need to answer.

ne is used with a declarative when the speaker believes the content of his utterance is not novel to his addressee (103a). It is used with an imperative when the speaker wishes to display weak deontic authority, turning into a suggestion what might otherwise be construed as a command and thereby 'softening' the utterance, in the terms of native speakers (103b). Finally, it turns a question into a statement of wonder, or musing (103c). The conceptual core common to these generalization is that with every sentence type, *ne* makes the resulting speech act less demanding of the addressee.[13] Ideally, the final analysis of these facts should capture this.

3.2.4 The distribution of *yone*

One understudied facet of the data on *yo* and *ne* is their ability to co-occur in a single utterance. In these cases, only the order *yone* is attested, and the complex particle has

[13]That said, *ne*-marked declaratives require a response, which weakens this claim. See §3.4.3.

some aspects of both *yo*'s and *ne*'s behavior. For example, a *yone*-marked utterance behaves mostly like a *ne*-marked one, but with additional "force" to the proposition, as shown below.

(104) *Ayaka is cooking eggs on the pavement.*

S: Kyoo, atui ?(*yo*) *ne*.
today hot ?(YO) NE
'It is (seriously) hot today, huh.'

(105) *Souta and Ayaka are watching videos of Boo the puppy on YouTube.*

A: Kawaii ?(*yo*) *ne*.
cute ?(YO) NE
'He's (super) cute, huh.'

The particle sequence occurs with both declaratives and imperatives. Like *ne*(φ), A declarative of the form *yo*+*ne*(φ) assumes φ is hearer-old and requests confirmation.[14] It also requires a response, like *ne*. *yo*+*ne*(φ) also strongly commits the speaker to the truth of φ (compared to an utterance marked with *ne* alone), to the point where it's sometimes odd without *yo*. This is shown in (104) above, where it would be pragmatically odd for Souta to need Ayaka's input to conclude whether it's hot out.

One interesting fact about this construction is that *yone* appears to occur more readily with 3rd person propositions, while *ne* is most natural with 2nd person:

(106) Hanako-wa kuru ?(yo) ne?
Hanako-TOP come ?(YO) NE
'Hanako's coming, right?'

(107) Onaka itai (?*yo*) *ne*?
Stomach hurts (?YO) NE
'Your stomach hurts, right??'

In (106), the sentence deals with Hanako, and is best with *yo*, while in (107) using *yo* is a bit strange when asserting something about the addressee.[15]

[14] McCready (2008b) notes that *yone* could be built compositionally as either *ne*(*yo*(φ)) or (*yo*(φ) $\wedge$ *ne*(φ)). I remain agnostic on this point for now; under my final analysis, it will not matter, although in adopting the overall structure proposed by Davis (2011), I will technically settle on the latter, nested view.

[15] Junko Ito (p.c.) points out that examples like (107) are possible with *yone* under very specific circumstances. For example, if a child and his mother are at the doctor's office and the child refuses to tell the doctor

With imperatives, *yone* is best in situations where the speaker is checking on a previous request. Compare the *ne*-marked utterance in (100) above to its *yone*-marked counterpart:

(108) *Ayaka is sick with the flu.*

A: Mizu, motte kite yo ne.
water bring come YO NE
'You're bringing me that water, right?'

Here, Souta has previously said that he would comply with Ayaka's request, but he has not yet done so. As indicated by the gloss, the result is that Ayaka's *yone*-marked request seems to take for granted that the request will be filled, but it is not a command, the way it would sound if the imperative were marked simply with *yo*.

Finally, two other facts about the construction are worth mentioning. First, *yo* does not occur with different intonational morphemes under *ne*. It was shown that rising intonation on *yo* is correlated with the so-called relevance requirement, so given the lack of the intonational morpheme in the combined *yone* case, it is unsurprising that there is no relevance requirement for *yone* the way there is with *yo* alone; a *yone*-marked utterance is primarily a confirmation-seeking act, and the conversational needs of the addressee to not figure prominently.[16] Also, given that falling intonation on *yo* seems to correlate with issuing a correction, it is not surprising that these two are incompatible. Asking for confirmation about something you are implying the addressee is misinformed about is an impossible situation; if the speaker believes the addressee is wrong about the world, he should not seek confirmation from her about the very issue at the locus of her misinformation. I conclude from these facts that *yone* is much more similar to *ne* than *yo*, and that if *yone* really is compositional, the contribution of *yo* in this case, devoid of the contributions of its intonational morphemes, reflects *yo*'s core character.

---

what the problem is, his mother might felicitously utter (107). Further, we can imagine contexts where the addressee is not consciously aware of his internal states, in which case *yone* becomes much easier to use (e.g., *You're thinking about the accident again.*).

[16] A *yone*-marked utterance *can*, of course, be relevant in discourse; (107–106) in particular are clearly intended to push discourse forward and so are relevant in a broad sense. Unlike *yo*-marked utterances, though, this kind of relevance is not strongly addressee-focused, nor is relevance required, as shown by (104–105), which are functionally small talk. See also the next footnote for my view of what counts as small talk.

With this in mind, the previous discussion has given rise to the following generalizations about *yone*-marked utterances:

(109) EMPIRICAL GENERALIZATIONS FOR *yone*

a. An utterance expressing *p* (declarative) or *p*! (imperative) is marked with *yone* iff:

   i. the content is **hearer-old**, and

   ii. the speaker is **more certain of** *p*, than if he'd used just *ne*.

   It requires **confirmation.**

b. A *yone*-marked utterance show a **preference for 3rd person**, while *ne*-marked ones are best with **2nd person**.

c. An utterance questioning *p*? is **infelicitous** when marked with *yone*.

The hearer-oldness requirement (109ai) and the need for a confirmation are carried over directly from the generalizations about *ne*, while (109aii) and (109b) contrast *yone* with *ne*. The relative certainty requirement in (109aii) in particular evokes the generalizations about *yo*, although at this stage the preference for 3rd person in (109b) is more mysterious. Still, enough of these components are familiar that it is desirable to derive the effect of the *yone* sequence compositionally, if possible. My analysis will do exactly this, deriving the meaning of *yone* from its constituent parts and their interaction.[17] It is worth noting that *a priori*, there is no reason that *yone* must be compositional; *yone* could be analyzed as a new particle with a distinct discourse effect. But given the intuitive relationship between its meaning and that of *yo* and *ne* independently, such an account would be unsatisfying. As will be shown, this is one place where my analysis diverges dramatically from existing treatments.

At this point it is important to clarify what exactly is meant by the claim that *yone*'s 'meaning' should be compositional. For one, I do not claim that the particles discussed here in any way directly encode the generalizations above. In fact, I do not claim that these particles carry semantic content of any sort; the literal meaning of an expression is unaffected by the presence or absence of discourse particles. Rather, discourse particles are markers that indicate already-existing arrangements of information in the context.

[17] I will not, however, derive their fixed order.

Under the view of discourse presented in this dissertation and following Farkas and Roelofsen (2012), I assume that the default effects that an utterance has on discourse is determined by the interaction of its form and the shape of discourse prior to that utterance. Discourse markers then introduce a separate layer of meaning where they indicate how a particular utterance's effects will diverge from the default effects. Discourse markers are therefore conventionally associated with particular non-default discourse effects, but they are neither necessary or sufficient to bring them about. To say that *yone* should be compositional therefore means that *yo* and *ne* should be explained in such a way that when the sum of the non-default effects they describe are met, *yone* is used.

3.2.5 Particles versus bare utterances

One remaining data point that deserves richer treatment is the difference between marking an utterance with any of these particles and simply leaving it bare. While bare utterances have innumerable uses, one in particular stands out based on the previous discussion. Saigo (2011) characterizes bare declaratives as "closing" utterances that indicate the speaker wishes to conclude discussion of a particular issue (or stop conversing altogether). This use can be exemplified most clearly in contrast to the small talk use of *yone* in examples like (104–105) above; if the speaker were to simply assert *it's hot* without *yone*, his addressee might conclude that he is talking to himself, or otherwise does not wish to converse about the issue.[18] This marks the final generalization to be accounted for:

(110) EMPIRICAL GENERALIZATION FOR bare utterances
Bare declaratives can be used in **closing** utterances.

Accounting for this fact is especially difficult because we clearly do not want to conclude that bare utterances are somehow the marked case. The fact that bare utterances can be used to signal the desire for a conversation to end should be a consequence of what it is that *yo* and *ne* mark when they are used. If the broad meaning of these particles involves tracking certain aspects of the context, this behavior of bare utterances makes some intuitive sense: By choosing not to engage in this aspect of context tracking, the speaker subtly reveals his unwillingness to engage at all.

[18] I construe "small talk" here as conversation that only affects the common ground in a trivial way, such as confirming the status of some proposition as mutually manifest.

3.3 Existing approaches

My analysis of the facts above owes a great deal to the previously mentioned analyses of McCready (2008b) and Davis (2009, 2011), so this section briefly summarizes those earlier accounts before launching into my own. In light of the generalizations above, I conclude that while the existing treatment of *yo*'s intonational contours offered by Davis (2011) seems correct, the accounts of the particles themselves fail to derive *yo*'s weak informativity requirement or the closing-response effect of bare declaratives. Further, the analyses do not leave room for an adequate analysis of *yone* that is built from the contributions of the two particles alone.

3.3.1 McCready (2008)

McCready's (2008b) account contrasts two possible analyses for the use of *yo* and *ne* with declaratives, one a dynamic account where the particles serve as instructions for information processing, and the other a utility-theoretic analysis that uses the particles to mark the relevance of an utterance relative to a particular context. The two approaches are eventually reconciled so that the particles are update instructions with relevance presuppositions that must be met for the updates to go through.

Under that final account, *yo* and *ne* differ in that *yo*(φ) 'strongly asserts' φ, where *ne*(φ) just asserts it. This means that both give the instruction to update the context (σ) to include φ, but in the case of *yo*, there is an additional instruction to remove $\neg\varphi$ from the context first, if it is present. More formally (where $\sigma||\varphi||\sigma'$ expresses instructions to update a context σ to σ' by adding φ and $\Downarrow$ means to downdate the context):

(111) a. $[\![yo(\varphi)]\!] = \sigma||sassert(\varphi)||\sigma'$
(Where $\sigma||sassert(\varphi)||\sigma' = \sigma||\varphi||\sigma'$ if $\sigma||\varphi \neq \emptyset$,
or $\sigma|| \Downarrow \neg\varphi; \varphi||\sigma'$ otherwise.)

b. $[\![ne(\varphi)]\!] = \sigma||\varphi||\sigma'$

What this means is that both *yo* and *ne* instruct the hearer to accept the truth of φ, and in the case of *yo*, the hearer is encouraged to do so even if he previously believed φ to be false.

The final account also includes presuppositions for both *yo* and *ne*, built on a utility-theoretic model where the goal of discourse is not only to build the common ground, but also explicitly to grow the common ground *in a particular direction*, in order to answer questions that one or more participants are facing. This conception is reminiscent of both the literature on the Question Under Discussion (QUD; Roberts 1996), and the Table as employed in the model presented in this dissertation. These presuppositions take the following form: $B_X Y$ means that some participant X believes some proposition Y; $IV_X(Z,Y)$ represents the informativeness value of Y, relative to some contextually salient question Z for participant X; and d is the 'contextually specified relevance threshold' that must be surpassed for an utterance to be informative.[19]

(113) a. Presupposition of $yo(\varphi) = B_S IV_H(Q,\varphi) > d$

b. Presupposition of $ne(\varphi) = B_S IV_H(Q,\varphi) < d$

In prose, $yo(\varphi)$ presupposes that the speaker believes φ is relevant in helping the hearer reach a decision on Q, while $ne(\varphi)$ will not help the hearer reach a decision. In the case of $yo(\varphi)$, φ's status as hearer-new seems to fall out fairly cleanly, as uttering something relevant to solving a discourse question requires that the one solving the question has not heard it before.

This account faces the following problems in light of the generalizations from the previous section. First, it does not predict felicity in the weak informativity cases exemplified by (75); if the addressee is already aware of φ, the act of confirming it with stronger evidence is not strictly speaking *informative*. This problem can be circumvented if the notion

[19]Before the analyses were combined, the dynamic approach featured alternative presuppositions, included below for completeness. Here, *must* is a deontic necessity model and $B_X Y$ represents that participant X believes Y.

(112) a. Presupposition of $yo(\varphi) = B_S \neg B_H \varphi; B_S must B_H \varphi$

b. Presupposition of $ne(\varphi) = B_S B_H \varphi$

In words, $yo(\varphi)$ is take to presuppose that the speaker doesn't believe the hearer believes φ but believes the hearer should come to believe it, and $ne(\varphi)$ presupposes the speaker believes the hearer already believes φ. The result is that the presupposition of *yo* will only be fulfilled when φ is hearer-new, in the sense that the hearer may not believe it, and the presupposition of *ne* will only be fulfilled when (according to the speaker) the speaker and the hearer are in agreement about φ.

that providing stronger grounds for believing φ is an informative act, but doing so requires a more nuanced representation of the common ground—along the lines of the model presented here. Either way, however, the core insight of weak informativity generalization was that the speaker and addressee must have an imbalance of evidence in favor of the speaker, and capturing this with instructions that the addressee give up a belief in the proposition's opposite is too strong.

Along similar lines, the presupposition of $ne(\varphi)$ as formulated suggests only that φ is irrelevant to Q, not necessarily that the hearer has not heard it before, as suggested by the non-novelty generalization in (103ai). The contrast between the felicitous and infelicitous contexts in (97) is lost here; in both cases, φ is irrelevant to any question of the addressee's, and yet only the case where the hearer can be presumed to know this independently (i.e., where φ is hearer-old) is attested. Further, it is not clear why a presupposition that the utterance is irrelevant to the addressee's goals should lead to the requirement that the addressee respond (103aii).

Finally, McCready's explanation of multiple particles is (admittedly) a bit odd. Assuming his compositional account, $ne(yo(\varphi))$ presupposes that the speaker believes that the proposition, 'The speaker believes φ is highly relevant to the addressee' is of low relevance to the addressee. While this is coherent, it is not consistent with the behavior of *yone* demonstrated in §3.2.4; the presupposition captures none of the generalizations in (109).

3.3.2 Davis (2011)

Davis's (2011) account focuses on separating the contribution of *yo* from the contributions of the special intonation morphemes that occur with it. Recall from §3.2.2 that *yo* occurs in two 'flavors', rising (⇑) and falling (⇓). Davis uses this to insightfully decompose the overall meaning of *yo* into three parts: dynamic update instructions encoded into *yo* itself, a presuppositional "guide to action" encoded in the rising intonational morpheme, and a "call for correction" encoded into the falling morpheme.[20] Based on the generalizations in the previous section, this division of labor seems essentially right, and I will adopt it as well.

[20]Note that this analysis differs from Davis (2009) not so much in content as in distribution. In that earlier work, *yo* encoded the guide to action, while the intonational morphemes provided normal and corrective update instructions.

That said, the specific detonation for *yo* provided here again does not account for its weak informativity requirement, and it obscures any possible connection between the meanings of *yo* and *ne*.

The analysis is presented in terms of Gunlogson's (2003) discourse model, a precursor to the models discussed and expanded in Chapters 1–2.[21] Simplifying slightly, under this model the conversational context C is a tuple of the *public beliefs* of each conversational participant. Assuming the set of discourse participants (abbreviated DC^C) has two members A and B, this would mean that $\text{C} = \langle PB_A, PB_B \rangle$, where each member of the tuple is a set of propositions that the relevant participant presents himself as believing to be true. This arrangement should be familiar as a version of the exploded or extruded form of the common ground introduced in Stalnaker (1978) and related work. The equivalent of the narrow meaning of the common ground must be computed by intersecting all of the public beliefs of the participants.

Under this account, $PB_A(\text{C})$ is the set containing Participant A's public beliefs in C. Asserting some proposition p adds it to the speaker's public beliefs, but the update instructions compose in two steps. First, Davis assumes, following Gunlogson's (2003), that the instructions connecting a declarative sentence form to an assertive speech act are syntactically split over a declarative force head DECL and a force modifier head, represented by normal falling declarative intonation (↓). (Note again that this fall is distinct from the fall on *yo*, notated ⇓). The DECL head takes a propositional argument p and includes underspecified instructions that some set of discourse participants $\mathbb{A}$ should have their public beliefs updated to include p. Falling declarative intonation then composes to specify that it is the speaker whose public beliefs should be manipulated. Formally:

(114) CCP FOR A DEFAULT ASSERTION

a. $[\![\text{DECL } p]\!] = \lambda\mathbb{A}.\left\{ \langle \text{C},\text{C}' \rangle \;\middle|\; [\![p]\!] \in PB_{\mathbb{A}}^{C'} \right\}$

b. $[\![\downarrow]\!] = \lambda S.\left\{ \langle \text{C},\text{C}' \rangle \;\middle|\; \langle \text{C},\text{C}' \rangle \in S(\{\text{S}_C\}) \right\}$

c. $[\![\text{DECL } p \downarrow]\!] = \left\{ \langle \text{C},\text{C}' \rangle \;\middle|\; [\![p]\!] \in PB_{\text{S}_C}^{C'} \right\}$

[21] One potentially confusing departure is that Davis calls Gunlogson's discourse commitments "public beliefs". He also discusses "public intentions," which correspond to Portner's (2007) *TODO* lists for commitments to action. I will make the same introduction in §3.4.2.

Here, $\langle \text{C,C}' \rangle$ is a mapping from input to output contexts, and S_C is the speaker. The result is that in the case of a normal assertion, the semantically defined portion of the complete speech act ends once the speaker is committed to the truth of p. The negotiations by which that commitment gets propagated to other participants and becomes a joint commitment of all of the discourse participants is left to the pragmatics.[22]

Davis claims that the more 'forceful' nature of *yo*-marked utterances comes from the fact that this usually-pragmatic step is specified directly in the semantics; these utterances encode instructions for all members of the discourse to update their beliefs, rather than just the speaker, as shown below.

(115) CCP FOR A *yo*-ASSERTION

a. $[\![yo]\!] = \lambda S.\left\{ \langle \text{C,C}' \rangle \;\middle|\; \langle \text{C,C}' \rangle \in S(DP^C) \right\}$

b. $[\![\text{DECL } pyo]\!] = \left\{ \langle \text{C,C}' \rangle \;\middle|\; \forall x \in DP^C : [\![p]\!] \in PB_X^{C'} \right\}$

Davis treats *yo* as an alternative to standard falling intonation, and modifies DECL to pick out all of the discourse participants. He therefore postulates that the structure of a declarative utterance in Japanese matches the tree below.

(116)

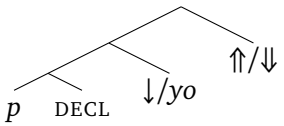

The greatest weakness of this account is that just like McCready's (2008b) account, there is no clear explanation for the the weak informativity generalization of *yo*, which is required to account for examples like (75), repeated below.

(75) *Souta and Ayaka both know that Hanako arrived in town today. Souta knows because he saw her a few minutes ago. Ayaka knows because someone told her so.*

[22] Already, we can see the shadowy outline of how an analysis in terms of Gunlogson (2003) will diverge from the one developed in this dissertation. Where the account above is phrased in terms of instructions for who should update their commitments, my account instead has interlocutors sharing details about the reasoning processes behind their own commitments, in order to facilitate the commitments of others. By taking this step back, from talking about instructions for discourse moves to the reasons behind those discourse moves, we gain a level of flexibility that more accurately mirrors the negotiations of naturalistic dialogue.

A: Hanako-ga tuita tte.
Hanako-NOM came QUOTE
'I hear Hanako has arrived.'

S: (N.) Hontou da yo.
(Yeah.) true COP YO
'Yeah, that's right.'

If *yo* encodes instructions for both Souta and Ayaka to update their commitments with the proposition that Hanako has arrived, *yo* should be aberrant here—Ayaka has already committed to this proposition prior to Souta's instructions.[23] This analysis does not capture the intuition that Souta is expressing that he knows this information better than Ayaka does.

3.3.2.1 *ne* under Davis (2011)

A criticism worthy of its own heading is that Davis' account leaves little room for a satisfactory analysis of *ne*. Most accounts of these particles try to capture the intuition that *yo* and *ne* are in opposition, meaning that their denotations should be in some sense inverses of one another (e.g., as in McCready (2008b)). This is not merely an aesthetic consideration; given their behavior as described in §3.2, these particles are descriptively in opposition, and if possible, this ought to be mirrored by the formalism that captures their behavior. By extension, it is unclear how Davis would analyze *yone*, especially with regard to explaining why the special intonational morphemes become unavailable. One possibility is that *ne* occupies the same syntactic position that ⇑ and ⇓ do, as discussed briefly below. This raises the question, however, of why both *ne* and *yone* occur with default declarative and rising question intonation.

3.3.2.2 A guide to action and a call for correction

Davis (2011) argues strongly for a division between the work that intonation and *yo* itself do in the derivation of a *yo*-marked utterance—a division I endorse, modulo the concerns

[23] Ayaka's utterance in (75) really does seem to publicly commit her to believing Hanako has arrived, even with the quotative marker *tte*. For a potentially conflicting discussion of *tte*, see (Itani, 1998).

above. While the discussion here is most interested in *yo* itself, it is worth briefly sketching the contribution of the intonational morphemes for completeness.

The rising intonational morpheme ⇑ carries a presupposition about the usefulness (i.e., relevance) of a particular proposition in settling a contextually salient *decision problem* in terms of van Rooy (2003). Under this view, uttering *yo*⇑(φ) indicates that if the addressee (Ac) commits to φ, it will allow her to recognize the *optimality* of some contextually salient action a out of the set of possible salient actions $\mathscr{A}_{AC}^{C'}$. In other words, uttering *yo*⇑(φ) indicates that p will settle the addressee's decision problem.

In order for some actions to be better than others, it is necessary to introduce an ordering of worlds, based on a contextually salient ordering source, in terms of Kratzer (1981). This allows for some actions to be better than others if they lead to the realization of preferred worlds. Davis' formal semantics for the ordering of worlds, the related idea of an optimal set of worlds, and his final semantics for a rising *yo* utterance follow.

(117) Partial Ordering of Worlds:
For all worlds $w_i, w_j \in \cap\, \mathrm{PB}_x^{\mathrm{C}}$, $w_i <_x^{\mathrm{C}} w_j$ iff
$\exists p \in \mathrm{os(C)}[p(w_j) \,\&\, \neg p(w_i) \,\&\, \forall q \in \mathrm{os(C)}[q(w_j) \rightarrow q(w_i)]]$
where os(C) is the *ordering source* in C.

(118) The *optimal set* of an agent x in a context C:
$$\mathrm{OPT}_x^{\mathrm{C}} = \left\{ w_i \in \cap\, \mathrm{PB}_x^{\mathrm{C}} \,\middle|\, \neg\exists w_j \in \cap\, \mathrm{PB}_x^{\mathrm{C}} : w_i <_x^{\mathrm{C}} w_j \right\}$$

(119) $$\left\{ \langle \mathrm{C},\mathrm{C}' \rangle \,\middle|\, (\forall x \in DP^C : [\![p]\!] \in PB_X^{C'}) \wedge (\exists a \in \mathscr{A}_{AC}^{C'} : OPT_{AC}^{C} \nsubseteq a(\mathrm{Ac}) \wedge OPT_{AC}^{C'} \subseteq a(\mathrm{Ac})) \right\}$$

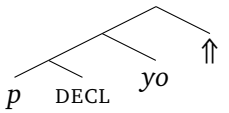

In words, *yo*⇑(p) describes a context change where all of the discourse participants commit to p (thanks to *yo*), and where there is some action a that the addressee can take which is not an optimal action in the input context, but has become one in the output context.

We might then ask, where does a come from? The simple answer is: pragmatics; if the speaker informs the addressee that her train has arrived, the optimal action involves acting on that information by boarding the train before it leaves.[24] Finding a is more

[24]While we might wish to formalize the search for the salient optimal action more rigorously, that mechanism goes beyond the purview of this paper, and indeed beyond linguistics in general.

straightforward in the case of imperatives, because the relevant action is the one described by the imperative itself.

The effect of the falling morpheme ⇓ is much easier to implement than the above. In short, ⇓ requires a proposition q that needs to be removed from the addressee's public beliefs. Formally:

(120) $[\![\text{DECL } \varphi \ yo \Downarrow]\!] = \lambda c.c'$ such that $PB^{C'}_{AC} = (PB^{C}_{AC} - q) \cup \{[\![\varphi]\!]\}$

Here, $-$ indicates subtraction, and q is resolved contextually. I now leave these calculations and return to the issue at hand: how best to treat *yo* (itself) and *ne* in order to account for the data.

3.3.3 Desiderata

So what do these analyses accomplish, and what do they lack? Consulting the generalizations from §3.2, we conclude the following. For *yo*, both analyses successfully capture the idea that a *yo*-marked utterance is a conversationally more forceful move than a a *yo*-less one, but they fail to account for the weak informativity of the particle. While McCready (2008b) puts aside the issue of intonation, Davis (2011) proposes to break down *yo* into three morphemes, *yo*, ⇑, and ⇓. The analysis of the intonational morphemes captures the relevance and correction generalizations, respectively, and I will adopt the same breakdown.

For *ne*, McCready's (2008b) attempts to capture the hearer-old generalization in terms of an anti-relevance requirement. This leads to an incompatibility between *yo* and *ne* which prevents satisfactory coverage of the generalizations for *yone*, that it behaves like a more forceful variant of *ne*. Davis (2011) also encounters problems with *ne* because the tools used to capture *yo* leave no room for compatibility between the effects of the two particles.

As long as we construe these particles in terms of informativity, these problems will persist, because utterances will always be either informative or not. Moving forward, I argue that the key to dealing with both the weak informativity of *yo* and the felicity of *yone* requires moving from a discussion of informativity to the factors that lead an utterance to be informative. This means comparing the commitments of participants, and especially the evidence behind those commitments, and building an account of the particles around those notions.

3.4 An evidential model of discourse

The analysis I propose in this chapter is couched in terms of the model of discourse developed in Chapter 2, as briefly summarized below. This section then discusses a few further expansions to the model necessary to capture the Japanese data, notably the inclusion of a Portner (2007)-style $TODO$ list to model the effects of core imperative speech acts on analogy with the commitment structure for questions and assertions.

3.4.1 Summary of previous commitments

I assume a commitment-based discourse structure built on that of Farkas and Roelofsen (2012; F&R), which features a register of issues called the Table and a set of discourse commitments DC_X for each discourse participant X.

The Table is a register of the issues participants raise in the course of a conversation, arranged in a stack (Kaufmann, 1997, 2000). An issue is a set of propositions under consideration for inclusion in the common ground. Note that F&R's structure makes use of the framework of Inquisitive Semantics, which I do not require. I therefore treat an issue as the proposition or set of propositions denoted by a sentence; a declarative sentence with content p raises the issue p, while a interrogative sentence's issue is the set of its answers (i.e., its Hamblin semantics).

To raise an issue is to propose that at least one of the proposition(s) that constitute the issue (termed the *alternatives*) be considered for inclusion in the common ground.[25] Because no single participant can change the common ground unilaterally, each participant in the discourse instead manipulates his own discourse commitments, which is intuitively comprised of the propositions each participant has agreed to behave as though he believes to be true. The common ground consists of shared background knowledge and the intersection of the participants discourse commitments. Under the Stalnakerian view that the purpose of discourse is to grow the common ground, increasing the size of this intersection constitutes progress toward that goal. When

As noted by Harnish (2005), however, commitment is not a monolithic notion in two

[25] Again, note that under the current view, an alternative is simply a classical proposition, i.e., a set of worlds. In terms of types this is actually equivalent to the Inquisitive Semantics notion of an alternative, even though a proposition is construed as a set of sets of worlds.

regards. **First**, commitments are made for "reasons," meaning simply that every proposition a participant commits to is underwritten by larger bodies of information. This is most clearly illustrated by expressions of evidentiality, which make explicit the kind of evidential base that underwrites the content of the evidentially-marked utterance:

(26) a. Juse iɾida di-manika-**ka**
José football 3sg-play-**rec.p.vis**
'José has played football (we saw it).' *visual direct*

b. Juse iɾida di-manika-**mahka**
José football 3sg-play-**rec.p.nonvis**
'José has played football (we heard it).' *nonvisual direct*

c. Juse iɾida di-manika-**nihka**
José football 3sg-play-**rec.p.vis**
'José has played football (we infer it from visual evidence).' *visual indirect*

d. Juse iɾida di-manika-**sika**
José football 3sg-play-**rec.p.assum**
'José has played football (we assume this on the basis of what we already know).'
inference

e. Juse iɾida di-manika-**pidaka**
José football 3sg-play-**rec.p.rep**
'José has played football (we were told).' *hearsay*

The examples from Tariana (Arawak) above demonstrate a five-way contrast in that language's system of evidentials (Aikhenvald, 2004, p. 2–3). Each of these utterances commits the speaker to the same propositional content, but a different evidential base underwrites the commitment in each case. In these situations, we say the commitment is *conditioned* on different evidence.

Second, commitments can be made at various "degrees" based on the strength of the reasons underwriting them. Gunlogson (2008) exploits this kind of variability in commitments to analyze the distinction between a falling declarative (*That's a persimmon.*) and a rising one (*That's a persimmon?*). She claims utterances of both types share a common semantic core, but that the latter involves a weaker commitment that is contingent on future discourse moves from the addressee.

F&R, citing Gunlogson (2008), model some of the ways that commitments can vary by dividing each participants's *DC* set into subsets 'tagged' for particular kinds of commitments. Strength of commitment is modeled as a binary distinction between *actual* and *conditional* commitments (originally called 'contingent commitments' by Gunlogson). As for capturing the reasons behind commitments, F&R model a binary distinction in *sourcehood*, which tracks whether a commitment is made on the speaker's own authority or on the authority of an interlocutor. These are commitments as *source* and *dependent*, respectively. The result is a four-way subdivision of each DC_X.

Chapter 2 argues that while this division may successfully capture the data these authors discuss, it does not 'scale up' to encompass other, arguably related notions in a transparent way; illocutionary evidentiality like that displayed in (26) must be treated independently, and conditional commitment, while a useful description, actually conflates two issues, *implicit addressee authority*, and *weaker speaker authority*, (reintroduced shortly).

Note that Japanese provides another clear example of the need to capture commitments of varying quality. Just like the English glosses suggest, the overall discourse effects of the utterances below—specifically in terms of how strongly they commit the speaker—is quite different.

(98′) *Souta is planning to cook for Taro.*

a. Waffuru-ga daisuki da yo.
waffles-NOM great.like COP YO
'He loves *waffles*, man!'

b. Waffuru-ga daisuki da ne?
waffles-NOM great.like COP NE
'He loves *waffles*, doesn't he?'

Uttering either sentence above would seem to commit Souta to the proposition *Taro loves waffles*, but later learning that he was wrong would be a more marked turn of events after the *yo*-marked assertion in (a) than after the more tentative, *ne*-marked one in (b). For example, it is much easier for the speaker to immediately distance himself from the *ne*-marked utterance than the *yo*-marked one:

(98″) *Souta is planning to cook for Taro.*

a. Waffuru-ga daisuki da yo. # Wakaranai kedo.
waffles-NOM great.like COP YO know.not however
'He loves *waffles*, man! ... But I don't know for sure.'

b. Waffuru-ga daisuki da ne? Wakaranai kedo
waffles-NOM great.like COP NE know.not however
'He loves *waffles*, doesn't he? ... But I don't know for sure.'

It is of course possible for the speaker of (98″a) to revise his original commitment, but doing so is understood as self-contradictory—the speaker 'walks back' from his commitment in the second clause. In (98″b), there is no such problem. The speaker can immediately indicate uncertainty, because his *ne*-marked utterance does not publicly bind him to the truth of his utterance as strongly.

Based on considerations like this, Chapter 2 provides a generalized structure for each participant's discourse commitment set to account for evidentiality, sourcehood, and conditionality with a single flexible mechanism, whereby a commitment is built of two components: its propositional content, the evidential base that underwrites the commitment.

(121) DISCOURSE MODEL

T = the Table, conversational workspace where issues are raised

DC_X = set of commitments that X subscribes to (for each X in the discourse)

(122) COMMITMENT

A *commitment* is an ordered pair $\langle p, E \rangle$, where E is the commitment's *evidential base*.

An *evidential base* is a set of propositions that constitute evidence for a given proposition. In absence of an explicitly introduced base, the *default base* comprised of a speaker's private beliefs and shared contextual evidence is employed. The character of a base can be modified or replaced by various illocutionary modifiers, including evidentials and expressions that indicate dependency (like the English rising declarative):

(40) DEPENDENT COMMITMENT

Let q stand for the fact of speaker B's independent commitment to p (or a closely related proposition). When a speaker A commits to p as dependent on speaker B, he conditions his commitment on the base E_{DEP} s.t.:

i. $q \in E_{\text{DEP}}$ and

ii. the speaker could not felicitously commit to p given a base $E_{\text{DEP}'}$ where $E_{\text{DEP}'} = E_{\text{DEP}} - q$.

The dependent base crucially contains the fact that the addressee is committed to the same proposition, or a salient, related one. Making a commitment that is conditioned on this base therefore indicates that the continued viability of the commitment rests on the addressee's willingness to avoid retracting the crucial commitment.

The "strength" of a commitment in terms of the discourse structure refers to how resistant to change the speaker indicates that its evidential base is. Commitments are by default *strong*, which simply means that the speaker does not foresee having to revise the commitment. Certain moves can also characterize the evidential base underwriting a commitment as *weak*. The weak base differs from the default base in that it is highly mutable, because it relies on the availability of evidence that the speaker does not yet have:

(123) (41) WEAK COMMITMENT

Let q stand for some proposition that serves as evidence for p. When a speaker weakly commits to p, he conditions his commitment on the base E_{WEAK} s.t.:

i. $q \in E_{\text{WEAK}}$ but

ii. the speaker could not felicitously commit to q (given any base) and

iii. the speaker could not felicitously commit to p given a base $E_{\text{WEAK}'}$ where $E_{\text{WEAK}'} = E_{\text{WEAK}} - q$.

The weak base carries an inherent expiration condition; if no additional support appears to strengthen it, the base will become inviable. This indicates a tenuous connection between the propositional content of the commitment and the base that underwrites it. Note that the dependent base in (40) and the weak base in (41) are both, in a sense, 'weak', in that they describe the conditions under which the base could fail. This is all that is meant by the notion of *weak* versus *strong* commitment here; while a speaker might in principle trust that his commitments are appropriate ones to various degrees, this gradience is not reflected in the discourse structure.

In a related point, it is important to keep in mind that that status of a commitment as conditioned a the default (strong) base or the weak base is a distinct notion from a

proposition's epistemic modality—although in some cases the two might go hand-in-hand; a speaker can make a strong commitment to *might-p*, indicating that she is staunchly certain that the modalized proposition is true, or she can make a weak commitment to *p*, indicating a categorical commitment to *p* would be relatively easy to give up, given future evidence to the contrary. These discourse moves might implicate one another in certain contexts, but their direct effect on the discourse structure are distinct.

The result of these considerations is a flexible system that can be easily expanded to account for other privileged distinctions among commitments.[26] The generalizations in this chapter will be shown to follow from one such distinction, and will require no further changes to the basic structure of a commitment, although existing notions will become clearer in the process.

3.4.2 Commitments to action

It has been observed that *yo* and *ne* occur not only with declarative and interrogative sentences, but also imperative ones. It is therefore important to pause and determine how the discourse model here might be expanded to properly reflect the discourse effects of utterances that employ these sentences. Note that the account provided here is not a full one; because the focus of this dissertation remains on the effects of assertions and questions, I include only enough about imperatives to deal with the Japanese particle data. In the face of a wider range of data on imperative speech acts, the specifics offer here would require much more nuancing. I proceed under the assumption that this nuancing would maintain the spirit of what is offered below.

Portner (2007) notes that the results of speech acts that employ imperative sentences do not lend themselves to analysis in the sort of framework introduced above. Consider the utterance below, which is either a suggestion or a command, depending on Crono's mother's intentions:

[26] As a preview, in Chapter 4, I argue that biased questions in English are best modeled in this system by introducing the notion that the *default base*—built from the speaker's private *speaker beliefs* and shared *contextual evidence* that grows over the course of a conversation—can be referenced at the current conversational index, or at one step previous (i.e., the *Hey wait a minute* state.) This earlier state of the default base, I term the *prior base*. Adding all these 'knobs' to the discourse model raises the worry of over-generation, but I demonstrate at the end of chapter that this concern is not warranted.

(124) *Crono's mother wants him to get up and go meet Lucca.*

M: Crono! Wake up, Crono!

a. ≠ Mom commits to 'Crono wakes up'

b. ≠ Mom commits to 'Crono will wake up'

Either way, however, this utterance crucially does not commit Crono's mother to any of the potential paraphrases above. This is because a *core imperative speech act* like a command or a suggestion does not publicize a speaker's commitment to act as though the world *is* a particular way (Farkas, 2011). Rather, it publicizes the speaker's stance on how the world *should come to be*.[27]

Note that it is insufficient to capture the effect of an imperative as merely a commitment to a desire that the world change in a particular way, as in (124c). Such a commitment would still ultimately rest on an inherently epistemic notion of evidence; this kind of evidence is appropriate for underwriting the default effects associated with declarative and interrogative sentences, but imperatives like (124) are at their core deontic.

Portner (2004, 2007) captures this difference by introducing the idea of the *TODO list*, an individualized register of *commitments to action*, on analogy with the discourse commitment list as a register of *commitments to truth*. Under such a conception, each participant X maintains a new discourse component, $TODO_X$, in addition to DC_X. While Portner (2004) conceives of the contents of $TODO_X$ as open propositions, I follow (Farkas, 2011) in taking commitments to action to be at their core propositional. The reader who disagrees with this move is invited to proceed with a different conception of the imperative's denotation; I will find it useful to maintain as strong a parallel between the structure of $TODO_X$ and DC_X as possible.

Under this view, $TODO_X$ contains all the propositions that X is publicly committed to bring about the truth of, and just like commitments to truth, commitments to action can be undertaken based on various sorts of justifications, which in many (but not all) cases are deontic. This suggests that the commitments in $TODO_X$ should be modeled analogously to those in DC_X, as the combination of the imperative's content and an (evidential) base that

[27] One potential paraphrase of (124), in line with Condoravdi and Lauer (2010), might be 'Mom wants Crono to wake up'. Under such a view, commitments to action could be more directly incorporated into the previous view. See Condoravdi and Lauer (2012) for a more nuanced view.

justifies it. The result is that when the addressee (B) accepts the command expressed by the speaker's (A) uttering "Do 50 pushups!," B adds the proposition *B does 50 pushups* to $TODO_B$, conditioned on a set of justifications that summarize why A can tell B what to do.

The legitimacy of the evidential base that underwrites a commitment to truth depends on some participant acting as the *source* for that commitment, a discourse-specific designation that applies to the participant who agrees to vouch for the truth of that commitment, on an assumed base of sufficient epistemic authority. Similarly, the legitimacy of the base underwriting a commitment to action rests on the parallel notion of a *deontic source*, as introduced by Farkas (2011). Farkas argues that "[d]eontic sourcehood is rooted in asymmetric power relations between the source and the dependent," meaning that for A to be the deontic source or *d-source* for B's commitment to action, A's power to tell B what to do must derive from a discrepancy in power between A and B (in A's favor) that both participants recognize as legitimate. Committing as *d-source* for a commitment to action therefore requires all participants to recognize the speaker's deontic authority to do so.

If the formal parallel between a speaker's DC and $TODO$ is to be made complete, then it should be the case that the deontic bases underwriting commitments to action can vary in their characterizations just like the evidential bases that underwrite commitments to truth. For example, is there such a thing as a weak deontic base? Arguably, the answer is *yes*; after all, a speaker who performs a command often has grounds to expect his command to be accepted with far greater certainty that a speaker who performs a request. The difference between these acts, which revolves around different expectations for compliance, could easily be reflected in the characterization of the deontic base; for example, we might think of the base underwriting a request as including the fact of the addressee's assent as a 'promissory note' on par with the crucial proposition that the weak base evokes. Whether this degree of parallelism is justified is an empirical issue which I put aside for the time being, although I will apply the characterizations introduced by discourse particles to both evidential and deontic bases.

In summary, I will adopt Portner's (2004) TODO list as the target of the discourse effects of core imperative speech acts. The structure of a participant's $TODO$ mirrors his DC; both are comprised of commitments, although $TODO$ lists commitments to action, while DC lists commitments to truth. A commitment to action is therefore a pair $\langle p, E \rangle$. These

impose different requirements on what counts as cooperative behavior for the speaker. A commitment to action is comprised of a proposition (following Farkas (2011)), and the justifications for that proposition. Every commitment to action requires a d-source, whose underwriting authority is deontic, rather than epistemic. Processes that augment or replace the evidential base of a commitment are by default agnostic to whether the commitment they modify is to truth or to action. This will allow *yo* and *ne*, which affect *E*, to operate in both domains.

3.4.3 The role of sourcehood

Gunlogson's (2008) notion of *sourcehood* is a central feature of F&R-style discourse models, but in Chapter 2 I suggested that it could be rolled up into a more general set of privileged distinctions that restrict the character of an evidential base. Sourcehood will play a pivotal role in the analysis of Japanese *yo* and *ne*, so this section reviews the immediate ramifications of this change. See Chapter 2 for a more detailed discussion of these issues.

3.4.3.1 Previous takes on sourcehood

Gunlogson (2008) proposes that every discourse commitment has a *source*, meaning its content is marked as having been sponsored by a particular discourse participant who has evidence of its truth. This follows from the fact that, in normal discourse, we expect participants to have a preponderance of evidence for the actual commitments they make public and attempt to convince others of.

A participant who is not the source of a proposition can nonetheless accept it and become *dependent* on the proposition's source. For example, consider the following exchange, drawn from Gunlogson (2008):

(15) A: The server's down.

B: i. Yes, I know. / Yes, that's right.

ii. #Yes, I didn't know that. / #Yes, is it?

iii. Oh (I didn't know that.)

As shown by the contrast between B's responses in (b) and (c), *yes* and *oh* correlate with different information states regarding the server's status. When B responds '*Oh*', she seems to

merely accept the truth of the server being down on A's word. B therefore becomes dependent on A for the truth of the proposition at hand. If B replies '*Yes*', she cannot felicitously go on to explain that she *didn't* in fact know that the server was down, meaning that a '*Yes*' response requires that B knew the server was down independently of A's utterance. In this case, when B accepts the utterance, she also indicates that she is a source of the proposition, not merely a dependent.

Farkas and Roelofsen (2012) model sourcehood by subdividing interlocutors' discourse commitments into those explicitly committed to as source DC^s and as dependent DC^d.[28] With this in mind, an informal characterization of the difference between *yes* and *oh* in the above example might look like the following:[29] (Recall that for F&R, DC is a repository of propositions, rather than the more complex objects that the model developed here uses.)

(125) Let $p = [\![\text{The server is down}]\!]$.

a. 'Yes' response: Adds p to DC_B^s (B commits to p as a source.)

b. 'Oh' response: Adds p to DC_B^d (B commits to p as a dependent.)

Just as certain kinds of utterances are associated with particular canonical additions to the Table, it seems that classes of utterances can also be associated with canonical sourcehood configurations. For example, regardless of B's response, A's use of a declarative sentence marks him as a source for the proposition that the server is down, effectively adding p to DC_A^s. Note that A cannot easily weaken or downplay his sponsorship of the proposition. Even if he attempted to make a more tentative assertion by adding rising question intonation ('The server is down?'), he would succeed in committing to p only conditionally, but his role as a source for p would not change.

3.4.3.2 Sourcehood in the evidential base

One vital note is that while the strength of evidence A might have for the server's being down is variable, it seems that sourcehood isn't; if a discourse participant wishes to add

[28] See Chapter 2 for a more complete picture; the sourcehood dimension interacts with the conditional/actual dimension as shown there in Fig. 1.

[29] This characterization differs greatly from F&R's final analysis. These characterizations are intended to illustrate a point about the basics of sourcehood and nothing more.

new at-issue content to the discourse, he must either commit to sponsoring it (i.e., be its source) or not. Consider the example below:

(126) A: Milo told me he's not coming back. ($= p$)

B: Yes/#Oh, I'd heard from Ellie.

Here, both A and B know that Milo isn't coming back. Even though B's evidence is weaker, he responds with 'yes', claiming co-sourcehood so that p is registered on both DC_A^s and DC_B^s. What goes unmarked and unrepresented is that A's evidence is stronger. This is in fact a central feature of the source-dependent distinction; either a participant has sufficient authority to cooperatively vouch for the truth of an utterance's content, or he does not and must rely on the say-so of others. This, and this alone, is what sourcehood seeks to capture.

F&R note that a consequence of this is that committing as source is the default, unmarked way that one commits. In other words, a speaker is generally taken to have independent justifications for his claims, unless he employs special forms to highlight his dependence on the commitments of his interlocutor. As previously mentioned, under the model developed here, this asymmetry between the status of committing as source and dependent is mirrored in the way the distinction is captured in the evidential base; a dependent commitment to p given the evidential base E means that E is characterized as crucially containing the fact of the addressee's commitment to p (or a related notion). The informal redefinition of *yes* and *no* below demonstrate the approach:

(125′) Let $p = [\![\text{The server is down}]\!]$.

a. 'Yes' response: B makes a strong commitment to p given E

b. 'Oh' response: B makes a strong commitment to p given E s.t.:

i. $\exists q \in E$ s.t. q is the fact of A's commitment to p and

ii. B could not strongly commit to p given $\{r \in E | r \neq q\}$

Without this specification, it is taken for granted that the addressee's previous commitment plays no special role in the viability of the evidential base.[30] Either way, however, the base makes no particular reference to how the addressee's justifications compare to the

[30] See §3.7.1 for a discussion of evidential implicature.

speaker's.[31]

The sourcehood distinction therefore picks out a very particular 'slice' of the variation in the authority relationships that interlocutors can have in their commitments; being source indicates a preponderance of authority for a proposition, and little more. For example, if speaker A is dependent on speaker B for the truth of a proposition, it will entail that B's evidence is superior to A's, but given two speakers who are co-sources, neither commitment indicates how the evidence underwriting that commitment stacks up against the evidence underwriting the other—or so it seems for English. In the following sections, the analysis of the Japanese particle data will argue that in fact the relative quality of interlocutors' authority plays a prominent role in the pragmatics of the language. Capturing this distinction is facilitated by the fact that the evidential base E can be characterized in an arbitrary number of ways.[32]

3.5 Measures of authority

So far, this dissertation has assumed an intuitive notion of epistemic authority (and authority more generally). The purpose of this section is to sharpen the intuitions involved. It looks at authority (and sourcehood) through the lens of the Japanese particles from §3.2, and concludes that while authority is the key to capturing what these particles mark, the sourcehood metric is insufficient.

3.5.1 Sourcehood (in)compatibility with *yo* and *ne*

If the Japanese particles from §3.2 involve authority as claimed above, it is instructive to look at how these particles constrain expressions of sourcehood, given that this too is an expression of authority. The data below can be summarized as follows: A *yo*-marked utterance reveals the speaker's assumption that his addressee will be dependent on him for the truth of the proposition it expresses, while a *ne*-marked utterance shows that the

[31]The result of this approach evokes Kamio's (1994) Territory of Information theory, but at one step removed; while the Territory of Information approach involves marking the content of an utterance as falling within the speaker or hearer's "domain," the present approach instead marks the justifications *behind* the primary content of the utterance.

[32]This openness is reigned in over the course of Chapter 4.

speaker believes his addressee will be able to commit independently as source.

So far, the data on *yo* and *ne* has been presented in abbreviated contexts, but looking at the possible responses to *yo*- and *ne*-marked utterances reveals an interesting contrast in felicity. Consider the contrast between the 'yes' and 'oh' answers to the *yo*-marked utterance below:

(127) S: Waffuru-ga daisuki da *yo*.
waffles-NOM great.like COP YO
'I love waffles, man.'

A: #N. / Soo ka. / #Nn.
#yes oh no
'#Yes. / Oh. / #No.'

The discussion above showed that the difference between 'yes' and 'oh' can be modeled as a difference in sourcehood. When a speaker accepts a proposition with 'yes', she makes a strong commitment to p given her default evidential base E, which includes her private beliefs and any appropriate contextual evidence. When she accepts with 'oh', she signals her dependence on her interlocutor for the truth of the proposition by making a commitment to p given an E that would not be an appropriate evidential base without her addressee's commitment as a part of it. Based on this, it would appear that felicitous use of *yo* requires that the *yo*-marked utterance will underwrite the addressee's eventual commitment to the utterance's content. In other words, *yo* blocks commitment as a co-source on the part of the addressee—only dependency is allowed.[33] The 'forcefulness' of a *yo*-marked utterance comes from the fact that A not only commits to p, but to the notion that B is dependent on

[33]Note that for this account to go through, it must be the case that *n*, glossed as 'yes', truly indicates independent commitment like its supposed English counterpart. The contrast below (based on (15)) illustrates that this is the case.

(128) S: Saabaa-ga daun site iru (*yo*).
server-NOM is down YO
'The server's down'

A: #N, siranakatta. / Soo ka, siranakatta.
#yes, did not know oh did not know
'#Yes, I didn't know that. / Oh, I didn't know that.'

him relative to φ and could not assert p herself. This is what Gunlogson (2008) calls the canonical act of *informing*.

In contrast, if we examine the possible responses to a *ne*-marked utterance, the opposite pattern emerges:

(97′) *The meteor approaches.*

S: Owari da ne.
end COP NE
'(This is) the end, isn't it.'

A: N. / #Soo ka.
yes #oh
'Yep. / #Oh.'

Here, Souta's felicitous use of *ne* precludes Ayaka's dependent response. *ne* appears to be licit only in contexts where the addressee can commit as (co)source. This is the arrangement Gunlogson (2008) calls the canonical act of *confirming*, which is an apt way to think about the function of *ne*.

Another way to phrase this characterization of *ne* is to say that in felicitous use, the speaker must be a source for p, and the addressee must have the potential to be an independent source for p. B's state in this case is one of being *implicitly authoritative*; Gunlogson (2008)'s definition is given below.

(129) An agent X is *implicitly authoritative* with respect to φ iff:

i. X is an implicit source for both φ and $\neg\varphi$.

(130) An agent X is an *implicit source* for φ iff:

i. X is not committed to φ; and

ii. It is inferable in the discourse context that if X commits to φ, X will be a source for φ.

Applying the above, we could reasonably describe the felicity conditions for *yo* and *ne* in the following way:

(131) *yo* AND *ne* IN TERMS OF SOURCEHOOD

a. Uttering *yo*(p) requires the addressee **not** be implicitly authoritative about p

b. Uttering $ne(p)$ requires the addressee be implicitly authoritative about φ

This is tantalizingly simple, but it does not fare well as the kernel of an analysis in itself. One need only remember the existence of *yone*-marked utterances to see why; there is no way for the hearer to be both a source (as required by *ne*) and dependent (as required by *yo*) at once. Further, the condition on *yo* is simply too strong, given the weak informativity generalization required to capture examples like (75), repeated below:

(75) *Souta and Ayaka both know that Hanako arrived in town today. Souta knows because he saw her a few minutes ago. Ayaka knows because someone told her so.*

A: Hanako-ga tuita tte.
Hanako-NOM came QUOTE
'I hear Hanako has arrived.'

S: (N.) Hontou da yo.
(Yeah.) true COP YO
'Yeah, that's right.'

Solving this puzzle will require stepping back from sourcehood and considering the nature of authority in more detail.

3.5.2 Epistemic and deontic authority

So far, the notion of authority has been discussed in broad terms, where being an authority means simply having good justifications for supporting a proposition. Authority can arise in different ways though. Already we have invoked the difference between epistemic authority on one hand, and deontic authority on the other. Of the two, deontic authority is perhaps the more straightforward notion; a speaker has deontic authority to the extent they are able to impose their desires on their interlocutors. Authority is not constant. The quality of a speaker's deontic authority will differ, depending on the imperative speech act it applies to and the context. For example, a linguistics professor might have the deontic authority to request that a student stand and answer a question, but she does not have the authority to demand he drop and do pushups when he fails to answer correctly. If the same student and teacher encounter one another at a coffee shop, the professor likely lacks the authority for either act. In any event, the speaker with the greatest deontic authority over an act determines whether it will happen (assuming commitments are followed through on).

Epistemic authority is similarly variable. A speaker's epistemic authority over a proposition p is the measure of how squarely he sits in a position to judge whether p holds, and this measure can fluctuate as the speaker learns new information or reinterprets old information. That said, there are at least two major founts from which epistemic authority can spring, and they are not equally variable; for the discussion that follows, it is helpful to distinguish between *inherent authority* on the one hand, and *external authority* on the other. Inherent authority derives from information that the speaker alone can truly know. Propositions a speaker has inherent authority over include facts about the speaker's own dreams, bodily states, and thoughts. Modulo any metaphysical waffling, a speaker's inherent authority is unimpeachable, meaning, for example, that if A claims that he dreamt about a rainstorm, it makes little sense to compare B's authority over that claim to A's. External authority, in contrast, covers all other cases where the speaker comes to know a thing, whether through observation or inference. It can be accepted, doubted, and—crucially—compared. If A and B are arguing over what time the bus is due to arrive and A is the only one of the two who can see the bus schedule (and they both know this), then A should be accepted as the greater authority over bus times, based on the external authority vested in him by the bus schedule.

Returning to the issue of discourse particles, the critical observation for the current discussion is this: When a speaker acts as though the proposition expressed by his utterance is informative (weakly or otherwise) or uninformative, he makes a claim about how his authority over that proposition compares to that of his addressee. One cannot be informed of what one is authoritative about.

3.5.3 Relative authority

The crucial generalizations about sourcehood from the previous section are that it is **binary** and **asymmetrical in markedness**. The first generalization states that for any proposition φ and any discourse participant X who has accepted φ, X is either a source or a dependent for φ. Therefore, even though the amount of epistemic or deontic authority a participant has over φ can in principle vary to an infinite degree, his ability to sponsor φ in discourse only depends on whether the evidential base he can muster to underwrite φ crucially contains his addressee's commitment to φ or not. The marked case is when the evidential

base must include the addressee's commitment to be viable. The unmarked case is when the addressee's commitment does not matter particularly to the speaker's ability to commit. In either case, however, commitment relies on the assumption that there is some objective threshold for whether a given base lends sufficient authority to underwrite a commitment, and that the evidential base the speaker employs beats that threshold. Sourcehood, therefore, captures a particular measure of a speaker's *absolute authority* over a proposition: It involves looking at a participant's authority over φ and deciding whether or not he can be authoritative independent of the addressee's prior commitment. Sourcehood is asymmetrically marked because committing as a dependent amounts to proclaiming, "I am making this commitment, even though my authority alone would be insufficient to do so."

But sourcehood calculations are not the only pressure on interlocutors' ability to make commitments cooperatively. Even if a discourse participant in a given context has a preponderance of evidence and so is implicitly authoritative about φ, conversationally the expectations he faces are more stringent; he must have enough of a surplus of authority, relative to the other discourse participants, to be able to make a claim and stand by it—even if it is challenged. These considerations are not restricted to discourse particles; consider the strange dialogue below:

(132) *Souta and Ayaka are sitting on a bench, doing nothing.*

S: #You're hungry.

A: No./#Oh.

The oddness of this discourse with either answer from Ayaka stems from the fact that in normal contexts, Ayaka is inherently authoritative about her hunger, so she will always be more authoritative it than Souta.[34] The only way Souta can assert that Ayaka is hungry is by ignoring this; he commits as a source for the proposition that Ayaka is hungry, even though she is the undisputed authority over her own hunger, so when she responds 'no', Souta's assertion seems strange—or at least strangely bold. Similarly, an 'oh' response, claiming dependence, is odd from Ayaka because by saying it, she recognizes Souta as more authoritative about her hunger than she is, counter to fact. In order to explain why

[34]Of course, if Ayaka already told Souta about her hunger, a dialogue like (132) may be is perfectly fine. Similarly, if Ayaka's stomach is growling, if there is some other evidence of her hunger, or even if Souta has detailed knowledge of Ayaka's eating habits, then (132) would be acceptable.

Souta's discourse move is uncooperative here, we require some way of talking about how one discourse participant is a *better* or *weaker* source than another, which indicates that this kind of comparison plays a role in the pragmatics of discourse. Note that English does not provide linguistic expressions that are sensitive to this distinction, leaving the issue open as to how directly it plays a role in the development of the formal structure of discourse over time. The next section, however, shows that the discourse particles discussed in this chapter mark exactly this kind of authority comparison, reenforcing the parallel with (grammaticalized) sourcehood.

3.6 *yo* and *ne* as relative authority markers

The problem with any attempt to analyze *yo* and *ne* in terms of sourcehood is that while the quantity being measured—namely authority—is correct, the metric it's examined with is the wrong one. This section presents a more nuanced way of looking at the contributions of *yo* and *ne* as markers of relative authority, meaning they place certain requirements on the relative quantities of epistemic evidence or deontic authority participants have. The authority relationships consistent with the use of *yo* and *ne* generally precipitate sourcehood configurations similar to those discussed above, but are not directly bound to those configurations. This allows for a certain amount of variation, which accounts for the data the previous analyses miss.

3.6.1 Reframing the generalizations

It may be useful to step back for a moment and consider the generalizations about the data one more time, focusing on relative authority. At least with declaratives, it seems that *yo* and *ne*, when used felicitously, do the following:

A *yo*-marked utterance containing a proposition φ (i) marks the speaker as greatest authority over φ, from which it follows that (ii) the speaker is a source and the addressee, if she accepts φ, should do so as dependent. This reformulates weak informativity.

A *ne*-marked utterance containing a proposition φ (i) marks the speaker as holding no more authority over φ than the addressee, from which it follows that (ii) either the addressee is a source who the speaker will be dependent on, or the speaker and addressee

are both sources (if the speaker is asserting φ). This reformulates hearer-oldness.

A *yone*-marked utterance containing a proposition φ (i) marks the speaker and addressee as equal authorities over φ, from which it follows that (ii) the speaker is a source, and while he seeks confirmation, his commitment is not contingent on the addressee's.

These generalizations allow us to take a step back from sourcehood, and say that the particular configurations that crop up with *yo*, *ne*, and *yone* fall out from facts about the participant's relative authority, as described above.

At last, a satisfactory analysis of *yo* and *ne* is at hand. In the rest of this section, I leverage the generalization above and the discourse model from §3.4 to explain *yo* and *ne* in terms of relative authority requirements that they impose on commitments. I argue that the particles characterize evidential bases that are consistent with particular relative authority configurations. When these characterizations are combined with the normal effects of an utterance, the patterns present in the data in §3.2 is predicted.

3.6.2 *yo* as a marker of maximal speaker authority

Rather than marking hearer dependence, *yo* indicates that no other discourse participants are more authoritative than the speaker. It encodes this requirement into the character of the evidential base that underwrites any commitments that derive from the *yo*-marked utterance, making the particle yet another kind of discourse-evidential marker.

This is captured in terms of counterfactual requirements on the propositions that constitute *yo*'s base, and authority that the various interlocutors have over them. To assist with this, it is necessary to be able to compare the authority potentially underlying the interloctuors' commitments. Let D stand for the set of active discourse participants, and $AUTH_X(\varphi)$ stand for the measure of authority which discourse participant X can muster for φ (given X's default base). *yo*'s contribution then, more formally, is:

(133) $[\![yo(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base E_{MAX} s.t.:
$E_{\text{MAX}} = \{q | \forall X \in D : AUTH_X(q) \leq AUTH_{Sp}(q)\}$.

This denotation reveals that *yo* does not change the at-issue content of an utterance it

appears in. Instead, it introduces the not-at-issue content that any commitments the speaker makes to the utterance's content will be conditioned on E_{MAX}. This base is characterized such that for each proposition q it contains, there is no participant with more authority over q than the speaker, based on their private beliefs and contextual evidence alone. When the conditions for felicitous use of *yo* are met, it will follow that the speaker is the strongest epistemic authority for φ, meaning that all others should accept φ on his word. Similarly with imperatives, the speaker promotes himself as having the most deontic authority to determine whether φ (the commitment to action) happens.

As discussed in Chapter 2, the evidential component of an expression like *yo* is not-at-issue content, but it is not necessarily presuppositional; it can in fact be new information that the speaker believes the relative authority situation is the way he claims, but because the evidential content is not-at-issue, challenging it is not straightforward. Note that the 'channel' the evidential base follows to enter the discourse is unique; because it forms part of his commitment, the speaker is responsible for its truth, and yet it is not the negotiable part of the commitment.

The critical test for this analysis is whether it can capture both the informative (72) and weakly informative (75) uses of *yo*, repeated below:

(72) *Souta sees Ayaka hasn't noticed that her train has arrived.*

S: Densya kita #(yo).
train came #(YO)
'Your train is here!'

(75) *Souta and Ayaka both know that Hanako arrived in town today. Souta knows because he saw her a few minutes ago. Ayaka knows because someone told her so.*

A: Hanako-ga tuita tte.
Hanako-NOM came QUOTE
'I hear Hanako has arrived.'

S: (N.) Hontou da yo.
(Yeah.) true COP YO
'Yeah, that's right.'

In both cases, the conditions for *yo*'s use are met; the speaker commits to the literal content of his utterance, underwritten by an evidential base whose content (he claims) his

addressee is no more authoritative about than he is. This covers both cases where the addressee is completely unauthoritative (72), and where the addressee is simply less authoritative than the speaker (75). In fact, as it stands, the account seems *too* weak, because it also licenses *yo* in situations where the speaker and addressee have equal authority. These are the situations where *yone* is to be employed, and a discussion of why only *yone* is used in these cases follows in §3.7.1.

The fact that *yo* only occurs with rhetorical questions falls out of this account directly. Because uttering *yo*(φ) indicates that there is no discourse participant more authoritative about φ than the speaker, it is infelicitous to then turn around and solicit the addressee's opinion on whether φ is true. Information-seeking questions are therefore impossible with *yo*.[35]

This account also explains why assumed-answer rhetorical questions sound snide or sarcastic when marked with *yo*. In normal cases, the answer to a rhetorical question is taken to be obvious to all the participants in a discourse, which is why it is not necessary to give the response. When a speaker adds *yo*, however, he is playing with the discourse by presenting himself as though he believes that he is in a better position to answer his own question than the addressee is, so that *yo* is licensed. In essence, a *yo*-marked rhetorical question says, "The answer to this question *should* be obvious, but isn't to you." Snide and sarcastic indeed.[36]

---

[35]Note that *yo* is also incompatible with quiz questions, which are uttered in contexts where the speaker and addressee both know that the speaker knows the answer:

(134) *The teacher is quizzing a student about the Japanese government and asks:*

T: Soori-daijin-wa dare desu ka #yo.
prime-minister-TOP who is QUOTE YO
'Who's the Prime Minister?'

Under the analysis of *yo* developed here, it is not clear why this should be the case. One possibility is that *ka* and *yo* sit in the same syntactic position. Another is that *yo* cannot host its intonational morphemes properly in a question because of standard question intonation.

Interestingly, the preferred way to ask a quiz question like this involves using the tentative copula *desyoo* in place of *desu*. This suggests that quiz questions involve feinting ignorance, which *yo* runs afoul of.

[36]The fact that only *falling yo* is compatible with rhetorical questions derives simply from the fact that the state of being ignorant of an obvious answer to a question is wrong. Failing to mark ignorance is itself an error, so the rhetorical question-asker is compelled to use falling intonation.

The account also predicts that *yo* will be felicitous in response to questions. Consider the relevant example, repeated below:

(78) *Ayaka wants to know whether Hanako went home.*

A: Hanako, kaetta no?
Hanako went home cop
'So then, did Hanako go home?'

S: N, kaetta #(yo).
(yes) went home YO
'Yeah, she went home.'

When Ayaka asks her question of Souta, she puts p (= *Hanako went home*) and $\neg p$ on the Table, commits to the trivial union of $p \vee \neg p$, but make no further commitments, presumably because she can't. Any helpful answer that Souta can provide, then, will occur in a context where Souta is maximally authoritative, so *yo* is licensed.[37]

One problem that any account that makes *yo*'s primary contribution not-at-issue is that it is difficult to explain why it is required even in situations where the information conveyed by the evidential base is obvious, presupposed, or otherwise easily accommodated in other languages that lack a similar particle. For example, if Ayaka has just asserted that the prime minister has died, as in (86), there should be no need to indicate her inability to be co-source when Souta corrects her—it simply follows from that fact that she was wrong.

The common 'solution' to this kind of problem is to invoke a conversational rule that says that whenever an appropriate evidential base can be invoked, it should be. This computation the common scalar implicature, but with non-asserted content. See Schlenker (2006) for an attempt to derive such an implicature directly from Gricean reasoning, or §3.7.1 below, which employs Vanderveken (1997)'s take on Quality, as discussed by Faller (2012).

[37]We must also not forget that relevance is a key component of the discourse effect of (rising) *yo*. As will be shown briefly below, I adopt Davis's (2011) solution of relevance wholesale, but already we can see that however relevance is formalized, this situation will support it. By asking her question, Ayaka makes it clear that 'whether p' is relevant to her (otherwise she wouldn't have asked). When Souta answers the question, then, the response will be relevant.

3.6.3 *ne* as a marker of minimal speaker authority

In contrast, *ne* indicates that the speaker be no more authoritative than his addressee:

(135) $[\![ne(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base $E_{\text{MIN}} = \{q | \forall X \in D : AUTH_X(q) \geq AUTH_{Sp}(q)\}$.

The evidential base E_{MIN} is filled with propositions that the speaker's interlocutors are at least as authoritative about as he is. Perhaps more so than with *yo*, it is important here to note that the identity of each q in E_{MIN} is never specified within the discourse structure. Speaker assumes that if enumerated, his interlocutors would be able to commit to them at least as well as he could, but in doing so he commits himself to none of them in particular. This is in part what makes the effect of characterizing the evidential base so clearly not-at-issue.

Note also that the characterization above, like the one for *yo*, combines with any other requirements imposed from within φ. For example, we know that assertions require the speaker to commit as source, i.e., using an evidential base whose effectiveness is not influenced by the addressee's commitments. These other requirements interact with those of the discourse particle in deriving the overall effects of the utterance. With a *ne*-marked assertion for instance, the speaker must have enough evidence to assert his utterance independently (by the nature of assertion), but by indicating that his addressee is at least as authoritative as he is (using *ne*), he leaves open the possibility of retracting his support if she disagrees. Both the speaker and addressee must therefore have sufficient authority to commit as source. The fact that the addressee can't felicitously commit as dependent is also explained; if the speaker's commitment (including its grounds) is accepted, then the speaker can't know better than her, so it would be odd for her to defer to his judgement.

With core imperative speech acts, the fact that the speaker indicates that his addressee has deontic authority equal to or greater than his transparently explains the 'softening' effect of *ne*; everything from the command to the suggestion is weakened by adding *ne*.

With questions, recall that *ne* always results in a musing 'mention' of a question that does not require an answer, never a default information-seeking one. It remains, however,

unclear why *ne* can't occur with default questions under this view. But there are a few possible solutions. One simple option would be to stipulate that *ne*'s base explicitly indicates that the speaker is an implicit source. This would rule out both information-seeking and rhetorical questions with *ne*, as the necessary requirements would never obtain.

3.6.4 Capturing *yone*

This account captures the facts about *yone* as well. Given the evidential base characterizations in (133) and (135) and a compositional approach to *yone* where the demands of both must be met simultaneously, the following must be true:

(136) $[\![yo+ne(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base $E_{MAX+MIN}$, which is be characterized as both:

i. $\{q|\forall X \in D : AUTH_X(q) \leq AUTH_{Sp}(q)\}$

ii. $\{q|\forall X \in D : AUTH_X(q) \geq AUTH_{Sp}(q)\}$

i.e., as $\{q|\forall X \in D : AUTH_X(q) = AUTH_{Sp}(q)\}$.

In other words, a *yone*-marked utterance indicates that the speaker and addressee will have an equivalent amount of authority over the utterance's propositional content, thanks to equal authority over the contents of the speaker's base. This captures the facts about the 2nd versus 3rd person declaratives in (106–107), repeated below:

(106) Hanako-wa kuru ?(yo) ne?
Hanako-TOP come ?(YO) NE
'Hanako's coming, right?'

(107) Onaka itai (?*yo*) *ne*?
Stomach hurts (?YO) NE
'Your stomach hurts, right??'

As shown above, *ne* is preferred when the proposition deals with the addressee, while *yone* is preferred when the proposition concerns a third party. Souta's uttering (107) certainly

does not lead to his authority over Ayaka's stomach ache equaling her own.[38]

3.6.4.1 Delineating *yo*, *ne*, and *yone* contexts

Note that the relative authority requirements of *yone* are contained within those of both *yo* and *ne* alone; whenever the speaker and addressee have equal authority (i.e., when they share equal access to a sufficient evidential base), *yo*, *ne*, and *yone* all introduce characterizations that describe the base and are, in some sense, live options for speakers. Why then, does only *yone* occur in these cases?

This issue is the first taken up in §3.7.1. In short, I suggest there is pragmatic competition among these three options, drawing from Faller's (2012) discussion of competition among evidential strategies, where a similar effect can be observed. Because *yone*'s evidential base is more strictly characterized than that of just *yo* or *ne*, failing to use *yone* implies that those stricter requirements are not met. In this way, the contexts that license *yo*, *ne*, and *yone* in practice do not overlap, even if their formal properties do: *yo* will be employed only when the speaker has the greatest authority among discourse participants, *ne* only when the speaker has less authority than his interlocutors, and *yone* when authority is equal.

3.6.5 Summary of the account

Under the account offered here, a *yo*-marked utterance requires that any commitments the speaker makes in the course of the utterance invoke an evidential base whose content the speaker is maximally authoritative on, relative to his interlocutors. A *ne*-marked utterance requires the speaker is minimally authoritative. Taken together, a *yone*-marked utterance requires the that speaker and his interlocutors be equally authoritative. The next section discusses a number of remaining issues for the account, including how pragmatic competition delineates the environments where each particle is used.

[38] Recall that examples like (107) are possible with *yone* whenever the speaker is willing to claim authority over the addressee's internal states for the purposes of discourse. (The example given was of a mother and a recalcitrant child.) I believe a pragmatic account of this is quite possible based on the discussion here.

3.7 Details and further predictions

This section continues the discussion of the discourse-evidential approach to *yo* and *ne*. It begins with a discussion of how these particles compete against one another, before entering a final empirical domain, exclamativity. While *ne* can be used in expressions of surprise, *yo* and *yone* cannot. We then return briefly to the issue of intonation.

3.7.1 Pragmatic competition

By the definitions offered in the previous section, any situation that licenses a speaker to employ *yone* in order to mark that he and his addressee are equally authoritative about the content of his commitment is also a situation where he is strictly licensed to use *yo* or *ne* alone. Nonetheless, *yo* in many situations implicates that the speaker is not only maximally authoritative, but that his addressee is *not* authoritative. This is clear from even the basic examples from §3.2:

(72) *Souta sees Ayaka hasn't noticed that her train has arrived.*

S: Densya kita #(yo).
train came #(YO)
'Your train is here!'

This utterance is a textbook example of the informative use of *yo*; Ayaka did not recognize that her train had arrived, and it is only thanks to Souta that she now does. Interestingly, however, this implicature disappears when a *yo*-marked utterance is offered in response to a *ne*-marked one:

(137) *Hanako sees a modest look book on a shelf and is alarmed to learn it's high price from the shopkeeper.*
She responds:

H: E, takai desu ne!
(pardon) expensive is NE
'How expensive!'

S: Soo desu ka/ne/yone/yo
that way is Q/NE/YONE/YO
'Is that so?/It is, isn't it./It sure is, huh./Yes it is!'

Here, Hanako has offered a *ne*-marked utterance, suggesting that the shopkeeper has at least as much authority over the proposition that the price is high as she does. In response, however, the full range of particles is licit, with the different effects approximated by the glosses. Crucially, *yo*-marked response is forceful, perhaps even boastful, but it does not claim that Hanako is an unauthoritative judge of prices—even if it does leave open the possibility that the shopkeeper is somehow more authoritative on the issue.

Given the difference in power relations between Hanako and the shopkeeper, a simpler example where the participants really should be equals perhaps illustrates more clearly:

(138) *Hanako and Souta are on day 3 of a skiing trip. It's been frigid the whole time, but this morning when the pair walk outside, it's amazingly warm. Hanako turns to Souta:*

H: Attataki ne
warm NE
'(It's) warm, huh'

S: Attataki yo↑
warm YO
'Yes it is!'

Here again, Souta's response does not suggest that Hanako does not know that it is warm, even though that is exactly what it would communicate under normal circumstances. Given that Hanako has already indicated with her declarative utterance that she has independent evidence that it is warm out, Souta cannot take that away with his use of *yo*.

The above effect appears to be a rather straightforward result of implicature cancellation. Given their overlapping denotations, *yo*–, *ne*–, and *yone*-marked versions of the same utterance are in competition: Using one implicates that the other evidential relationships do not hold. The main requirement for this argument to go through is that differences in the evidential base be treated as part of the informativity of the overall utterance. A *yone*-marked utterance carries the evidential information of both *yo* and *ne*, by which it follows that *yo* implicates not-*ne* and vice versa.

Alternatively, it is possible to adopt Vanderveken's (1997) generalized version of Grice's (1975) Quality maxim, which rephrases the informativity requirement in terms of simply "strength:"

(139) ILLOCUTIONARY Q-PRINCIPLE:

Let your speech act be as strong as required (that is, neither too strong nor too weak) to achieve your current linguistic purposes in the context of each utterance!

Faller (2012) uses this argument to derive the competition between evidential markers in Cuzco Quecha and German, by ranking the evidential markers by the "strength" of the commitments they license.[39] The issue here is slightly different, but the same general program applies: If Japanese speakers expect their interlocutors to track the relative authority they have over their commitments, then an act that specifies the relative authority relationship more precisely is 'stronger'.

In a language with many discourse particles beyond the two discussed here, this immediately raises the question of how bare (particle-free) utterances in Japanese fit into the emerging picture. How do these compare in terms of of the Illocutionary Q-principle? In short, the answer is that they are quite weak; if the context is such that the speaker has any particular expectations about how his authority compares to his addressee's, he is better off using a particle-marked utterance. This accounts for the 'closing' use of bare utterances, as well as the inattentiveness that is conveyed by failing to use particles in appropriate contexts (especially *yo*): They are uncooperative discourse moves.

The above discussion suggests that pragmatic competition between discourse particles is rampant in Japanese. This seems to be true to a surprising degree. For example, recall that even in situations where the relevance of a speaker's utterance is obvious from context (e.g., when answering a question), Japanese speakers are barred from omitting rising *yo* to mark that relevance. This is a fact that any analysis of these particles will eventually contend with, and therefore does not constitute a weakness of the current account.

3.7.2 Exclamativity

As a final test for the analysis developed here, I turn to a new speech act: the exclamative. Japanese exclamatives can be followed by *ne*, but not *yo* or *yone*. I show that this is expected by the account.

[39] Note that "strength" here means roughly the amount of authority the speaker claims, not the mutability of a commitment.

3.7.2.1 Exclamative mini-primer

Rett (2011) shows that "Exclamation" refers to at least a few constructions:

(140) (Wow,) John bakes delicious desserts! *Sentence exclamation*

(141) a. (My,) What delicious desserts John bakes! *wh-exclamative*

b. (Boy,) does John bake delicious desserts! *inversion exclamative*

c. (My,) The delicious desserts John bakes! *nominal exclamative*

The common factor among these utterances is that they express *broken expectations*, better known as surprise. Exclamation can be expressed using most sentence forms, including the normal declarative in (140), so long as they are accompanied by special exaggerated intonation. The parenthetical surprised particles shown above also help with ensuring a surprised interpretation.

An *exclamative*, on the other hand, is a sentence form with a distinct syntax and, arguably, semantics. Chernilovskaya et al. (to appear), summarizing recent findings in the field, breaks down the meaning of an exclamative into its descriptive and expressive content:

(142) How many people took part in the rally!

a. Many people took part in the rally. *descriptive content*

b. The speaker is impressed/amazed/surprised/awed...
by the number of people who took part in the rally *expressive content*

Note that neither of these components is foregrounded (i.e., put on the Table). This can be show by the fact that neither the descriptive or expressive content proffered using an exclamative can be directly challenged:

(143) A: How many people took part in the rally!

B: Well / ? No, most of the people were bystanders.

(144) A: How many people took part in the rally!

B: # Well / #No, you are not surprised/impressed/...

Are you sure you're surprised/impressed/...?

Unlike assertions resulting from declarative utterances, an exclamative used felicitously can also never be false:

(145) A: What a crowd that is!

B: # That's a lie. / # I don't believe you.

Taken together, this suggests that the commitments that accompany an exclamative are simply not up for debate. Another way of stating this is that because the purpose of an exclamative is to publicize the speaker's surprise, the speaker's authority on that surprise is unassailable.

3.7.2.2 Japanese exclamatives

It is therefore interesting to learn that Japanese features basically the same set of exclamation strategies as English, and that both of them can occur with *ne* (Itani, 1992):

(146) Baka (ne)! Anta-wa hitori-ni narya sinai
stupid NE you-TOP alone-DAT not be
'How silly (of you to have such an idea)! You won't be left alone.'

(147) Nante takai n desyoo (ne)!
how expensive the-case is NE
'How expensive!' (slightly edited from Itani (1992))

(146) exemplifies a normal declarative used to express exclamation. The gloss is given as a true exclamative, but a more accurate one might be 'You're silly, aren't you!'. (147), on the other hand, is clearly an exclamative, as shown by the presence of the exclamative wh-word *nante*.[40] *ne* is generally optional in exclamations of either sort, but it is not the case that *ne* is always allowed. For example, in (148) below, *ne* is not possible because the speaker has no specific addressee.

(148) Zisin (da)!
earthquake (be)
'(It's an) Earthquake!'

[40]Note that widely cited judgments offered by Uyeno (1971) claim that *ne* is not allowed in exclamations.

This suggests that even in cases of exclamation, *ne* offers something akin to its usual meaning: that the speaker's interlocutors have equal access to the grounds for his utterance. A related—and novel—observation is that even true exclamatives cannot support *ne* when the addressee lacks access to the source of the speaker's surprise:

(149) *Hanako picks up a used book and together with Souta, looks at the price on the inside front cover. Hanako is shocked by the price and says with amazement:*

Nante takai n desyoo (ne)
how expensive the-case is NE

'How expensive!'

(150) *Hanako picks up a used book and looks at the price on the inside front cover. She is shocked by the price and says with amazement to Souta, who is watching her reaction from a few feet away:*

Nante takai n desyoo (# ne)
how expensive the-case is NE

'How expensive!'

The contexts in (149–150), differ in whether Hanako's addressee Souta can see the price of the book that leads her to exclaim. In (149), the price is available, and *ne* is licit. In (150), Souta can infer that the price of the book is what has caused Hanako's surprise, but *ne* is nonetheless anomalous.

This behavior is predicted by the analysis offered here. This is because even though Hanako's utterance does not make any at-issue claims, it still commits her to acting as though the content of her exclamation is true:

(151) Nante takai n desyoo (ne). # Ammari takakunai kedo.
how expensive the-case is NE that-extent not-expensive however
'How expensive! # But it's not really all *that* expensive.'

Because Hanako makes a commitment, tagging the utterance with *ne* simply requires that the commitment she makes is made relative to E^{MIN}; the fact that this commitment is not tied to an at-issue component of the utterance's discourse effect is immaterial.

Note that Itani's (1992) claim, while not precise in its meaning, still manages to evoke the analysis presented here:

(152) ITANI'S (1992) CHARACTERIZATION OF *ne*:

The speaker desires to establish the assumptions communicated by the utterance as common ground

The "assumptions communicated" involve the speaker's relative authority over the propositions that make up the evidential base that underwrites his utterance. Only when that content is shared (e.g., the price in (149)) is *ne* is licit.

3.7.2.3 Obvious evidence vs shared surprise

One potential objection to the analysis presented in this chapter is that the propositions that form the evidential base underwriting a commitment might in many cases be rather tangentially related to the speaker's authority for the proposition itself. Fortunately, exclamatives offer evidence that this indirectness is a boon.

While the examples above are potentially ambiguous on the point, (153) below unequivocally shows that it is the obviousness of the evidence, rather than an expectation of shared surprise, that licenses *ne* in an exclamative:

(153) *Hanako is shopping at a used book store alone. She sees a modest-looking book on the shelf behind the counter. She asked the shopkeeper its price, and* ***the shopkeeper responds matter-of-factly with a figure that Hanako finds shocking****. She says:*

Nante takai n desyoo (ne/#yo/#yone)
how expensive the-case is NE/YO/YONE

'How expensive!'

Given that the price originates from the shopkeeper (who shares it without emotion), Hanako is clearly informing the shopkeeper of her surprise in this example. Her *ne*-marked exclamative therefore does not serve to indicate her belief that the shopkeeper is at least as authoritative as Hanako is about whether the price is surprising. Rather, she simply indicates that the facts leading to her surprise are facts that the shopkeeper knows as well as she does. This is because what *ne* requires that the addressee merely share the grounds by which the speaker makes her commitment, and emphatically *not* the commitment itself. The shopkeeper need not be surprised by the price to appreciate the fact that Hananko *is* surprised, and that his access to the reasons for her surprise at least equal hers.

Part of the effectiveness of this strategy is that an exclamative very easily conveys "feinted incredulity" (Farkas, p.c.). The speaker can therefore use *ne* as a way of minimizing her own authority—after all, if the exclamative's content was surprising to the speaker, she can't have had a significant body of evidence pointing toward that conclusion. The overall effect is that Hanako seems to be goading the shopkeeper to share her view, without in any way imposing her surprise. She communicates something akin to: "*Surely* this book is expensive (right?). I'm shocked."

3.7.2.4 Infelicity of *yo* and *yone*

To complete the paradigm, observe that neither *yo* or *yone* can follow exclamatives of either sort, although the results differ:

(154) Baka yo/yone! Anta-wa hitori-ni narya sinai
stupid YO/YONE you-TOP alone-DAT not be
'You're stupid! You won't be left alone.' / '#How silly of you...'

(155) Nante takai desyoo (*yo/yone)!
how expensive is YO/YONE
'How expensive!' (Itani (1992); *yone* data added)

The default assertion in (154) is not ungrammatical, but it fails to communicate exclamativity; the closest description that native speakers will use is "surprised indignation." With the true exclamative in (155), however, *yo* and *yone* are both unavailable.

At first glance, this seems problematic for the analysis presented in this chapter. After all, the previous the explanation for *ne* lends itself to a conversationally 'forceful' mirror in *yo*. (155) with *yo* seems like it could communicate something like, "This book is rather expensive, man, and I'm shocked." This is in fact close to what the *yo*-marked assertion in (156) can express, with a sufficiently animated delivery. But in fact this is not possible. Consider the original (149–150) examples, expanded to include all the particles under discussion:

(149′) *Hanako picks up a used book and together with Souta, looks at the price on the inside front cover. Hanako is shocked by the price and says with amazement:*

Nante takai n desyoo (ne/#yo/#yone)
how expensive the-case is NE/YO/YONE

'How expensive!'

(150′) *Hanako picks up a used book and looks at the price on the inside front cover. She is shocked by the price and says with amazement to Souta, who is watching her reaction from a few feet away:*

Nante takai n desyoo (ne/yo/yone)
how expensive the-case is NE/YO/YONE

'How expensive!'

Even when Hanako's evidence is unavailable to the addressee (meaning that she is a much strong authority on the book's price than Souta), the particle is not licensed.

The argument for why *yo* and *yone* are illicit with exclamatives follows from the nature of what exclamation conveys. Looking again at the examples above, we see that the true exclamative always features 1) a special wh-word, *nante*, and 2) the copular form *desyoo*.[41] As it happens, even non-exclamative utterances containing *desyoo* are infelicitous with *yo*:

(156) Kono hon takai yo
this book expensive YO
'This book is *expensive*, man.'

(157) Kono hon takai desyoo/daroo (#yo)
this book expensive is YO
'This book is expensive, probably.'

This is because, as the glosses suggest, *desyoo* is not a simple copular form. Rather, it is a tentative, hedging move.[42] Itani (1995, 1998) glosses *desyooo* as 'is' or 'will' and the 'probably' gloss is due to Davis (2011). More generally, Hara (2010) argues that it encodes "evidentiality or epistemic (un)certainty."

From this we derive both a question and an answer. The answer is that if *desyoo* involves weakened commitment, then its incompatibility with *yo* and *yone* follows directly; if the speaker both takes himself to be unauthoritative (in absolute terms) and at least as good an authority as his interlocutors, then he proposes to make a commitment that he thinks no one in the discourse is really qualified to make. The question that is raised in

[41]The direct equivalent of *desyoo*, *daroo*, is also licit, though depending on the speaker it can sound as though the speaker is talking to herself. In such cases, *ne* is of course infelicitous.

[42]See Chapter 2.

return is: Why does the exclamative require a hedge? Again, the solution might be to lean on the "incredulity" aspect of surprise.

Returning, at last, to the issue of evidentiality versus epistemic modality from Chapter 2, note that *yo* is licit with modifiers that affect the modality of a proposition, rather than the evidential base of its commitment:

(158) Kono hon tabun takai (yo)
this book probably expensive YO
'This book is expensive, probably.'

Unlike *desyoo*, *tabun* is uncontroversially a modal expression. Because it changes the proposition that the speaker commits to, rather than the strength of the commitment, there is no incompatibility between it and *yo*.

In sum, this subsection has shown that the analysis of *yo*, *ne*, and *yone* offered in this chapter is consistent with the particles' behavior with acts of exclamation. It was shown that *ne* is licit in exclamatives and requires that the addressee have equal-or-better authority over the factors which lead to the speaker's surprise, but thanks to the indirectness of *ne*'s requirement with regard to the actual proposition committed to, at no point does the speaker's use of *ne* suggest that the addressee will share the speaker's surprise. *Yo* and *yone*, on the other hand, do not occur with exclamatives because they are incompatible with the incredulity that these utterances express.

3.7.3 Interactions with intonation

The analysis advanced in §3.6–3.7 improves on the basic denotation for *yo* offered by Davis (2011), but it cannot be the whole story, because it fails to account for the intuition that a *yo*-marked utterance must be relevant to the addressee's conversational goals when paired with rising intonation (⇑), or that it must indicate a correction when paired with falling intonation (⇓):

(ex:decyodown) *Ayaka and Souta have conflicting information about whether the prime minister has died, and Ayaka has just said, "The prime minister died."*

S: Sinde-nai yo⇓/#yo⇑!
died-not YO⇓/#YO⇑
'He did NOT die!'

Independent evidence of the need to maintain some kind of relevance component for *yo* comes from Cook (1990). This corpus study of spoken, Tokyo-dialect Japanese finds that *yo*-marked utterances are used to introduce new topics in discourse only 1.8% of the time. Only 1 out of 55 new topic introductions in the corpus is marked with *yo*.[43] This is a surprising finding given just the account of *yo* presented in §3.6–3.7, which would predict hearer-new information, marked with *yo*, could easily mark the beginning of a discussion whenever the speaker wishes to start off by informing the addressee of something she doesn't know. It is therefore necessary to retain Davis' general account of the intonational morphemes, even under this new view of the particles themselves.

It appears that beside the fact that the intonational morphemes are hosted on *yo*, they do not really interact with *yo* in any deep way; the core meaning of *yo* as a maximal authority marker combines quite naturally with the notions of relevance and correction. As previously noted, *yone* does not host these particles, which suggests that *ne*, too, might occupy the same syntactic position as the intonational morphemes. At first blush, it might seem as though this breaks the parallel between *yo* and *ne*, but this is exactly the level at which the parallel should be broken; if *yo* sits in the position sometimes occupied by the question particle *ka* and *ne* alternates with the intonational morphemes, the patterns of concurrence are covered.

3.8 Extension: Singlish *lah*

As a final extension before closing this chapter, this lengthy section departs from Japanese in order to introduce a a preliminary analysis of the Singlish particle *lah*, which I argue functions as a marker of what might be termed 'shared background assumptions'. Based on a review of existing data, I conclude that *lah* is very similar to Japanese *ne*, with a few differences which the present model is able to capture successfully, especially when we keep in mind that the 'purpose' of discourse-evidential marking of the sort discussed throughout

[43] Unfortunately, the corpus itself is no longer available, so I was unable to examine the aberrant *yo*-marked topic introduction. Cook (p.c.) offered an additional example of a topic-introduction marked with *yo*, but noted that it is most natural when introduced by the phrase *soo ieba*, 'Now that you mention it…'/'Speaking of that…'. In other words, the felicity of (rising) *yo* appears to be dependent on the existence of preceding context.

this dissertation is to assist in the negotiation of commitments.

Note that Singlish features a vast array of discourse particles and there is some indication that the particles interact with tone (see below), so the conclusions about *lah* itself should be taken as tentative. The availability of particles as minimally different as *ne* and *lah*, however, suggests that the overall view of discourse presented here is on the right track. Because Singlish is less well-known than the other languages discussed at length in this dissertation, I begin with a short profile of it.

3.8.1 Language profile

Singapore Colloquial English (SCE; Singlish) exists on a continuum with Singapore Standard English (SStdE), which approximates British English (Deterding, 2007). This continuum—especially the Singlish end—constitutes the *lingua franca* of Singapore; the large ethnic Chinese, Malaysian and Indian populations speak Mandarin, Malay, and Tamil, respectively, but all groups use Singlish and SStdE publicly. Originally a contact dialect, Singlish is heavily influenced by these other languages, and has often been misclassified as a pidgin, although this is possibly a terminological issue. The word *Singlish* can refer both to the "highly informal" English used by native speakers, as well as the non-native English used by second language learners in Singapore (Gupta, 2006). While both Singlishes are deserving of study, I will focus here on the native dialect.

Deterding (2007) notes that Singlish occupies an uncomfortable sociolinguistic position in Singaporean society. It is consistently considered a part of Singaporean national identity, but it is also under siege for exactly this reason; its mutual intelligibility with other dialects of English is low, and many of its distinctive features, including lack of agreement morphology and tones, imbue the language with low (overt) prestige. To complicate matters further, the Singaporean government is renowned for the efficacy of its language policies; government programs are responsible for convincing the ethnic Chinese population to largely abandon minority Chinese languages in favor of Mandarin in a single generation. The government's long-running "Speak Good English" Movement has similar aims for Singlish, which speaks to the importance of studying the language now. There is a growing body of work on the language, although it is mostly descriptive; formal linguists should take heed and begin collecting ungrammatical examples immediately.

For present purposes, the most interesting facet of Singlish is the prevalence of sentence-final discourse particles. These particles are one of the most salient features of Singlish. They also occur in all but the most formal levels of SStdE. These particles were likely borrowed from Hokkien or Cantonese, either directly or via Malaysian English, a close relative of Singlish (Deterding, 2007).

3.8.2 Data: *lah*

The most studied particle in Singlish is *lah*, exemplified below. Note that *lah* is also the most salient particle in public perceptions of the dialect.

(159) B: Where (are) you now?

K: a. In my office.

b. In my office *lah*.

≈ *I'm in my office, obviously...* (adapted from Kwan Terry 1978)

This example was originally employed to show that *lah* marks "impatience" or "annoyance" (Kwan Terry, 1978). It also seems to evoke the German unstressed *doch*, which conveys propositional content as given or uncontroversial, often paired with a 'corrective' attitude on the part of the speaker (Grosz, 2010). In fact, though, the picture with *lah* is much more complicated. Ler (2005, 2006) collects a bevy of uses for the particle:

Author	**Characterization of *lah***
Richards and Tay (1977)	"Solidarity and support"
Kwan Terry (1978)	Stressed: persuasive, conciliatory, or explanatory Unstressed: impatience or annoyance
Bell and Ser (1993)	Same split as above, called "long" and "short;" opposite characterizations
Lock and Low (1988)	Nine tonal variants, distinct "pragmatic functions"
Platt and Ho (1989)	Same functions as above, but tonal variation is just nuclear pitch movement

Table 3.1: Summary of analyses of *lah*

The variation in these analyses points to the fact that *lah* appears to be a highly flexible particle. What remains unclear, however, is what portion of this flexibility is located in the direct discourse effects of *lah*, and what is pragmatic reasoning on top of a more general discourse effect. The situation is not helped by the fact that the authors above do not employ negative examples. For example, even in (159), it is not whether the particle conveys obviousness, annoyance, or both. Gripes aside, these papers provide a wealth of data about *lah*, which makes a preliminary look possible.

3.8.2.1 The functions of *lah*

Ler (2005, 2006) provides a functional typology of *lah* based on the above work, concluding that *lah* is used to convey **solidarity**, **emphasis**, **obviousness**, **persuasion**, **friendliness**, and **hostility**. There is a great deal of variety across these uses (especially the last two), which suggests either the particles is polysemous, or it has a single core meaning that, when combined with pragmatics reasoning, derives all the above effects. I will conclude that the second option is correct. Note that among these six functions, some are incompatible even with each other. That said, it is not unusual for a lexical item with strong positive or negative associations to change polarization (c.f. *terrific*), so it is not particularly problematic for a functional approach that friendliness and hostility both appear on this list.

Solidarity *Lah* can convey solidarity in a way similar to American Engilsh *y'know*. It conveys that there is some shared premise underlying the utterance. Ler's (2005) approximations of what those premises are for the examples below are given in brackets:

(160) Don't be shy lah. [We are friends.]

(161) No use trying to hide our roots lah. [We are Singaporeans.]

Note that 'solidarity' here appears to mean both that the speaker and addressee share a property, and that they share knowledge of that shared property. In terms of arriving at a basic meaning for *lah*, this appears to be too strong, though, because not all examples involve this kind of solidarity:

(162) *A mother (A) and her daughter (B) had a disagreement on who is to buy Mandarin oranges. (It is customary for the Chinese to exchange Mandarin oranges when visiting*

during the Chinese New Year).

A: Then after that it's the Lunar New Year special lah.

B: So?

A: Ya lah, then during that period we can go what?

B: Cannot lah. Aiyah, when I wash my hair, I don't want to go out. Dirty my hair lah.

A: You bring one of them lah. (Ler's ICE-SIN-S1A-007)

In this example, *lah* is employed several times. The crucial use is B's second reply (*Cannot lah*), which seems to express the *opposite* of solidarity; the daughter has disagreed with her mother, but still uses *lah*, though toward what end remains unclear at present. In the discussion of the other functions below, similar counterexamples appear; one need only look at the subsequent function to find one.

Emphasis One function that is particularly opaque is what Ler (2005, 2006) calls "emphasis," although it appears to be closer to contrastive topic focus:

(163) Do you want to go? I'm not going lah. (Kwan Terry, 1992, p. 69)

(164) Normal doctors lah who are on our medical panel. [I.e., not specialists.] (Ler's ICE-SIN-S1B-073)

Notice in (164) that *lah* needn't occur sentence-finally. It is also able to mark noun phrases, in which case the phrase is the target of *lah*'s effect.

Obviousness In (159), *lah* seemed to express both obviousness and impatience or annoyance, but these two needn't be coupled together, as shown by (166):

(165) They generally don't take beef lah. [It's obvious; everybody knows that.] (Ler's ICE-SIN-S1A-023)

(166) I mean of course it changes lah. (Ler's ICE-SIN-S1A-065)

While (165) clearly expresses both obviousness and annoyance, (166), especially if paired with an appropriately 'soft' delivery, is not at all rude. This suggests that obviousness is more closely tied to *lah*'s meaning, and that annoyance is simply a common emotion to express when conveying the obvious.

Persuasion Persuasive uses of *lah* range from polite suggestions (167) to mild arm-twisting (168):

(167) Go to Chinatown lah. [Why don't you?] (Ler's ICE-SIN-S1A-007)

(168) Come with us lah. [Won't you?] (OED online, 2000)

Note also that these examples reveal that *lah* can occur with imperatives.

The persuasive uses of *lah* are clearly connected to the obvious uses. After all, acting as though something is obvious is often a successful way to persuade, don't you agree?

Friendliness Ler (2005) considers friendliness as a distinct function, but it is not clear how much these examples differ from the 'solidarity' function:

(169) Okay, doesn't matter lah. [It's all right; we're friends.] (Ler's ICE-SIN-S1A-091)

(170) Quite nice lah. [I'm your friend; consider my opinion.] (Ler's ICE-SIN-S1A-023)

If there is a difference, it is that the fact that the speaker and addressee share a property seems less important here than the idea that there is shared information. Alternatively, friendliness could simply follow from solidarity. That is in fact what the bracketed information wears on its sleeve in both cases.

Hostility Finally, the examples below show that *lah* can express outright hostility:

(171) If you want then it should be after this week lah. [Not earlier!]
(Ler's ICE-SIN-S1A-091)

(172) I don't want to eat lah. [Don't force me!] (Ler, source unknown)

Hostility is of a kind with annoyance and obviousness. These examples express that the speaker expects the addressee to know the reasons behind the utterances' content.

3.8.2.2 Problems for a functional analysis

There are at least two major objects that can be raised against stopping at the functional analysis above. First, Ler (2005) notes that each function can be conveyed even without *lah*. For example, both (173a–b) convey obviousness:

(173) a. I mean of course it changes lah! (Ler's ICE-SIN-S1A-065)

b. I mean of course it changes!

In a related point, recall from the table in (3.8.2) that Lock and Low (1988) claim that there are nine tonal variants of *lah*, each with its own function. In contrast, Platt and Ho (1989) note that every member of this functional inventory can be conveyed with 'tone' (i.e., intonational contour) alone in English. If these functions can be readily connected to tone, it breaks the idea that any specific function is inherent to *lah*.

A final argument against stopping at the functional view of *lah* is that many examples of *lah* show several of these functions active at once (Besemeres and Wierzbicka, 2003). Consider their example (174) below:

(174) *S has just received theater tickets from a colleague, K. The show starts soon, but S starts talking to K about reimbursement. K interrupts and says in a friendly but firm tone:*

Not now lah!

This utterance is friendly, but also expresses some annoyance. It also highlights the obviousness of the fact that S should be heading off to the show now, rather than dealing with the money.

In summary, a 'multifunctional' analysis of *lah* where each function is analyzed independently is not particularly enlightening, because it seems that *lah* only has a partial role in bringing about each of these function. Further, there does seem to be at least the hint of a common core to these functions, especially obviousness and solidarity. This core is that there is some additional meaning that the speaker expects his addressee to understand, and which serves as a crucial part of his ability to make his *lah*-marked utterance. Previous attempts at deriving a unified meaning for *lah* illuminate the rationale for this approach, so the next section briefly discusses two of these.

3.8.3 Previous approaches to particles

3.8.3.1 Paraphrastic account of *lah*

Besemeres and Wierzbicka (2003; B&W), citing Leibniz (1949), argue that testable hypotheses about language meaning "can be achieved only though substitution in context."

They explain that the meanings of expressions like discourse particles should be captured by paraphrases that are first-person and use vocabulary simple enough to be shared by all native speakers of the language. For the widespread English discourse marker *y'know*, for example, the authors offer the following paraphrase:

(175) B&W's PARAPHRASE OF *y'know*
I think that if you know this you can think about it
I want you to think about it now

I admit that I find this paraphrase unhelpful, but the authors choose it over formal accounts, as well as more complex paraphrases like the preliminary one offered by Jucker and Smith (1998) in advance of their final (formal) proposal:

> [...I]nstead of saying 'I assume that you could draw this inference yourself, but I just want to make certain you have it readily accessible because you will need it to interpret the rest of what I will say,' a speaker may say, *you know*.

B&W's argument against such a paraphrase is that speakers might not know what an *inference* is, or what *accessible* means. Presumably their distaste for formalism stems from a similar attitude. But unless we adopt an extreme level of linguistic relativism, there is no reason that a speaker's inability to paraphrase the meaning of a word precisely should preclude the speaker from knowing the word. Similarly, I can throw a ball without knowledge of how its trajectory decays, and I can navigate a three-dimensional world without knowing the exact distances between obstacles or how depth perception works.

Nevertheless, B&W's proposed paraphrase of *lah* is promising start for an analysis from the perspective of discourse offered in this dissertation:

(176) B&W's PARAPHRASE OF *lah*
I think you can know what I want to say

According to this paraphrase, the speaker who uses *lah* expresses a belief about his addressee's ability to understand the motivations behind his utterance. The specifics here are vague, but the analysis clearly points the way toward a discourse-evidential meaning for *lah*. But recall that the discourse model presented here has two significant that commitments can vary: they can be made relative to different evidential bases, and they can be made

a different strength ratings. The paraphrastic account seems to suggest that employing a special evidential is the way forward, but it worthwhile to consider the possibility that *lah* affects commitment strength as well.

3.8.3.2 Gupta's (1992) assertiveness scale

The most insightful work on *lah* tends to discuss the particle as part of a broader system of particles. Given, for example, how the meanings of *yo* and *ne* interact in Japanese, it is not surprising that this approach works. Singlish particles do no co-occur, but nevertheless seminal work by Gupta (1992) seeks to make sense of a broad range of Singlish discourse particles by positioning them on a scale of assertiveness.

Gupta's (1992) scale places 11 different particles (all with an array of functions as broad as *lah*'s) into three categories determined by how 'assertive' an utterance with the particle is:

(177) GUPTA'S (1992) CATEGORIES OF ASSERTIVENESS

i. **Contradictory** - 'Strong' assertion

ii. **Assertive** - 'Regular' assertion

iii. **Tentative** - 'Weak' assertion

The discussion of these terms is impressionistic, but some generalizations can be drawn. *Contradictory* particles occur with so-called 'Strong' assertions, not to be confused with how the term is employed in this dissertation. Rather, a strong assertion is one that is made baldly in the fact of contrary evidence, or even contradictory assertions from ones interlocutors. Falling *yo* in Japanese would likely be classified in this way (see Ch 3). *Assertive* particles occur with 'Regular' assertions, which approximate my *strong commitments*, or Farkas and Roelofsen's (2012) *actual commitments*. This is the default kind of assertion. Finally, *tentative* particles occur with 'Weak' assertions, which evokes my notion of *weak commitment*, but drawing a parallel between the two ideas would risk mischaracterizing the author's intent.

On the basis of a lengthy discussion of the functions for *lah* introduced in §3.8.2 above, Gupta (1992) places *lah* near the middle of her scale, as a squarely *assertive* particle. This is perhaps surprising, given that its obviousness and hostility functions seem rather 'forceful'.

For present purposes, however, the important point is that *lah* is not tentative marker because it does never "solicit[s] assent" from the addressee, a defining feature of the tentative group, and also of weak commitment under the discourse model proposed here. The way is therefore paved: *lah* introduces a shared base of evidence, but does not seem to alter commitment strength.

3.8.4 Discourse-evidential solution

The idea of a 'shared' base immediately evokes the Japanese particle *ne*, but its behavior in a number of environments, including rhetorical questions and imperatives, suggests a slightly different approach is warranted. The analysis is introduced below, building off the analysis for *ne*.

Lah and *ne* both invoke an evidential base that can be informally described as being shared with the addressee, but the formal characterizations are distinct. Recall the analysis for *ne* from §3.6.3 above:

(135) $[\![ne(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base $E_{\text{MIN}} = \{q | \forall X \in D : STR_X(q) \geq STR_{Sp}(q)\}$.

According to this definition, the base introduced by *ne* is one where every discourse participant could commit to the propositions in the base at least as well as the speaker can. This means the speaker's grounds for committing to φ are 'shared' in the sense that the base as a whole is assumed to be available to the addressee, should he want to use it. Further, the base encodes the idea that the addressee has minimal *relative authority* by leaving open the possibility that the addressee in fact has superior access to the base.

By invoking this base, the speaker of a *ne*-marked utterance publicizes an expectation that the addressee could—or perhaps even should—commit to the same proposition, given the same base or a similar one. This is why a *ne*-marked utterance always elicits a response, and why that response is expected to be agreement

This story, however, will not carry over to *lah*, which does not seem to involve relative authority at all. Consider (161) again, where the bracketed sentence is the additional

content communicated by *lah* in the specific context from which Ler (2005) culled the example:

(161) No use trying to hide our roots lah. [We are Singaporeans.]

Crucially, consultation with native speakers reveals that there is no strong-than-usual expectation here that the addressee *agrees* with the speaker's asserted content, the way we find with *ne* in Japanese. Rather, the 'shared background' here is rooted in the bracketed expression itself. *Lah* seems to highlight a single background assumption prominently, rather than characterizing the interlocutors' relative authority over the whole base.

The discussion of *lah* in previous literature often supports this view of *lah*. In the lead-up to their paraphrastic analysis, B&W note that (Luke, 1990, p. 95) claims that the Cantonese particle *la* is "a device for displaying the assumption of common ground." Note that the Cantonese particle is likely the origin for *lah* in Singlish. Even more revealing is Luke's later discussion, which neatly mirrors the purpose of discourse-evidential marking as argued by this dissertation:

> One way to [proceed is to] regard *la* as the grammaticalization of a solution to a particular kind of problem in the organization of a conversation, namely, the need to negotiate and achieve common ground. (Luke, 1990, p. 118)

Given the hypothesis that in a discourse-based commitment model, the way the common ground is negotiated involves managing commitments and settling issues, it appears that a discourse-evidential approach to *la(h)* is apt.

With the justifications in place, I now offer a preliminary characterization of *lah*. The particle deserves a more careful treatment, and in light of future evidence, the denotation provided here very likely will need to change. I suggest that unlike *ne*, where the characterization of the 'minimal authority' base E_{MIN} indicates the relative status of every proposition in the base, *lah* instead targets just one, particularly important proposition, designated here as q.

(178) $[\![lah(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base E^{LAH}, and $\exists q \in E_{\text{LAH}}$ s.t.:

i. The speaker could not commit to φ given a base $E^{\text{LAH}'}$
where $E_{\text{LAH}'} = E_{\text{LAH}} - q$, and

ii. The speaker and addressee are taken to be committed to q

According to this definition, the proposition q that is part of the evidential base *lah* introduces is very significant in two ways. First, the base is only viable as a basis for the speaker's commitment if it includes q, which means q is unusually important. Second, and more significantly, the speaker takes it that he and his addressee are already committed to q—in other words, q is common ground.

Given that the exact identity of q (i.e., its content) is obviously left unspecified by the *lah*-marked utterance, it is reasonable to doubt whether the proposition could be so directly targeted within the evidential base. There is precedent for this, however, in my treatment of dependent commitment, where the fact that the addressee has committed to the same proposition is the key component of the evidential base that makes it viable. The dependent base is therefore extremely similar to the *lah*-base, except that for the dependent base, the important proposition can be derived directly from the speaker's utterance, while for the *lah*-base, the search for the appropriate proposition is left to pragmatic reasoning.

This appeal to reasoning is not a weakness; if we wish to capture the wide variety of functions that *lah* can mark, the nature of the background assumption that *lah* targets must be flexible. By specifying that the proposition is something the speaker takes for granted should be treated as common ground, and that the proposition is the lynchpin in underwriting the speaker's commitment, the task is greatly simplified. Note that utterances with *lah* can be and often are informative. This is because even if the background assumption that *lah* highlights is shared, the at-issue content of the *lah*-marked utterance needn't be.[44]

[44]More generally, it is useful to have strategies that point out when there is common ground that the interlocutors are perhaps not regarding as equally valuable. Farkas (p.c.) suggests that the response particle *duh* perhaps serves this function. While *yes* commits the responder in the default manner (independently) and *oh* commits him dependently, *duh* highlights the fact that the responder already took the antecedent commitment to be common ground.

3.8.4.1 Specific comparisons

Next, this section provides specific comparisons of Japanese *ne* and Singlish *lah*, and how the analysis above captures the differences between them. A small amount of original data is included here, based on reports from native Singlish speakers living in Los Angeles, California.

Backchannel-eliciting use One extended use of *ne* is as a tool for eliciting backchanneling. When employed this way, *ne* can occur at any phrase boundary, unlike the normal case discussed here, where it must be sentence-final:

(179) *Souta is excitedly telling a story about his morning.*

A: Asa ne? Wafferu-ga tabetakatta kara ne? Syowaa ato-de...
this morning NE? waffles-NOM wanted to eat because NE? shower after-P...
'This morning, right? I wanted to eat waffles, right? So after showering I...'

B&W suggest that *lah* can also be used to elicit backchanneling. While *ne* prompts the addressee to agree that the speaker should continue his utterance, *lah*'s backchanneling use is more tightly connected to its formal analysis. Functionally, *lah* used in this way is extremely similar to the previously-mentioned case of (non-Singaporean) English *y'know*; it urges the addressee to remain involved in the conversation by constantly reaffirming the abundance of common ground between the interlocutors.[45]

Imperative uses Recall that core imperatives with *ne* in Japanese are usually perceived as 'softer'. Interestingly, (180) reveals that this is really only the case when deontic authority rests with the speaker:

(99) *Ayaka is tutoring Souta on inquisitive semantics.*

A: Tyanto benkyoo site kudasai ne./?
properly study do please NE
'Please study properly, ok?'

(180) *Souta's time is almost up.*

[45]It is unclear whether there is still one particularly important proposition with common ground status under this backchanneling use.

S: Moo sukosi jikan-o ataete kure ne./?
more a little bit time-ACC give request NE
#'Give me a little more time time, ok?'
'I need a little more time, alright?'

Native speakers report that while (99) is a polite request, (180) is, "a little off-putting." The relative authority analysis of *ne* easily spins a story to account for this effect: When the speaker uses *ne* with an imperative, she indicates that her deontic authority is lower than her addressee's. In a situation where both participants know this arrangement is false and that the imperative user's authority is higher, this approach often comes off as soft, or kind. Alternatively, if the *addressee* believes she has more authority than the imperative user, as in (180), the use of *ne* is presumptuous.

Similarly, core imperatives with *lah* in Singlish often convey a pleading tone:

(181) Give me more time *lah*!
'Please give me more time, you know I need it.' (B&W, p. 23)

As indicated by the gloss, the additional meaning conveyed by *lah* here clearly does not indicate that anyone has equal or lesser authority than anyone else. Rather, the utterance acts more like a request for understanding by publicizing the expectation that the addressee should already understand that the request is a valid one. This follows from the analysis of *lah*, because all that the evidential base requires is that the major premise behind the asker's utterance is know to all.

Use with rhetorical questions Recall that *ne* is infelicitous with information-seeking questions (182) and rhetorical questions (183):

(182) Konna hon, kau ka *ne?/.
this.kind.of book buy Q *NE
'Are you going to buy a book like this (*huh)?'

(183) Konna hon, dare-ga kau ka *ne?/.
this.kind.of book who-NOM buy Q NE
'Who the hell would buy a book like this, (*am I right)!?'

More fully, (182) is a licit string, but it does not express a question. Rather, the speaker is taken to be musing to herself, e.g. *I wonder if (he's) going to buy a book like that.*. Still, there is a basic incompatibility between *ne* and true questioning acts.

Interestingly, B&W claim that rhetorical questions with *lah* are felicitous—but the native speakers I've consulted say otherwise:

(184) ? What to do lah!

The first issue here, however, is that this utterance is *not* a rhetorical question. Rather, it appears to be an exclamative. These utterances do not raise issues the way that questions do, and without propositional alternatives placed on the table, it is misleading to call it a rhetorical question.

When provided with more standard examples of rhetorical questions with *lah*, native speakers universally reject them. Regardless of whether the implied answer is negative (185a–b) or positive (185a–b), the result is the same.

(185) a. You think I'm stupid #lah?

b. Have I ever let you down #lah?

c. Don't you trust me #lah?

Here, the reasoning for why *lah* and *ne* cannot appear with rhetorical questions converges. I have claimed asking a rhetorical question is a highly idiosyncratic act that involves the speaker pretending that the question's issue is not settled, even though all those involve know that it really is. The result is that the speaker asks the question, raising the issue as normal, but then the issue is immediately resolved with no need for a response. "Raising the issue as normal" in the case of a *ne*– or *lah*-marked utterance, however, involves introducing an evidential base that undermines the question being asked.[46] The act of asking a rhetorical question and feigning ignorance is therefore incompatible with formally marking a discourse move that entails lack of ignorance.

[46]The discussion of rhetorical questions in Chapter 4 presents a more detailed view of this intuition. Specifically, I argue that when the speaker asks the question, he also commits to the union of its possible answers, in a step I call the *commitment to openness*. This commitment is otherwise trivial, especially for a polar question where the commitment to openness is a commitment to $[p \vee \neg p]$. It would therefore be rather strange for the speaker to condition such a trivial commitment on a special 'shared' base of either sort discussed here. At best, it is a redundant action.

Finally, I conclude this short analysis of *lah* with a hint of the next step: While *lah* is not licit with rhetorical questions, using the discourse particle *ah* is in fact the most natural way to do so, as in (186a) and (186c):

(186) a. You think I'm stupid ah?

b. Have I ever let you down or not?

c. Don't you trust me ah?

Also, note the rhetorical use of the alternative polar question in (186b). This is not felicitous as a rhetorical questions in American or British English. This has interesting ramifications for the discussion of biased polar questions presented in the next chapter.

3.8.5 Conclusions about *lah*

This section has provided a preliminary analysis of the Singlish particle *lah* in order to compared it with Japanese *ne*. It was shown that *lah* introduces an evidential base whose most important premise is taken to be common ground, which contrasts with *ne*, where it is the proposition expressed by the utterance itself that is marked as shared. Both particles support the view of discourse advanced throughout this dissertation, namely that discourse markers existing to facilitate the management of commitments and issue-settling, as a way toward the Stalnakerian goal of growing the common ground.

That said, there is much more to say about *lah* specifically, and discourse particles as a class of lexical items. For *lah* itself, recall that Platt and Ho (1989), at least, suggest there are different tonal variants of *lah*. This might easily lead to an elaboration along the lines of analysis of Japanese *yo* adopted here. More broadly, the competition between *yo*, *ne*, and *yone* in Japanese, and the fact that *lah* is not licit with rhetorical questions while *ah* is preferred both suggest that a language's discourse markers should be investigated as a system. Given that particles compete, it is extremely likely, for example, that two languages might feature particles that are formally identical, but whose behaviors are slightly different because they compete in different environments. Such variation would be welcome under the view advanced here. Singlish is likely an excellent place to start this investigation—and of course Japanese still has a plethora of discourse-evidential markers waiting to be examined.

3.9 Conclusions and next steps

Empirically, this chapter has presented a characterization of some of the data concerning the Japanese sentence-final particles *yo* and *ne* that had previously gone undiscussed in the English-language formal pragmatic literature on these particles. Novel findings includes the fact that *yo* does not require absolute hearer-newness to be felicitous, and that the particle string *yone* exhibits some but not all of the properties of what appear to be its constituent parts. Building on McCready (2008b) and Davis (2011), a new analysis was proposed to capture these findings within the model developed in Chapter 2 that ties the particles to the notion of *relative authority*. I have treated *yo* and *ne* discourse evidentials that characterize evidential bases in terms of relative authority, and from this derived and the majority of the novel facts about felicitous and infelicitous use of these particles presented above.

Under the final account, a *yo*-marked utterance requires that any commitments the speaker makes in the course of the utterance invoke an evidential base whose content the speaker is maximally authoritative on, relative to his interlocutors. A *ne*-marked utterance requires that the speaker is minimally authoritative. Taken together, a *yone*-marked utterance requires that the speaker and his interlocutors be equally authoritative. Pragmatic competition delineates these situations, and also explains cases where *yo* can be used without implicating the speaker is more authoritative than his interlocutors.

Along the way, I clarified Gunlogson's (2008) notion of *sourcehood*, which I argue is just one way that language collapses the gradient notion of authority into a binary system of source and dependent. While polarity particles like *yes* and *oh* illustrate sourcehood at work, *yo* and *ne* illustrate relative comparisons of interlocutors' authority.

The discussion points to the more general idea that notions like epistemic and deontic authority are important factors in the reasoning behind the discourse moves of a conversation, but these gradient, non-linguistic quantities should not themselves appear in a formal discourse structure. Instead, languages grammaticalize specific interpretations of gradient information into discrete sets. Nonetheless, the logic by which the gradient information is collapsed *is* a linguistic quantity, as formalized here as the character of an evidential base. Coming full circle, this discussion highlights the distinction noted at the end of Chapter 2 about differentiating between the content of a base, which is linguistically opaque, and its character.

Still, questions remain. Besides authority, what kinds of characterizations of evidential bases do discourse evidentials permit? This dissertation will not arrive at an exhaustive list, because likely no such list is possible. Rather, as highlighted by the tentative comparison between *ne* and the Singlish particle *lah*, what seems to unite evidential bases is their function in the management of discourse. Chapter 5 discusses this generalization about how bases may deviate from the default more fully.

A more pressing question, though, arises from the fact that we have still not arrived at a clear notion of what the default evidential base is, beyond the fact that implicates authority independent of a speaker's interlocutors. Chapter 4 takes up this issue with reference to biased polar questions in English.

Chapter 4

An evidential approach to biased questions

4.1 Introduction

English provides a number of ways to ask about whether a proposition or its negation holds. The most common of these are *polar questions*, speech acts that employ the interrogative sentence form for this purpose. In this chapter, I contrast the discourse effects of a few kinds of these polar questions: Positive polar questions (PPQ), high negation polar questions (HNPQs), and various kinds of tag questions (TQs), as exemplified below.[1]

(187) *Speaker and addressee observe from afar as Jack studies a map.*

a. Does Jack know where Barcelona is? *PPQ*

b. Doesn't Jack know where Barcelona is? *HNPQ*

c. Jack knows where Barcelona is, doesn't he? *TQ*

It is generally agreed that while polar questions like these differ in their syntax and overall discourse effects, they share a common semantic core that induces a choice between two mutually-exclusive alternative propositions. In the examples above, these alternatives are that Jack knows where Barcelona is, and that he doesn't. Crucially, however, the various

[1]Note that I will exclude from discussion the class of so-called low-negation polar questions (LNPQs), which employ constituent negation (e.g., "Does Jack not know where Barcelona is?"). See Farkas and Roelofsen (2012) and especially Anderbois (2011) for a detailed account of these sorts of questions. Also, PPQs are included only for their relative simplicity in terms of bias compared to HNPQs; I treat PPQs as the baseline case which other polar questions differ from.

questions present the asker as holding different expectations about the alternatives; while (187a) gives no clues about the asker's expectations, in (187b–c) she seems to expects that Jack does know—or at least *should* know.[2] 'Favoring one alternative over another' is an informal characterization of *bias* as I will use the term here.

Formalizing bias, which will involve making explicit what it means for a speaker to 'favor' an answer, is the immediate goal of this chapter, but it is important to keep in mind the commitments I have made throughout this dissertation, because it shapes the kind of analysis we expect should be possible. In particular, I have argued that while the default discourse effects of utterances derive from the utterance's literal content, non-default discourse effects result from changes to the evidential base that underwrites the commitments resulting from an utterance. This was demonstrated in the case of default versus rising declaratives in Chapter 2, and bare versus particle-marked utterances in Japanese in Chapter 3. Returning to English, then, the expectation this approach imposes on the present investigation of biased questions is that if the differences in the discourse effects of biased questions have anything to do with evidence (which I argue they do), then those effects must be amenable to an analysis presented in terms of differences in commitments, and especially the evidential bases underwriting those commitments.

That said, biased questions like (187b–c) do not intuitively seem to introduce any unusual evidence sources, which immediately raises the question of how the commitments they spawn will vary. By investigating what underlies their biasing commitments, the characteristics of the *default evidential base* will be revealed, lending further support to the idea that characterizing evidential bases in terms of the ways they diverge from the default base is on the right track.

This chapter unfolds as follows. §4.2 delves into two kinds of bias which appear to underly the special discourse effects of biased questions. In §4.3–4.4, I focus on HNPQs like (187b), which have been studied extensively and offer the clearest window into what it means for a question to be biased. This section elaborates the discourse model introduced

[2]What it means to 'expect' an answer is a complicated notion. This issue will return shortly, but for now it is enough to note that we must be careful to make a difference between an epistemic expectation based on the likelihood of a particular outcome, and a more bouletic flavor of expectation (van Rooy and Safarova, 2003). As will be shown for HNPQs, I can desire one outcome while granting that the other is more likely to be accepted.

in the previous chapter and culminates in a novel analysis of these questions which preserves the insights of previous analyses and accounts for some novel data. In §4.5, I show that the bias conveyed by TQs like (187c) is not identical to that of HNPQs. The discussion illuminates the ways that bias in polar questions can vary. In §4.6–4.7, I contrast the conclusions about HNPQs and TQs, and discuss the ramifications for my discourse model. Briefly, while the dimensions of variation introduced in §4.3–4.5 seem overly permissive at first glance, the model in fact predicts exactly the set of licit bias configurations that are attested in English. §4.8 concludes.

4.2 Kinds of bias

The literature on polar questions recognizes at least two major sources of bias, and that the biases that result from these sources are, in some cases, independent (Ladd, 1981; Büring and Gunlogson, 2000; Romero and Han, 2004; Anderbois, 2011; Roelofsen et al., 2012). Consider the contrast in the reasoning behind the bolded tag questions in the examples below, featuring two kinds of tag questions (to be discussed in depth in §4.5).

(188) *Happy hour approaches; Zabi and Kazuko need a drink. Zabi believes the brewery is on Swift Street.*

Z: Do you remember which side of town the brewery is on?
It's on the Westside, isn't it?

(189) *Happy hour approaches; Zabi and Kazuko (who are on the UCSC campus, near the Westside and far from downtown) need a drink. Zabi has no idea where the brewery is.*

Z: Do you remember which side of town the brewery is on?

K: There's no rush, we can get there in like 5 minutes.

Z: (Ah, so) **it's on the Westside, isn't it**.

Both of the bolded utterances are intuitively biased for the positive answer, but for different reasons. In (188), Zabi is biased for p because of pre-existing beliefs about the world (specifically, the layout of Santa Cruz). In (189) by contrast, Zabi reaches his conclusion

based in part on his knowledge of Santa Cruz, but the direct catalyst is Kazuko's utterance. That new information—that the brewery is only a short drive away—guides Zabi's bias.

To express speaker bias is to make public one's expectations about the future of the discourse, but those expectations can be conditioned on different evidence. Based on the contrast in (188–189), I follow Büring and Gunlogson (2000), Reese (2007), Sudo (2013), and Roelofsen et al. (2012) in claiming that there are in fact at least two kinds of bias, *speaker* and *contextual*.[3] I depart from these authors, however, in stressing that exploring what these two kinds of bias share is equally as important as determining what makes them distinct. In other words, a satisfactory analysis of bias should explain not only what the different sources of bias are, but why it is those sources in particular that need characterizing in the first place.

4.2.1 Speaker bias

Speaker bias concerns a speaker's beliefs about what is true, normal, or desired in the context. This bias is especially evident in (190b–c) below.

(190) *The speaker wants to watch some home movies.*

a. Does Milo have a VCR? *PPQ*

b. Doesn't Milo have a VCR? *HNPQ*

c. Milo has a VCR, doesn't he? *Rising TQ*

d. Milo has a VCR, doesn't he. *Falling TQ*

While all four of these sentence can be employed to elicit a response from the addressee about Milo's entertainment system, only (190b–c) reveal that the speaker has some private reasons for thinking that the response will be positive based on knowledge that he brings to the conversation. (190a) is neutral on this point, while (190d) is most natural as the conclusion of an inference process.

It should be noted that while speaker bias has been called 'epistemic bias', this is something of a misnomer. As demonstrated by the biased utterances below, the source of the

[3] Sudo (2013) revealingly calls contextual bias "evidential." I avoid this term because I consider both kinds of bias are in fact based on evidence and are formalized akin to evidentiality. I also adopt the term 'speaker bias' over Sudo's (2013) 'epistemic bias' in part because contextual bias will also be shown to involve the speaker's epistemic state.

speaker bias can sometimes be deontic or bouletic (Reese, 2007; Huddleston and Pullum, 2002).

(191) a. Aren't you ashamed of yourself?[4]

b. Don't you like it?

(191a) expresses a judgement about the addressee, while (191b) can in some contexts express the speaker's hopes. What these examples share with the biased examples in (190b–c), however, is that in all cases, the bias crucially originates within the speaker; it is the speaker's preexisting knowledge, beliefs, or desires that drive the need to express bias.

4.2.2 Contextual bias

Contextual bias, on the other hand, is not rooted in speaker's private beliefs about the world. Instead, it is 'read' off of the context itself.

(192) *B hands A some right-handed scissors to bring to B's friend Ellie.*

A: Oh, Ellie can't use these.

B: She's left-handed?

Here, B's bias toward the proposition that Ellie is left-handed is based on what he has gleaned from A's utterance. Of course, for B to arrive at an appropriate conclusion, she must interpret A's utterance against a body of background knowledge. The bias B expresses, however, could not have arisen without contextual support, and is therefore not *rooted* in speaker's epistemic state the way speaker bias is. Based on examples of this sort, Büring and Gunlogson (2000) argue that contextual bias is based on *contextual evidence*:

(193) CONTEXTUAL EVIDENCE

Evidence that has just become mutually available to the participants in the current discourse situation.

This definition encodes a notion of immediacy not found in the spekaer bias cases. This has lead some authors (notably Gunlogson 2003) to construe contextual bias not as a quality directly connected with a speaker's utterance, but a state of the context. So while speakers

[4]This question is most naturally rhetorical, which perhaps weakens Huddleston and Pullum's (2002) claim.

can remark on this bias and call attention to it, it is not their own in the sense that speaker bias is.

Gunlogson (2003) claims that a context is biased for p just in case p's truth is an active issue in the discourse that is not yet settled (i.e., p is *controversial*), adding p to the common ground would result in a monotonic update, and adding $\neg p$ would be non-monotonic. In other words, a context is biased for p if p needs to be settled and only settling it in the affirmative would avoid having to revise any existing commitments. This is a very strong view of contextual bias, because it requires that the truth of p must be entailed by the evidence available in the context. It therefore also comes with the assumption that if all interlocutors are in agreement about the evidence, they must also be in agreement about the biased nature of the context.

I construe contextual bias in more individualistic terms, where a speaker's expression of contextual bias only reveals his individual reasoning about the context. More explicitly, I will argue that contextual bias introduces an evidential base that crucially includes what the speaker construes as relevant contextual evidence, without which the base is not viable.[5] There are reasons to think this view of contextual bias is a more natural position than the alternative. Consider, for example, a possible continuation for the situation above:

(192′) *B hands A some right-handed scissors to bring to B's friend Ellie.*

A: Oh, Ellie can't use these.

B: She's left-handed? ($= p$)

A: No, she's allergic to whatever they make these kinds of handles out of.

Here, it is not the case that the context requires Ellie to be left-handed, as shown by A's final response. The context is therefore not biased in Gunlogson's (2003) sense, because updating the common ground with either p or $\neg p$ will be monotonic. And yet, B does manage to show some form of contextual bias with his rising declarative utterance. His bias involves reasoning over the information available in context, and clearly reflects his individual conclusions, rather than broader conclusions about the context, precisely because he is wrong.

[5]This characterization evokes my notion dependent commitment; in both cases, speaker signals that the basis for his commitment is not entirely under his control.

As a further illustration, the examples below demonstrate an interesting contrast between an infelicitous HNPQ, and a felicitous one ending with *then*:

(194) *Terra's house has a bathroom between the kitchen and the living room, and Cassie knows this.*

T: My house flooded. The kitchen and the living room are wrecked.

C: a. # Isn't the bathroom flooded?

b. Isn't the bathroom flooded then?

The oddness of (194a) arises from the fact that it seems to suggest that Cassie has some independent source of information concerning the state of the bathroom. (194b), on the other hand, is fully felicitous, but the inclusion of *then* makes it a (surface) question about Cassie's reasoning, rather than the state of the bathroom itself. In (194b), Cassie clearly expresses some contextual bias based on Terra's utterance, but the context itself is still not 'biased' in Gunlogson's (2003) sense. Note that the final analysis will readily derive the contrast above.

As a final clarification before moving on toward the analysis of HNPQs, however, note that the evidence required to express contextual bias can come from the context at large, not just the linguistic discourse:

(195) *Ellie is painting with her left hand.*

B: Ellie's left-handed?

A: Nope, ambidextrous.

Just as before, B expresses contextual bias for Ellie's being left-handed that turns out to be wrong. His evidence comes from direct observation of the context, however, rather than from any leading utterance of A's.

4.3 Bias in high negation polar questions

The discussion above has employed an amalgam of rising declaratives, HNPQs, and TQs. In this section, I focus in on HNPQs, determine how the two biases above interact, and propose a formal account of their discourse behavior in that light. Note that the discussion here does

not address the classic inner– vs outer-negation ambiguity popularized by Ladd (1981). I focus for now on the more common outer-negation reading. See §4.4 for an introduction to Ladd's (1981) ambiguity, reviews of existing approaches to it, and a discussion of how that issue fits with the one at hand.

4.3.1 HNPQ and the two biases

Sudo (2013) and Roelofsen et al. (2012) argue HNPQs convey both speaker and contextual bias, which is to say that they are fully licensed only when evidence of both sorts can be brought to bear on the utterance. Reiterating, these evidence sources are *speaker belief*, which is generally rooted in the speaker's prior epistemic state, and *contextual evidence*, which involves induction over the information recently made salient in the context of utterance.

To express bias of either kind is to present oneself as having evidence of the appropriate sort, either in favor of a proposition or against it. Intuitively, a particular evidence source could lend support to either a proposition or its negation. In the former case, the evidence is *positive*, and in the latter it is *negative*. I also suggest that an evidence source could be *neutral*, a concept we will return to later.

Sudo (2013) shows through a series of carefully constructed contexts that HNPQs have very specific requirements for the polarity of the of evidence from speaker beliefs and contextual evidence. These tests are adapted below. In (196), the speaker's HNPQ is followed by several continuations, each of which expresses a different kind of speaker belief. Note that for this and all future examples involving HNPQs, '*p*' will refer to the proposition without cliticized negation; for example, for *Doesn't Milo have a VCR?* below, *p* is *Milo has a VCR*. Borrowing from the tag question literature, we can refer to this proposition as the HNPQ's *anchor*.

(196) HNPQs WITH VARYING SPEAKER BELIEF

Doesn't Milo have a VCR?

a. ✓. . . I think he does. *pro-p* ✓

b. # . . . I have no idea either way. *neut-p* #

c. # . . . I don't think he does. *anti-p* #

What the contrasts above reveal is that HNPQs are only licit with **positive** speaker belief for the anchor p. Another way to phrase this is that HNPQs themselves *publicize* positive speaker belief.[6]

In terms of contextual evidence, felicitous use of HNPQs requires that the contextual evidence be **non-positive**:

(198) *A and B are crafting a list of left-handed people to invite to an unusual party.*

a. A: So we have Ellie, Reuben, and Faye so far. Oh, look over there—Milo is writing with his left hand.

B: # Isn't Milo left-handed too? *pro-p #*

b. A: So there's Ellie, Reuben, Faye. . . who else?

B: ✓Isn't Milo left-handed too? *neut-p ✓*

c. A: So we have Ellie, Reuben, and Faye. We're done!

B: ✓Isn't Milo left-handed too? *anti-p ✓*

Contextual evidence is provided in this examples via A's utterance, which needn't even be accepted before it can serve as evidence of the appropriate sort. Given that the dependent base, as introduced in Chapter 2, relies crucially on the fact that the addressee has made a commitment (rather than on the literal content of the commitment), it is unsurprising that the fact that a discourse move has been made is sufficient to act as evidence.

To summarize the generalizations thus far, Sudo (2013) and Roelofsen et al. (2012) argue for the following:

(199) EMPIRICAL GENERALIZATION FOR HNPQS, #1: **Bias conditions** (to be revised)

a. HNPQs require **positive** speaker belief for the anchor proposition.

[6]Note that these examples are worse if the admission of bias comes prior to the question. This initially seems to be at odds with the claim made here, but I argue it is actually unproblematic because of the kinds of bias involved. This is evident from the contrast below:

(197) a. ? I think Milo has a VCR. Doesn't he have one?

b. ✓I thought Milo had a VCR. Doesn't he have one?

In (197a), the speaker asserts current speaker belief, and the result is degraded compared to (197b), which instead expresses *prior* speaker belief. This will follow from the final analysis presented here.

b. HNPQs require **non-positive** contextual evidence for the anchor proposition.

The next section nuances this view.

4.3.1.1 Relating speaker belief and contextual evidence

Sudo (2013) and Roelofsen et al. (2012) present the above findings as two separate requirements. I suggest taking the idea further, by claiming that the two bias requirements aren't independent; what's needed is a *conflict* between them. Consider the example below, where Ned possesses the necessary speaker belief to license his question, and Prof. Hemanti's exclamation presents no particular evidence against this claim:

(200) *Ned but not Prof. Hemanti knows that Jessup is habitually late to class. There's still time until the appointed start time, but class can't begin until everyone is there.*

H: We have such a busy day ahead!

N: *(Turning to a classmate)* # Won't Jessup be late?

The bias conditions as described above—speaker beliefs for the anchor and a lack of contextual evidence against it—seem to be met here, and yet Ned's question is not particularly natural (although see (205) below). This is because the fact that it's going to be a busy day in class has little bearing on Ned's question. We can, however, replace the professor's utterance with a similarly 'neutral' one in terms of contextual evidence, and make the question licit:

(201) *(Same context as above.)*

H: I want to start on time today.

N: *(Turning to a classmate)* Won't Jessup be late?

From Prof. Hemanti's perspective, her utterance above accomplishes nearly the same goal as her utterance in (200). For Ned, however, the professor's utterance is now incompatible with the anchor proposition, thereby licensing his HNPQ.

More generally, Prof. Hemanti can produce any of a number of superficially similar utterances about the start of class, but the HNPQ is only licit if in doing so, she provides contextual evidence suggesting she does not share Ned's (private) pessimism about Jessup's punctuality:

(202) *(Same context as above.)*

H: I hope we can start on time today.
ps: starting on time is possible

N: *(Turning to a classmate)* Won't Jessup be late?

(203) *(Same context as above.)*

H: I wish we could start on time today.
ps: starting on time not possible

N: *(Turning to a classmate)* # Won't Jessup be late?

As demonstrated by Anand and Hacquard (2013), among others, *hope* in (202) contrasts with *wish* in (203) in terms of presuppositions about the possibility of the complement. Where *hope* requires that its complement is possible, *wish* is counterfactual and so implicates that its complement is not possible. To make Ned's question licit, Prof. Hemanti's public desires must conflict with Ned's expectations, i.e., Ned must be able to conclude that the professor at least entertains the possibility that Jessup will arrive on time. Note that *want* in (201) differs from *hope* and *wish* in that it makes no demands of the speaker's expectations. Nevertheless, it conflicts sufficiently with Ned's speaker beliefs to license his HNPQ, unlike the basically irrelevant (but neutral) contextual evidence in (200).

This discussion leads us to revise the generalization in (199) in favor of (204) below:

(204) EMPIRICAL GENERALIZATION FOR HNPQS, #1: **Bias conditions** (revised version)
HNPQs require a **conflict** between **positive** speaker belief and **negative** contextual evidence for the anchor proposition.

This generalization strengthens the negative contextual evidence requirement from needing non-positive evidence to negative evidence, with the understanding that in some cases, (e.g. (201)), evidence that seems neutral can negatively affect the speaker's expectations.

Further evidence that HNPQs require a salient conflict is that these questions are infelicitous "out of the blue" whenever the nature of the conflict cannot be accommodated. Compare (205a–b), which might both occur with no prior conversation in the halls of Stevenson College.

(205) *Jorge walks into the department office and asks Ashley:*

a. ✓Isn't Jim on campus today?

b. ?? Isn't Jim going to the supermarket today?

(205a) is a natural question in this context because we can easily imagine Jorge wandering the halls looking for Jim and not finding him, which constitutes negative contextual evidence for Jim's presence. (205b) is in fact potentially licit as well, but it requires elaborating the context so that the general facts surrounding Jim's trip to the supermarket (at least) are known to Ashley. In other words, (205b) becomes licit as soon as Ashley has enough shared background knowledge to make a reasonable guess about the conflicting evidence Jorge has encountered.

These data make it clear that a conflict in bias is central to the felicity conditions for HNPQs. By way of comparison, we can imagine a simpler view where HNPQs and other biased questions merely express a level of commitment lower than that of an assertion, and the dueling bias requirements are just one way that the necessary conditions for reduced commitment could be met. This straw man cannot stand, however, because it would predict that HNPQs would appear more widely; any of the examples above should be potentially licit.

4.3.1.2 Bias is weak commitment

Before moving on, one more point bears mentioning. The intuitive notion of bias I have adopted so far in this section is couched in terms of conflicting speaker expectations; speaker beliefs push the speaker one direction, and contextual evidence the other. While this is a useful shorthand, we must be careful not to take it too literally. This is because expectations alone are a weak driver of discourse moves; there is nothing wrong with being unsure what way the discourse will progress—especially given that HNPQs are presumably information-seeking questions at their core, and answering the question will resolve uncertainty. To actually capture the complete discourse effects of HNPQs and other biased questions, I take a stronger view of what a conflict in bias entails. I argue that a bias conflict does not just involve revealing uncertain expectations; instead, it involves making incompatible, evidential commitments about the speaker's confidence that a given proposition is true.

In Chapter 2, it was shown that commitments can either be *strong* or *weak*, where

strong commitments are the default case, and weak commitments are those underwritten by the *weak base*, whose defining characteristic is its reliance on evidence that the speaker does not (yet) command:

(206) (41) WEAK COMMITMENT

Let q stand for some proposition that serves as evidence for p. When a speaker weakly commits to p, he conditions his commitment on the base E_{WEAK} s.t.:

i. $q \in E_{\text{WEAK}}$ but

ii. the speaker could not felicitously commit to q (given any base) and

iii. the speaker could not felicitously commit to p given a base $E_{\text{WEAK}'}$ where $E_{\text{WEAK}'} = E_{\text{WEAK}} - q$.

The weak base is defined as crucially containing a piece of evidence that the speaker lacks the authority to commit to, meaning that a commitment conditioned on the weak base is untenable if it is not backed up by further evidence in support of the proposition committed to. The base therefore anticipates both future directions for the discourse where the commitment has been bolstered, and future directions where it has been invalidated. In contrast, a speaker who makes a default (strong) commitment performs a discourse move that only anticipates the continued viability of that commitment. In other words, a default commitment is expected to be relatively immutable in light of future discourse moves, while a weak commitment is fragile due to the inherent mutability of the base.

It is therefore immediately clear that between these two options, weak commitment is the more appropriate sort of commitment to use to express bias. To capture bias using a default commitment would be to deny that expressing bias for p has distinct discourse effects from asserting p—a poor outcome. Rather, bias involves commitments characterized as weak.

The fact that weak commitments project the possibility of their own invalidation imbues them with one further, useful property in light of the bias conflict generalization, namely that as long as the bases are not identical, a speaker can weakly commit to both p and its negation. When a speaker makes conflicting weak commitments, he presents himself as nearly able to commit to either possibility, if the appropriate evidential base could be bolstered. This would be inconsistent if both commitments were underwritten by the

same evidence, but if the two commitments are made given different sources of evidence, the speaker instead presents himself as being of two minds, and 'on the hunt' for evidence to bolster either base. As will be shown, this puts the speaker's DC into a kind of local crisis and drives the full discourse effects of these questions.

Note that some form of commitment is crucial to overcoming the problem mentioned above. We can imagine an alternative to the sorts of commitments discussed so far, whereby a speaker mentions a proposition without staking any sort of claim about whether he will act in accordance with it. This kind of 'negligible' commitment would indicate a functional lack of commitment, while still associating the commitment's propositional content with the speaker. At first glance, this is appealing as a way to express bias, but it is too weak to serve as a driver of future discourse moves in exactly the bias conflict cases discussed here. This is because negligible commitment as just described is no more forceful than the preceding discussion of "expectations;" just because a speaker has mentioned p and $\neg p$ and suggested that this mention was spurred by two different sources of evidence does not lead straightforwardly to the requirement that the evidence is strong enough to be conflicting. The conflict results from the fact that speaker's continued ability to commit is at put stake.

4.3.2 A wrench: SB and CE are not bases

The most straightforward solution for modeling the bias expressed by a HNPQ questioning p, then, involves contrasting the speaker's weak commitment to p given an evidential base of his speaker beliefs (E_{SB}), and his weak commitment to $\neg p$ given an evidential base of contextual evidence (E_{CE}). I.e., the speaker expresses positive speaker bias and negative contextual bias. Unfortunately, more needs to be said on this point because of the fact that expressing contextual bias generally requires a good deal of speaker belief as well; the two bases cannot be so easily juxtaposed.

The examples repeated below illustrate the distinction between contextual evidence and speaker belief. In both examples, the bolded phrase expresses bias for the truth of the anchor, but we understand the evidence behind that bias to be quite different. Note that in both cases, the bolded utterance is a falling tag question; the felicity of these utterances would differ with rising intonation.

(207) *Happy hour approaches; Zabi and Kazuko need a drink. Zabi can't easily recall where the brewery is.*

Z: Do you remember where the brewery is? Nevermind; **it's on Swift, isn't it**.

(189) *Happy hour approaches; Zabi and Kazuko (who are on the UCSC campus, near the Westside and far from downtown) need a drink. Zabi has no idea where the brewery is.*

Z: Do you remember which side of town the brewery is on?

K: There's no rush, we can get there in like 5 minutes.

Z: (Ah, so) **it's on the Westside, isn't it**.

In (207), Zabi's uncertain speaker beliefs form the knowledge base underwriting her bias; we understand the purpose of her question is to 'check' the validity of her pre-existing belief. In (189), the purpose of the question is to check the validity of a more immediate inference; when Kazuko says the brewery is close by, Zabi uses this information to arrive at her inference.

Crucially, however, Zabi's pre-existing beliefs are also involved in this second utterance; it is only by combining the immediate evidence (i.e., the information contributed by Kazuko's utterance) with world knowledge about the geography and commercial zoning of Santa Cruz that the inference about the brewery's location is possible.

In fact, it seems speaker beliefs are always assumed to be actively involved in inferences, unless this assumption is overtly denied, as in the first clause of (208a):

(208) *Jessa is advising Shoshanna about a potential suitor's online dating profile*

J: a. **Based on this alone, Dylan sounds nice enough**—but he's probably a crazy stalker.

b. # **Based on this alone, Dylan is nice enough**—but he's probably a crazy stalker.

In the bolded portion of Jessa's utterance in (208a), she indicates that at least some of her speaker beliefs—presumably those concerning the general characteristics of men on internet dating websites—do not underwrite the claim that Dylan sounds nice. We understand her continuation, however, to be based on a broader swath of her beliefs, even though this

is never mentioned. Comparing (208a) to (208b), however, we see just how difficult it is to exclude a speaker's larger set of beliefs. In context (where Jessa knows nothing about Dylan besides what she's just read), the bolded portion of (208b) should unambiguously be based on the dating profile, but it cannot be understood as such. In other words, we assume that the totality of an interlocutor's speaker beliefs are 'active' by default, and overriding this default requires great effort—in the above case, the verb *sound*, used to express periphrastic evidentiality.

In many ways, the conclusions above are unsurprising. We assume our interlocutors will be consistent, which is a task greatly simplified if their background knowledge and assumptions are not in constant flux. This means that for all the ways that the evidential base of speaker beliefs is similar to other, more accepted evidential bases, the way it is invoked in discourse is quite different.

We can therefore conclude that a base E_{SB}, comprised solely of the speaker's private beliefs, should not be part of the inventory of bases that speakers can invoke. What, then, about the potential base of contextual evidence, E_{CE}? From the example below, it would seem that this base is similarly out-of-bounds:

(192′) *Sora hands Taji some right-handed scissors to bring to Sora's friend Ellie.*

T: Oh, Ellie can't use these.

S: She's left-handed? *✓S suspects Ellie is left-handed based on SB+CE*
 # S suspects Ellie is left-handed based on just SB

(209) *Sora and Taji are shopping for their mutual friend Ellie. They are looking at a display of scissors for sale. Sora vaguely recalls that Sora is left-handed, but thinks Taji, Ellie's best friend, likely knows better.*

S: Ellie's left-handed? *#S suspects Ellie is left handed based on SB+CE*
✓S suspects Ellie is left handed based on just SB

In both (192′) and (209), Sora's rising declarative indicates bias toward the truth of the claim that Ellie is left-handed, but the difference in contexts forces us to accommodate different reasons behind Sora's bias. In (192′), Taji's utterance provides contextual evidence in favor of Ellie's left-handedness. We therefore understand Sora's inference to be built on a) Taji's utterance, and b) Sora's general knowledge about scissors, such as the fact that most

people can use them, and that they come in two orientations. We cannot, however, understand Sora's question to follow merely from his preexisting beliefs, ignoring the contextual evidence. In (209), the opposite pattern holds. Because there is no contextual evidence regarding Ellie's handedness, we understand Sora's bias to be based solely on her (uncertain) beliefs about Ellie. Note also that the above examples demonstrate that the rising declarative is compatible with either positive speaker bias or positive contextual bias—a quality predicted if the weak base that rising declarative invoke does makes no distinction between the two, but not shared by the non-default polar questions discussed in this chapter.[7]

In sum, this evidence reveals that utterances not only *can* be supported by evidence stemming from speaker beliefs and contextual evidence, but that when such evidence is available, one's interlocutors assume it *will* be used. This requirement evokes Faller's (2012) discussion of evidential scalar implicature (as summarized in Chapter 3), whereby using a weaker evidential marker implicates the stronger ones are not available to the speaker. In this case, however, there is no selection for the speaker to make; by default a speaker's commitment is based on the totality of the evidence available to him, be it his private speaker beliefs or shared contextual evidence. These bodies of information, together, appear to underwrite all commitments, which is problematic for the idea of invoking E_{SB} and E_{CE} separately for modeling the effects of HNPQs.

Based on the above, it is tempting to completely collapse speaker belief and contextual evidence into a single default reasoning base which participants mutually assume their interlocutors will employ when making commitments. And yet, there is one crucial way that speaker belief and contextual evidence differ: By any appropriate characterization of a base E_{CE}, the content of that base would be mutually manifest and readily reconstructable by the speaker's interlocutors, while the exact content of a base E_{SB} would be usually unknown (and perhaps unknowable). It is therefore useful to continue considering the default base E_{DEFAULT} as a combination of private speaker beliefs and pubic contextual evidence, keeping in mind that over the course of a conversation, changes in the contextual evidence would be evident to all, while changes in speaker belief (if any) would not. The analysis of HNPQs presented in the next section exploits this very difference.

[7]Examples along the lines of (209) are mark as infelicitous by Gunlogson (2003), but subsequent literature (especially Poschmann 2008 and Gunlogson 2008) have argued convincingly otherwise.

4.3.3 The prior base

Unlike standard evidentiality, which involves introducing an evidential base of the appropriate evidence source, bias marking appears to be underwritten by speaker beliefs and contextual evidence, which together form the default base that conditions nearly all of the commitments a speaker makes. These two sources of evidence are difficult to divide, so invoking a base of speaker beliefs and a separate one of contextual evidence is problematic. How then, to proceed, given that these sources of information seem to be at odds in HNPQs?

The solution involves keeping in mind the crucial difference between the propositions that comprise a speaker's beliefs and those that are contextual evidence, namely that contextual evidence is public and shared. As the interlocutors engage in conversation, the bases they invoke for their commitments change to include the newly-shared contextual evidence; consider any normal discourse:

(210) *Ark looks up as Britta enters the room.*

A: I fear clowns.

B: That's understandable.

A: Yet you insist on tormenting me by dressing like that.

Ark makes two commitments in the dialogue above, one from each assertion. Neither is marked for evidentiality, dependence, weakness, or any other special discourse effect, so the default base of Ark's speaker beliefs and contextual evidence is employed in both. While the characterization is identical, though, the content of the base in each case is slightly different. This is because by the time of his second commitment, his default base contains the fact that Britta has committed to Ark's fears being understandable. Like the proverbial river that can never be stepped into twice, a speaker's default base changes as the conversation flows and is different every time he invokes it. How the base has changed, however, is apparent to everyone. This makes it straightforward for the speaker to step back in time and invoke an earlier version of his default base, or his *prior base*:

(211) PRIOR COMMITMENT

Let t be a time in the discourse prior to the current state. When a speaker A makes a prior commitment to p, he conditions his commitment on the base E_{PRIOR} s.t.:

i. $E_{\text{PRIOR}} = E_{\text{DEFAULT}}$ at t.

The characterization of the prior base does not specify what the prior time is. This is intentional, for reasons to be explored through the lens of HNPQs. Finally, note that the introduction of this base automatically makes prominent another quality of the default base (which it had all along). Just as we can talk of a default commitment as 'strong' in comparison to a weak commitment, it will be useful to think about default commitment as 'current' in comparison to a prior commitment, as shown below.

4.3.3.1 Formal analysis for HNPQs

As previously indicated, HNPQs share a semantic and pragmatic core with other polar questions. When a speaker utters a polar question questioning p (of any sort), she places the issue it raises, $\{p, \neg p\}$, on the Table. The immediate goal of discourse is then to settle this issue, which canonically occurs when the participants have all committed to either p or $\neg p$.

Bias marking involves the speaker making further, not-at-issue commitments that do not raise issues but nonetheless publicly commit the speaker to behave in a manner consistent with the commitments. Specifically, uttering a HNPQ questioning the anchor p has the two additional effects below:

(212) NOT-AT-ISSUE EFFECTS OF UTTERING A HNPQ QUESTIONING p:

i. Speaker makes a commitment to p, conditioned on $E_{\text{PRIOR+WEAK}}$

ii. Speaker makes a commitment to $\neg p$, conditioned on E_{WEAK}

Under this view, the asker of a HNPQ commits to both possible answers, but on different grounds. The commitment to p is based on an evidential base that is characterized as both prior and weak. This means that it differs from the default base in that any recent additions are absent, and that the base itself as the grounds for a commitment is highly mutable and requires immediate bolstering to remain solvent. The commitment to $\neg p$, on the other hand, is based on an evidential base that is characterized as weak, but otherwise derived from the default base like all evidential bases are.[8] In comparison to the commitment to p, it is similarly weak, but current. Note that the terms prior and current as used here apply

[8] I.e., the characterization of E_{WEAK} is equivalent to the characterization of $E_{\text{DEFAULT+WEAK}}$.

crucially to the provenance of the evidence, not the status of the commitment itself: A prior commitment does not somehow commit the speaker retroactively. Rather, the commitment is simply based on an un-updated body of evidence. In sum, (212) weakly commits the speaker to p given an earlier state of the context, and $\neg p$ given the current state.

This state of affairs triggers some pragmatic reasoning on the part of the addressee. The addressee can reason that if the speaker was willing to weakly commit to p based on outdated evidence, but to $\neg p$ based on current evidence, then the difference between the prior base and the default current one must make p less credible. This in itself is unremarkable; propositions become more or less credible all the time as evidence fluctuates. What makes the HNPQ case unique is that the speaker has chosen to alert the addressee to this change in his ability to commit. The addressee can therefore reason: Why does the speaker bother invoking an outdated body of evidence, unless he is not ready to give up on p?

4.3.3.2 Data revisited

I now return to some familiar examples below, followed by a series of new ones. Recall the generalization about the evidence requirements for felicitous use of HNPQs:

(204) EMPIRICAL GENERALIZATION FOR HNPQS, #1: **Bias conditions** (revised version)
HNPQs require a **conflict** between **positive** speaker belief and **negative** contextual evidence for the anchor proposition.

This generalization was based on examples (200–203), repeated in condensed form as (213) below, where Prof. Hemanti's utterance must supply contextual evidence that casts doubt on Ned's beliefs for him to felicitously utter his HNPQ. In this version of the example, the professor's utterances are marked with hashes and checks, rather than the question itself, to indicate that the entire dialogue is infelicitous.

(213) *Ned but not Prof. Hemanti knows that Jessup is habitually late to class. There's still time until it starts today.*

H: a. # We have such a busy day ahead!

b. ✓I want to start on time today.

c. ✓I hope we can start on time today.

d. # I wish we could start on time today.

N: Won't Jessup be late?

Under the proposal in (212), a context will support a HNPQ if the speaker can make the following commitments: **First**, he must be able to weakly commit to the anchor proposition p (here, that Jessup will be late) given a prior version of his default base. This requirement is clearly met above, no matter what the professor says; Ned comes to the discourse with knowledge of Jessup's habits (but not his actual actions), which is appropriate for a weak commitment given his base at that time. **Second**, he must be able to weakly commit to $\neg p$ given his default base, i.e., at the time of utterance. This depends on there being some change in his base between the start of the discourse and his utterance that bolsters his ability to commit to $\neg p$. Whenever Prof. Hemanti's utterance provides the contextual evidence that 'Prof. H believes starting on time is possible,' this condition is met and the question is licensed.

The same requirement explains why HNPQs are odd 'out of the blue' as in the repeated example below:

(205) *Jorge walks into the office.*

a. ✓Isn't Jim on campus today?

b. ?? Isn't Jim going to the supermarket today?

These examples are slightly different because they involve accommodation of contextual evidence. In (205a), it is easy to imagine what might have changed between the time of Jorge's prior base and his default current one to license the question; Jim was nowhere to be found. (205b) is licit only to the extent that we can guess at the change in contextual evidence.

The proposal offered here evokes Murray and Rett's (to appear) work on mirative evidentials in Cheyenne, where a *recency restriction* on felicitous use of such an evidential with propositional content p requires that the "learning event" of p and the "speaking event" follow in quick succession. From this it follows that p is *surprising* to the speaker. As mentioned previously, the issue of time itself is less crucial in the current account; there is no requirement that the time referenced by the prior base immediately precedes the current state. Rather, what's important is that all the interlocutors can identify *some* prior state,

such that more recent evidence has plausibly elevated the likelihood that $\neg p$. The idea that $\neg p$ is surprising is less crucial than in the mirative case.

As a final point in favor of the prior-evidence approach, note that HNPQs are extremely productive when used in lieu of accepting moves in conversation to indicate previously-unpublicized reservations:

(214) *Shelly and Harold are on a cross-country road trip. They recently stopped for lunch, and are now back on the road. Shortly, signs for a rest stop appear.*

S: We've only been going half an hour. Let's skip this one.

H: Actually, I have to use the restroom.

S: Oh, but didn't you go at lunch?

Given that there is no way that Shelly could hold on to her beliefs about Harold's bodily states in the face of what he tells her about them, the simplest response she could offer to his admission is a simple *oh*. Instead, she uses the HNPQ to indicate that she is unwilling to give up her convictions so easily.[9]

4.3.4 Novel data and solutions

While the analysis presented here fits well into the overall model of discourse proposed in this dissertation, the next major section will offer a comparison of the approach to existing ones. Before that, however, this section offers a few novel details about the uses of HNPQs which the present analysis highlights.

4.3.4.1 What counts as appropriate evidence?

The examples repeated below demonstrate an interesting contrast between an infelicitous HNPQ, and a felicitous one ending with *then*. In addition to serving as evidence against Gunlogson's (2003) sense of objective contextual evidence, the example reveals two important details of my analysis. First, it is possible to leverage the distinction between prior and current evidence to exclude cases where the speaker's ability to commit may have changed, but not in a way that is consistent with the bias configuration of a HNPQ. Second, it reveals

[9] Shelly's response is less good without *but* (*Oh. Didn't you go at lunch?*), and more combative if *oh* is omitted altogether.

that in some cases the difference between two states of CE can be a *lack* of expected change. Consider the contrast below:

(194) *Terra's house has a bathroom between the kitchen and the living room, and Cassie knows this.*

T: My house flooded. The kitchen and the living room are wrecked.

C: a. # Isn't the bathroom flooded?

b. Isn't the bathroom flooded then?

Recall that the oddness of (194a) arises from the fact that it seems to suggest that Cassie has some independent source of information concerning the state of the bathroom. (194b), on the other hand, is fully felicitous, but the inclusion of *then* makes it a (surface) question about Cassie's reasoning, rather than the state of the bathroom itself.

The badness of (194a) follows from the fact that HNPQs require a weak commitment to p given the prior base. Assuming that the current state of the context is after Terra's utterance and the prior state is before it (a reasonable default assumption), it is clear that Cassie's prior base did not license Cassie to commit to p (i.e., that the bathroom was flooded).

It is possible to make (194a) licit by changing the context so that Cassie did in fact think the bathroom was flooded. The relevant change between the prior and current states of her base, interestingly, is that Terra *didn't* mention the bathroom at all. In this way, it is possible for evidence against p to arise by omission.

This negative-by-omission evidence is even clearer in (194b). Here, the prior and default bases are almost inseparable; at the time the prior base hails from, Terra makes her utterance and Cassie can conclude that it is reasonable to infer that the bathroom is flooded. By the current time, however, Terra has said nothing about the bathroom (even though she did mention other rooms), which casts doubt on Cassie's inference. This example is illuminating because nowhere does positive speaker belief factor directly into licensing the HNPQ—in fact, Cassie has no positive speaker beliefs on the issue at all. Rather, Cassie contrasts a prior base that (in this case) depends on contextual evidence with her current default base that also depend on contextual evidence, but for the opposite proposition. Nevertheless, a contrast of the right sort exists, so the question is licensed, in accordance with the data.

This demonstrates an improvement over any analysis that attempts to pit private beliefs and public evidence against each other directly; both the new and old evidence, it seems, can be contextual. See the discussion of tag questions in the second half of this chapter for more evidence that prior and current evidence are the correct level of analysis, rather than explicit evidence type.

4.3.4.2 Actionability of evidence/degrees of felicity

Finally, recall that one of the clearest ways to investigate the discourse effects of an utterance is to look at the possible responses to that utterance. This was first demonstrated in the discussion of sourcehood in Chapter 2, where *yes* and *oh* responses could diagnose the source/dependent distinction.

In a similar vein, strong evidence for the analysis of HNPQs outlined here comes from considering the most felicitous responses to these sorts of questions. This is because the naturalness of a response reveals how the HNPQ forecasts the future of the discourse. In particular, it seems that for an affirmative answer to a HNPQ to be fully helpful, it must not only answer the at-issue question, but also address the aberrant change in evidence that the asker publicized:

(215) *Gertrude notices one fewer place settings than expected, given that Teddy was supposed to be visiting tonight, and says to Maxine:*

G: Isn't Teddy coming over?

M: a. ? Yes, (he is). *No target*

b. ✓Yes, but he won't get here until after dinner. *Targets CE*

c. # Yes, we were both there when he accepted the invitation. *Targets SB*

(215a), the bare or almost bare polarity particle response, is not entirely helpful. This is because Gertrude's HNPQ makes salient that her confidence that Teddy is coming over is weakened by the fact that there are too few place settings, and just answering her question does not assuage that conflict. (215b–c) attempt to address Gertrude's question more fully, but only (215b) is successful, because it specifically addresses the confusing contextual factors. (215c) also attempts to allay Gertrude's concern, but does so by bolstering her outdated evidence for p. This is therefore a less effective move than (215b), and comes

off as a snarky dismissal of Gertrude's conflict. In other words, neither (215a) or (215c) are blocked as possible responses to Gertrude's question because they do answer the question. They are not fully cooperative responses, however, because they do not address the particular factors that lead Gertrude to employ the HNPQ to publicize her conflict in bias.

In sum, my account is that by using a HNPQ, the speaker indicates that an anomaly in CE is affecting her ability to commit to one of the question's answers. Simply settling the at-issue content of the question ignores its broader effects, which suggests that modeling the bias conflict in a way that places the asker into a dilemma is an appropriate strategy for capturing the direness of the epistemic situation.

4.3.5 Interim conclusions

This section has presented a view of HNPQs that combines the semantics and core pragmatics of a default polar question with a pair of weak commitments to both answers. These commitments are made for conflicting 'reasons', with the positive commitment based on an old state of the default base, called the prior base, and the negative commitment based on its current state. This combination of effects requires the addressee to respond to the question normally, but also invites him to help deal with the discrepancy in the asker's commitments. This provides a second major generalization about the behavior of these questions, now accounted for:

(216) EMPIRICAL GENERALIZATION FOR HNPQS, #2: **Cooperative responses**
To be fully helpful, the response to a HNPQ must **address the aberrant evidence** that prompted the question.

The addressee can only do this if the offending change in the default base is recoverable, meaning that HNPQs will only be licensed when it is clear to the asker's interlocutors what evidence is responsible for casting doubt on her conviction in the positive answer. This leads to one final revision to the generalization about the bias conditions for HNPQs:

(217) EMPIRICAL GENERALIZATION FOR HNPQS, #1: **Bias conditions** (final version)
a. HNPQs require a **conflict** between **positive** speaker belief and **negative** contextual evidence for the anchor proposition.
b. The source of conflicting evidence must be **recoverable** by the addressee.

This is a behavior that any account of HNPQs will need to account for, and the present one has the advantage of tying the explanation back to more general facts about the negotiation of commitments.

Along the way, the character of the default base was investigated: It combines a discourse participant's private beliefs with the shared contextual evidence that evolves over the course of the conversation. Because changes to contextual evidence are by definition mutually manifest, I then claimed it was possible to make commitments based on an earlier state of the context. Note that this additional choice greatly increases the space of possibilities for what kinds of commitments a speaker could potentially make when using a polar question (or any non-evidential utterance, for that matter). In §4.5, I therefore show that introducing the prior base distinction is needed for than just HNPQs; it also plays a crucial role in differentiating between different kinds of tag questions.

Before this, however, the next section briefly examines existing analyses of HNPQs, which tend to center on a slightly different issue surrounding an ambiguity that the questions have been claimed to allow. While reviewing these analyses, it is worthwhile to keep in mind that the holy grail for analyses of HNPQs is successfully explaining why it is *negation*, and not some other marker, that conveys the form's rather complicated bias conditions. The most successful account to date in this regard is Romero and Han (2004), where negation is involved in a scope ambiguity and straddles the boundary between the sentential realm and the illocutionary one.

Among accounts that do not involve a scope ambiguity, the general approach to this question follows van Rooy and Safarova's (2003) utility-theoretic argument, where use of negation signals that the negative answer is less expected and hence more informative. The account presented here tells a version of this story; the negative answer is technically *more* expected in terms of subjectively likelihood, because it is based on more current evidence, but the speaker's unwillingness to adopt that new evidence (as signaled by a commitment to the positive answer based on old evidence) underscores that the negative answer is less desired. The speaker's reservations revolve around whether she must really adopt the negative answer. To do so would require giving up her skepticism, which makes the negative answer more informative. More on the issue of why the HNPQ conveys the bias that it does appears in §4.7, which explores the typology of polar questions.

4.4 Aside: A related ambiguity in HNPQs

The discussion so far has focused on the discourse effects of HNPQs that separate them from 'default' unbiased polar questions, but it is important to note that HNPQs have presented a long-standing syntactic, semantic, and pragmatic puzzle even without direct comparison to other kinds of polar questions; as noticed by (Ladd, 1981), out of context, HNPQs seem to be ambiguous between 'biased for p' and 'biased for $\neg p$' readings:

(218) Outer negation reading ('biased for p')

A: You guys must be starving. You want to go get something to eat?

B: Yeah, **isn't there a vegetarian restaurant around here?** Moosewood or something like that?

(219) Inner negation reading ('biased for $\neg p$')

A: I'd like to take you guys out to dinner while I'm here.

B: But there's not really any place to go in Hyde Park.

A: Oh, really, **isn't there a vegetarian restaurant around here?**

Ladd's terms 'outer–' and 'inner negation' stem from the intuition that in the cases like (219), negation is, in the words of Horn (1989), "of the clause but not in it", occupying a higher position in the sentence's structure.[10] This difference in the relationship of negation to the rest of the clause between the two readings can be drawn out with the use of negative and positive polarity items (NPIs;PPIs), as shown below.

(220) a. Those options sound good, but **isn't there a vegetarian restaurant around here, too?**

b. No Thai, no Indian. . . **isn't there a vegetarian restaurant around here either?**

The fact that the presence of the NPI *either* versus the PPI *too* disambiguates the readings suggests an actual structural difference between the sentences, under the view that NPIs are licensed by sentential negation. As will be shown, the existing approaches to explaining the ambiguity exploit this to varying degrees of success.

[10]Horn's (1989) quote here refers to the general phenomenon of metalinguistic negation, but as Reese (2007) notes, Ladd's intuition about negation in these examples is closely related. This is consistent with the idea that the 'anchor' answer for these questions is p, not its negation.

With some notable exceptions (e.g., Büring and Gunlogson 2000; Romero and Han 2004), these analyses abstract away from the complicated facts about conflicting bias explored in §4.3.1, but regardless of whether the HNPQ is used to convey its 'biased-for-p' or 'biased-for-$\neg p$' reading, the conflict in bias argued for in the previous section is present. What differs between the two readings is which bias the speaker seems to expect should win out. This has the effect of suppressing the 'weaker' bias. Given the felicity conditions for these questions as presented here, however, ignoring the fact that two different sources of bias must be at odds and focusing on just the stronger source cannot lead to a satisfactory analysis. Note that in terms of Ladd's (1981) ambiguity, the analysis presented so far focuses on the outer-negation reading only. Whether this reading is in fact the only reading or just the primary one is explored at the end of this section.

4.4.1 Previous approaches to high negation

There have been many attempts to derive the above behavior of HNPQs, a non-exhaustive sample of which are sketched below. The approaches reviewed here have been selected either because they are widely known, or because the accounts present illuminating contrasts with the account I argue for. See Reese (2007) and Krifka (2012) for more detailed reviews of the literature.

4.4.1.1 Verum focus (Romero and Han, 2004; Anderbois, 2011)

Arguably the most influential view of HNPQs is due to Romero and Han (2004), who argue that these questions exhibit *verum focus*, which Höhle (1992) describes as "emphasis" on determining the correct truth value for an expression. They instantiate verum focus by means of an epistemic operator VERUM that takes a propositional argument p and returns true just in case the interlocutors share a high degree of certainty that p belongs in the common ground. Focusing on the truth value has the effect of changing the issue that the question raises. Under this view, HNPQs exhibit a syntactic scope ambiguity between negation and VERUM. The questions expressed by each reading of the HNPQ can thus be schematized as in (221) below.

(221) a. $\neg\text{VERUM}(p)$ = 'It is not for sure that p should be added to the common ground'
Outer-HNPQ

b. VERUM($\neg p$) = 'It is for sure that $\neg p$ should be added to the common ground.'
Inner-HNPQ

Where the account offered in §4.3 captures the expression of bias as a unique form of weakened commitment, Romero and Han (2004) instead make bias a pragmatic inference that results from the way that the speaker frames the issue that he raises. Ladd's bias contrast, in this view, arises via implicature under the assumption that speakers move to confirm what they already are predisposed to believe; for example, when a speaker asks explicitly about the common ground status of p with the outer reading of the HNPQ, it implicates bias for p.

Note that the NPI facts follow directly, under the assumption that when negation scopes over VERUM, it is no longer able to interact with any propositional content and so cannot license NPIs. This further predicts the acceptability of PPIs with outer-HNPQs, given that the proposition is itself positive.

Romero (2005), Reese (2007), and Anderbois (2011) note however that by changing the issue that the HNPQ raises, Romero and Han's (2004) analysis makes some incorrect predictions about the distribution of these questions. For example, the analysis exhibits what I will call the *response-to-certainty problem*; if the questions alternatives really target the certainty of p, Marline's answer in (222) is predicted to be licit:

(222) RESPONSE-TO-CERTAINTY PROBLEM

G: Isn't Teddy coming over for dinner?
≈ 'Is it not for sure that Teddy is coming over for dinner?'

M: # No, but he probably is

Further, as Romero and Han (2004) themselves note, the analysis stipulates the presence of VERUM, rather than tying it in a meaningful way to the fact that the only overt marking of the HNPQ involves negation. In other words, the fact that it is negation, and not some other morpheme, that differentiates this question form is not derived from the analysis.

In response to these facts, Anderbois (2011) offers an alternative verum-based analysis, starting from a different view of what it means to 'emphasize' the question of a proposition's truth value using verum focus. Anderbois's (2011) analysis employs a rich, two-tiered

Inquisitive Semantics whereby the discourse effects of utterances can specify not only the main issue that the utterance raises, but also *sub-issues*. These sub-issues indicate further expectations for the evolution of the discourse, and can be introduced by a number of operators. For example, the indefinite *a cake* in (223) below raises the sub-issue of which cake (or rather, what kind of cake) John is baking:

(223) A: Is John baking a cake?
B: Yes, chocolate.

B's answer is more cooperative than a simple *yes* precisely because it addresses this sub-issue. In contrast, Anderbois (2011) claims that B's answer is not particularly good after an otherwise-parallel HNPQ from A:

(224) B: I want to get some fruit or something for after dinner.
A: Isn't John baking a cake?
B: # Yes, chocolate.

Anderbois (2011) suggests that this is because the effect of the VERUM operator is to suppress any sub-issues that would otherwise pop up in its scope—thereby focusing in on the main issue, the truth of p (or $\neg p$). This approach avoids the response-to-certainty problem exemplified in (222) that dooms Romero and Han (2004), and also offers a satisfactory explanation of why the discourse effects of a HNPQ are marked with negation: the suppression of sub-issues is effected in Inquisitive Semantics via double negation.

Under either formulation, the specific bias conditions for the felicitous use of the HNPQ are treated as a separate issue. Citing Büring and Gunlogson (2000), these authors note that HNPQs with either reading require prior speaker belief, and in at least some cases, negative contextual evidence. When these conditions are met, the VERUM operator is licensed, and the bias is communicated via implicature.

The infelicity of the reply in (224) deserves further discussion, however, in light of the generalization #2 from §4.3.5:

(216) EMPIRICAL GENERALIZATION FOR HNPQS, #2: **Cooperative responses**
To be fully helpful, the response to a HNPQ must **address the aberrant evidence** that prompted the question.

The infelicity of (224) fits with this generalization under the assumption that the fact that the cake is chocolate is insufficient to resolve the asker's confusion. Note, however, how easily the response is made licit:

(224′) B: I want to get some fruit or something for after dinner.

A: Isn't John baking a cake?

B: ✓Yes, he's making a chocolate cake and I'm allergic

Similarly, if the original (224) is uttered in a context where A and B both know that A is aware of B's chocolate allergy, the response is felicitous. In either case, all that is required is that the response address the offending evidence to make the continuation licit.

Even allowing for (224′), what is problematic about the verum-based accounts discussed here is that it is not clear why verum focus should be sensitive only to the specific conflict in evidence explored in previously:

(217) EMPIRICAL GENERALIZATION FOR HNPQs, #1: **Bias conditions** (final version)

a. HNPQs require a **conflict** between **positive** speaker belief and **negative** contextual evidence for the anchor proposition.

b. The source of conflicting evidence must be **recoverable** by the addressee.

Whether VERUM is invoked to manipulate the main issue a question raises or suppress subissues, the resulting 'emphasis' on discerning truth is compatible with a range of pragmatic explanations. It is therefore unclear why only the contrast between positive prior evidence and non-positive current evidence is ever implicated. Consider (224″):

(224″) B: I want to get some fruit or something for after dinner.

A: Isn't John baking a cake?

B: # Yes, he was very adamant about when he told us earlier.

Under either verum focus analysis, it is unclear why this response should be problematic, given that it directly bolsters the idea that the issue should be settled in the positive direction.

Finally, note that the same basic pattern (and hence, problem) appears with negative answers to HNPQs as well:

(225) B: I want to get some fruit or something for after dinner.

A: Isn't John baking a cake?

B: i. # No, (he isn't.)

ii. No, he was just saying that to sound good in front of Maria.

iii. # No, he's definitely not. He can't even microwave popcorn.

The verum-based accounts predict that the simple answer in (225a) should be sufficient, given that it focuses on the truth value of the response and avoids sub-issues. (225c) should similarly be licit, given that it addressees the certainty of the answer. Instead, only (225b), where the additional information deals with the conflict A is experiencing, is licit.

In a sense, the above could be rephrased in terms of these verum focus analyses as follows: I have argued is that rather than either suppressing or changing the issues a question would otherwise raise, using a HNPQ allows a speaker to promote a new 'sub-issue', centered on the conflict in evidence he is experiencing. The result, similar to Anderbois's (2011), is that other sub-issues are suppressed, but the suppression derives from the way that the (independently-needed) evidence conditions directly affect the discourse structure. While the analysis does not yet derive Ladd's ambiguity, it offers a more transparent explanation of the interaction between evidence conditions and the limits a HNPQ imposes on the future of the discourse.

4.4.1.2 VERUM and FALSUM (Repp, 2012)

Repp (2012) offers an analysis that builds on Romero and Han by adding a second "common ground managing" operator called FALSUM. FALSUM, as its name implies, takes a propositional argument *p* and returns true just in case the degree of certainty that *p* should be in the common ground is zero. Under this view, it is FALSUM that outer negation contributes in the Outer-NPQ case, so that B's utterance in (198b) can be understood to ask whether 'it is the case that there is zero evidence that *p* should be added to the common ground'.

This view offers a solution for Romero and Han's (2004) response-to-certainty problem (in (222)), while maintaining the view that the main issue raised by the HNPQ is what is altered by verum focus. Given than the analysis still employs verum focus formally and leaves bias to implicature, however, the criticism raised in the previous subsection still

applies; the main data generalizations, about the conflict in bias and how it plays prominent in felicitous answers, has no explanation beyond what Romero and Han (2004) can offer. Also, FALSUM fits awkwardly in cases where the outer-HNPQ is used in a context that is neutral:

(198c) *A and B are crafting a list of left-handed people to invite to an unusual party.*

A: So there's Ellie, Reuben. . . who else? *p-neutral*

B: Isn't Milo left-handed too?

Here, 'is it the case that there is zero evidence that Milo is left-handed?' seems like a much weaker sentiment than B is actually expressing. Consider, for example, that B cannot felicitously follow her utterance with an admission of ignorance:

(198c)′ *A and B are crafting a list of left-handed people to invite to an unusual party.*

A: So there's Ellie, Reuben. . . who else? *p-neutral*

B: Isn't Milo left-handed too? # I have no idea.

If B's question merely asks about whether there is a lack of evidence, it is not clear why she is not licensed to indicate that she in fact lacks such evidence herself.

Finally, unlike Anderbois's (2011) analysis, which provides an explanation for why the effects of HNPQs is tied to negation, Repp's (2012) analysis further complicates the issue; negation must now sometimes introduce VERUM, and other times FALSUM.

4.4.1.3 Pragmatic utility (van Rooy and Safarova, 2003; Krifka, 2012)

van Rooy and Safarova's (2003) analysis is unique in straying from Ladd's (1981) intuition that the negation in outer-HNPQs is outside of the clause and instead treats it as normal propositional negation in both cases. Their account instead relies on the notion that bias in HNPQs is a matter of "pragmatic utility," whereby a speaker is expected to choose to form his polar question so that the affirmative answer is the less expected option. In other words, speakers set up their polar questions to check if they are *incorrect*. This argument is predicated on the claim that the less expected an answer is, the more useful it is to the speaker, and that in general, an affirmative answer is less marked than a negative one.

The general schema of this view resonates with generalizations about the conflict in evidence (217) and licit responses (216), but as Krifka (2012) and others note, the specific implementation is inadequate with regard to the issue of NPI licensing; if both readings of a HNPQ involve simple propositional negation, there is no explanation for the distribution of NPIs and PPIs across the readings.

In terms of the bias conflict central to the analysis advanced in this chapter, pragmatic utility therefore meets the same problem as verum focus; there is no way to limit the use of a HNPQ to those cases where the reasons behind the speaker's bias involve aberrant contextual evidence.

Krifka (2012) makes use of pragmatic utility as well as scope ambiguity, and so is something of a portmanteau of earlier approaches. Similar to Romero and Han (2004), the approach allows negation to act above the level of the proposition. Unlike Romero and Han (2004), however, negation is allowed to reach all the way to the level of the speech act, where it interacts with a new realm of 'meta-speech act' operators. In addition to a fairly standard ASS(ERT) operator that takes a propositional discourse referent φ and returns an assertive speech act corresponding to a context update with content φ, the account introduces a higher $\text{REQUEST}_{S1,S2}$ operator that registers a request from discourse participant S1 that another participant S2 perform a given speech act.

The REQUEST operator is notable for two qualities. First, it adds a new layer of abstraction to the discourse structure, because it formalizes the process of eliciting a response, rather than leaving it to higher pragmatic reasoning about turn-taking—the more usual assumption. Second, the fact that the operator can scope under negation is extremely powerful; participants can not only request that their interlocutors make certain future moves, but that they *refrain* from making certain moves. Krifka (2012) terms this *speech act denegation*, after Searle (1969).

Under this view, B's utterance of the outer-HNPQ "Isn't Milo left-handed too?" registers B's request that A perform the denegation of asserting "Milo is left-handed." (i.e., $\text{REQUEST}_{S1,S2}(\sim\text{ASS}(p))$).[11] At this point, pragmatic utility must come to bear. In uttering the outer-HNPQ, B has requested that A refrain from asserting that Milo is left-handed as a way of expressing that B in fact expects that Milo *is* left-handed. This follows only if B is

[11]The author notes that his approach is similar to Reese and Asher (2009) in that it allows negation over speech acts in this way. See below.

framing his question in such a way that the positive answer would be more surprising for him. Note that again, the account cannot transparently predict that the reason behind the speaker's denegation request involves a conflict in evidence.

It is possible to construe this account as successfully capturing at least some portion the response generalization for HNPQs, repeated below:

(216) EMPIRICAL GENERALIZATION FOR HNPQS, #2: **Cooperative responses**
To be fully helpful, the response to a HNPQ must **address the aberrant evidence** that prompted the question.

Under this view, when a speaker utters a HNPQ with the anchor p, he requests that his addressee refrain from asserting p. Given (224′):

(224′) B: I want to get some fruit or something for after dinner.
A: Isn't John baking a cake?
B: ✓Yes, he's making a chocolate cake and I'm allergic

This would mean that A's question requests that B not assert that John is baking a cake. When B *does* go on to assert it, he must therefore offer an explanation for why he has gone against A's request. The downside to this account, however, is that the reader will note that even though B has rejected A's request in the example above, the response is marked with *yes*, which is generally taken to be a marker of acceptance. In this case, Krifka (2012) claims that *yes* instead refers directly to the positive proposition, rather than the asker's question.

No negative answers to questions, the results are less clear:

(225) B: I want to get some fruit or something for after dinner.
A: Isn't John baking a cake?
B: i. # No, (he isn't.)
ii. No, he was just saying that to sound good in front of Maria.
iii. # No, he's definitely not. He can't even microwave popcorn.

It seems that if the B is going along with A's request not to assert that John is baking a cake (p), further discussion should not be needed. At the very least, it is unclear there should

be a problem with response (225c), which shows how adamantly B opposes the idea of asserting p.

Finally, because this solution employs a syntactic ambiguity, the NPI facts are again straightforwardly accounted for; the negation of speech act denegation is quite literally metalinguistic, and so does not interact with propositional content, which means NPIs are unlicensed in outer-HNPQs, while the normal sentential negation in inner-HNPQs does license them.

4.4.1.4 Conditional assertion (Reese and Asher, 2009)

Similar to the counterarguments offered above, Reese and Asher (2009) argue against verum focus and pragmatic utility accounts based on the premise that because the range of acceptable contexts for these questions is limited, "conversational goals—" i.e., discourse moves that explicitly move the discourse in a particular direction—are not the right level of description to capture the effects of HNPQs.

Further, Reese and Asher (2009) move away from the assumption that HNPQs are at their core identical to PPQs and other polar questions. They instead treat Outer-HNPQs as complex speech acts, built out of an assertion and a question, with a pragmatic "glue logic" that composes the acts. The final speech act is ASSERTION • QUESTION, within the framework of Segmented Discourse Representation Theory (SDRT; Asher and Lascarides 2003). The arguments for this hybrid fact derive from the idea that, in my terms, HNPQs show a certain level of commitment on the part of the asker. Bias is therefore captured as an outright assertion, counter to the earlier approaches where bias is an implicature calculated off a non-default question act.

While the spirit of this approach is of a kind with the one advanced in this chapter, there are at least two major weaknesses which limit its usefulness. First, it is problematic to treat bias as 'derived' from assertion, because it muddies the concept of a strong commitment. If an assertion followed immediately by a question results in an expression of bias, it is not immediately clear why the utterances below are not equivalent:

(226) a. The Major's returned from her mission?

b. Hasn't the Major returned from her mission?

c. The Major has returned from her mission, hasn't she?

d. The Major has returned from her mission. # Has(n't) the Major returned?

The (226d) utterance is clearly self-contradictory in a way that (226a–c) are not. This is because the commitments expressed by (226a–c) are mutable, unlike (226d)'s. If the situations in which the glue logic can combine an assertion and a question are limited, more needs to be said about why that is the case. Also, note that while a tag question like (226c) might look easily derivable from a combination of two sentences, the forms in (226a–b) are more transparently declarative and interrogative, respectively.

A second concern is that while the approach illuminates a number of insightful truths about HNPQs, its formal explanatory power is low because the inventory of discourse relations in SDRT is both fine-grained and high level. The relations used to capture the effects of polar questions include *Correction*, *Counterevidence*, *Acknowledgement*, or *Confirmation* depending on the context for a specific question. This is a dubious improvement over the analyses couched in terms of conversational goals. The analysis advanced in this chapter does without the granularity of SDRT moves.

A final concern is that negation is a fairly minor marker on which to pin the idea that an utterance that employs the HNPQ form is a hybrid of two speech acts.

The advantage of the approach, on the other hand, is its empirical coverage. The account clearly predicts that HNPQs requires a contrast between prior bias (which licenses the assertion) and current bias (which licenses the question). Given this, the account also predicts the felicity of responses that address this conflict. That said, a potential failing lies in the infelicity of bolstering answers:

(224″) B: I want to get some fruit or something for after dinner.

A: Isn't John baking a cake?

B: # Yes, he was very adamant about when he told us earlier.

If the purpose of B's response here is to help A resolve the fact that he wishes to retain the result of his assertion, then B's answering the question in a way that strongly settles the question-half of A's hybrid utterance is incorrectly predicted to be felicitous.

4.4.1.5 The universal problem

With the notable exception of Reese and Asher (2009), what these accounts have in common is that the bias conflict that I have claimed is the defining feature of (outer-)HNPQs is derived via implicature, based on the way that the speaker presents expectations for the future of the discourse. Reese and Asher (2009) give primacy to the expression of bias, but go too far and lose explanatory power.

In each case, one of the primary goals of the analysis is to derive the NPI licensing facts, which most of the accounts manage successfully. It is crucial to note that the analysis advanced here does *not* derive these facts. The natural move at this juncture would therefore seem to be extending my analysis to handle the inner/outer distinction, but I will not pursue this for the reasons outlined below.

4.4.2 A missing reading

Ladd's (1981) inner/outer ambiguity is not universally accepted in the literature (van Rooy, 2003; Anderbois, 2011; Roelofsen et al., 2012). Three objections are given below. First, a surprising proportion of speakers outright reject the ambiguity and the evidence that supports it. This points toward the possibility that for many speakers, the outer (biased for p) reading is in fact the only one.

Sailor (2013) finds that the inner (biased for $\neg p$) reading licenses NPIs poorly; as shown below, *either* is the only NPI that seems to be licensed with any reliably, and even then, only for speakers of Canadian and British English. For example, (227), one of the test contexts from Sailor (2013), presents a situation where Mary's prior beliefs about the set of foods that her father-in-law is allergic to come into conflict with overwhelming contextual evidence that her beliefs are wrong. Nonetheless, the inner-HNPQ, marked with *either*, is not licit for most speakers:

(227) *Mary's in-laws are coming over for dinner, and she wanted to make dessert for them. The recipe calls for walnuts, though, and her father-in-law is allergic to walnuts, so Mary planned to use peanuts instead. However, her husband said, "I don't think that's a good idea".*

M: a. % Can't your father eat peanuts either?

b. Can your father not eat peanuts either?

An acceptable way to express something close to the intended meaning is with the low negation polar question in (227b), although see Anderbois (2011) for arguments that the low negation polar question and the inner reading of the HNPQ are distinct. When presented with the second option, speakers uniformly prefer it to the inner-reading of the HNPQ.

Other NPIs, like *until* below, fail to disambiguate the HNPQ as Ladd claims. Rather, they result in an ungrammatical question. Note that the example below is Sailor's (2013), although the context is mine.

(228) *Barnaby and Zoë are discussing a party Zoë attended last weekend, which Barnaby had to miss at the last minute. Barnaby knows that Cristian had planned to leave the party early, so as to avoid running into his boss, Sarah.*

Z: Christian wound up leaving through the bathroom window to avoid Sarah.

B: a. * Didn't Christian leave until Sarah arrived?

b. Did Christian not leave until Sarah arrived?

c. Wasn't Christian going to leave before Sara arrived?

In this context, Barnaby is presented with compelling contextual evidence that speaks against his prior beliefs, but the attempted inner-HNPQ in (228a) is unacceptable as a way to express the question in (228b). Note that an appropriate outer-HNPQ (228c) is licensed perfectly here. Because the presence of the NPI in (228a) supposedly disambiguates the HNPQ, it is unclear under any of the accounts presented so far why the question is so unacceptable. More generally, this pattern is problematic for any view of HNPQs that treats the negation in inner-HNPQs as normal sentential negation, given that the NPI is clearly not licensed here.

Perhaps more damning evidence comes from experimental work by Hartung (2007), which shows that speakers only rarely accept the inner reading of a HNPQ. The study asked participants to judge the naturalness of HNPQs with *either* or *too*, and compared these judgements to those of Low Negation Polar Questions (LNQPs). The results showed that inner-HNPQs were judged extremely unnatural relative to their LNPQ counterparts, suggesting that only under duress is the inner reading even available.

These results suggest that Ladd's (1981) ambiguity may have less to do with the nature of HNPQs and more to do with the specific NPI *either*. This is the position taken by Anderbois (2011), for whom HNPQs by default express a desire to retain a prior belief:

(229) ANDERBOIS'S (2011) DEFAULT BIAS PRINCIPLE FOR HNPQS
A speaker who utters a HNPQ conveying a prior belief is taken by default to have a (strong) bias towards retaining that belief.

The inclusion of an NPI like *either* can, for some speakers, override this default bias configuration in order to express that the speaker expects to divest his prior belief. This means that the inner-HNPQ reading is only available when an NPI is present, counter to both Ladd (1981) and Romero and Han (2004). I happily adopt a version of this solution for the analysis advanced here, as well. For some speakers, the inclusion of (certain) NPIs in a HNPQ has the effect of indicating that the speaker wishes to merely register his prior bias for the positive answer, rather than retain it. This derives from the factors that govern felicitous use of NPIs, however, rather than a special discourse effects that should be attributed to the HNPQ, which I maintain involves the publicization of discourse-evidential factors as described in the previous section.

4.5 Bias in tag questions

In the previous section, I showed that the bias expressed by HNPQ is complex because it derives from conflicting evidential sources: the speaker expresses a willingness to commit to p given prior evidence built from preexisting beliefs and previous contextual evidence, held back by a reluctance due to recent evidence that argues against p. This conflict is formalized as a pair of weak commitments, one to p given a previous state of the speaker's default base, and the other to $\neg p$ given its current state. These commitments, plus pragmatic reasoning about why a speaker would publicize commitment based on an outdated base, convey the overall message that the speaker is unwilling to readily give up her pro-p position.

Tag questions (TQs) provide an interesting extension to this picture because they show sensitivity to the same kinds of evidential bases, but are used to express bias in different kinds of contexts. Below, I introduce basic data about the kinds of TQs I will be considering, both in terms of their form and their discourse effects. The result shows a picture remark-

ably similar to HNPQs, but with some distinguishing features, chief among which are less stringent bias conditions. While the bias of particular kinds of tag questions derives from certain configurations of evidential bases, these configurations do not involve a conflict of the sort displayed by HNPQs. Rather, different tag questions split the work of marking just a single kind of bias each. The goal of this section is to provide a formal analysis of TQs that both captures their properties and highlights how they differ from HNPQs.

4.5.1 A brief typology of tag questions

A *tag question* is defined by a combination of its syntax, semantics, and pragmatics. Each of these is considered in turn below.

4.5.1.1 Invariant vs variant tags

A tag question is encoded in the form of a *tag interrogative*, which has two parts: A declarative *anchor*, followed by a short *tag*. While the anchor can vary in any of the ways a declarative sentence normally can, the space of variation for the tag is more circumscribed. At the most obvious level of tag variation, tag questions break into those featuring *invariant tags*, as in (230a–b), and *variant tags*, as in (230c–e) below. A final question mark indicates rising intonation, while a period indicates falling intonation.

(230) a. Milo owns a VCR, right?

b. Milo owns a VCR, huh.

c. Milo owns a VCR, doesn't he?

d. Milo owns a VCR, doesn't he.

e. Milo owns a VCR, does he?

Kimps (2007) notes that the ubiquity of variant tag questions in English is at odds with a broader cross-linguistic trend toward the use of invariant tags. This is perhaps because the dividing line between what should be categorized as a tag interrogative and what is actually just a declarative sentence with a sentence-final particle is empirically fuzzy; recall from Chapter 3 that the most natural translation for many declaratives that occur with the final particle *ne* in Japanese involves a tag question:

(97) *The meteor approaches.*

S: Owari da ne.
end COP NE
'(This is) the end, isn't it.'

This is not to say that *ne* in this example and the tag question in its translation contribute the same effect—I will argue that they do not—but the fact that there are empirical commonalities between tags and discourse particles cross-linguistically constitutes yet another fact that accounts of these phenomena should be able to explain. In the case of the present analysis, the common denominator is the structure of discourse commitments and the way that non-default discourse effects are encoded in evidential bases.

The example above highlights an important difference between Japanese *ne* and the tags, as well; namely, the Japanese sentence was analyzed as a modified assertion, while the proposal advanced here is that TQs are questions. The fact that the form of one speech act is translated as another speaks to the difficulties in establishing the difference between a non-default assertion with weakened commitment on one hand, and a non-default polar question with strengthened commitment on the other. This is because diverging from either a default assertion (with full commitment) or a default polar question (with none) results in an expression of weak commitment, i.e., bias (Farkas and Roelofsen, 2013).

In order to provide a measure of clarity, the discussion offered here will therefore be limited to the class of tag questions that most clearly involve an interrogative element: variant tag questions. The approach here should be amenable to expansion for the invariant tags as well, but more empirical work on their distribution is needed beforehand.

As discussed by Sailor (2009), variant tags are always reduced interrogative clauses that 'match' the anchors they occur with. In order to match, the tag must minimally include a pronominal subject coreferent with the anchor subject and the same tense and mood as the anchor clause:

(231) a. Milo owns a VCR, doesn't he?
b. # Milo owns a VCR, doesn't Milo?
c. # Milo owns a VCR, doesn't she?
d. # Milo owns a VCR, shouldn't he?
e. # Milo owns a VCR, mightn't he?

f. # Milo owns a VCR, doesn't he?

Note also that including more of the anchor's lexical content, such as the matrix verb, leads to an infelicitous tag question, as shown in (232) below.

(232) a. Milo owns a VCR, doesn't he?

b. # Milo owns a VCR, doesn't Milo own a VCR?

c. # Milo owns a VCR, doesn't he own a VCR?

While (232b–c) are felicitous utterances, their prosody must be that of two distinct phrases. The resulting discourse effect seems to be different as well; the fuller tag-like phrases seem to cancel the assertion in a way that is not the case for the true tag question, which do not give the impression that the speaker is contradicting himself.

(233) EMPIRICAL GENERALIZATION FOR TQS, #1: **Speaker consistency**
(Variant) tag questions are not understood as self-contradictory, even though their forms suggest they would be.

The above observation suggest that tag questions do not involve the speaker making two, contradictory speech acts. This is in fact one of the arguments against a view of tag questions in the style of Reese (2007) that involves chaining the semantics of a question after the semantics of a declarative; it's unclear why the unreduced form should pattern differently from the reduced one.

4.5.1.2 One clause or two?

Given the discussion above, a major theme in the study of tag questions centers on discovering how tags are formed, and by extension, how closely related the two clauses are. Some prominent analysis of tag questions, notable Reese and Asher (2009), treat tag questions as a simple combination of a declarative sentence and an interrogative one, encoding either both an assertion and a question, or a hybrid speech act that is explicitly derived by combination of the two.[12]

[12]Reese and Asher's (2009) analysis deals specifically with reverse variant tag questions, as introduced below. They employ a similar analysis for HNPQs, as mentioned above.

Tags are generally taken to be the remnants of an ellipsis process, which is sometimes used as evidence for this view. See Sailor (2009) for details, but also King (2011), which nuances the view.

4.5.1.3 Intonational contour on the tag

The form of tag questions vary systematically in at least four other ways, two of which are polarity-based and the other two of which are intonation-based. These distinctions are introduced below, with an eye to what each dimension could contribute to the meaning of the tag question compositionally.

One of the most salient ways that tag questions can differ is that the final intonational contour can either be falling or rising, as represented below with a period and a question mark, respectively.

(234) a. *Gertrude notices an extra place setting and says:*
Teddy's coming to dinner, isn't he.

b. *Gertrude notices one fewer place settings than expected, given that Teddy was supposed to be visiting tonight and says:*
Teddy's coming to dinner, isn't he?

Note the difference in contexts, which helps disambiguate the utterances; switching the intonation in either case results in an aberrant utterance. The way that the context leading up to a tag question affects its felicity evokes the previous discussion of HNPQs, and indeed the bulk of this section will be devoted to capturing the rising versus falling tag distinction. While I will eventually arrive at an analysis that I believe captures the difference between these kinds of questions in an enlightening way, it is worth noting that some researchers question the basis for such an approach. In particular, Reese (2007) argues the effects of variation in intonation is too multifaceted across sentence types and speech acts to allow for a unified characterization of, e.g., the final rising tune. The present discussion does not diverge from Reese (2007) on this point, but nevertheless I believe that if we limit our attention to the difference in final intonation specifically within the context of tag questions, it is possible to arrive at a useful characterization, albeit one that cannot be decomposed further into distinct meanings for tag question and the intonational contour it

carries. Rather, the form of a tag question is paired with intonation idiomatically.[13]

Caveats aside, there does seem to be an intuitive difference in the level of commitment that rising and falling tag questions express, which must somehow be attributable to the intonation itself. The difference is slight, but here I briefly introduce a new test for commitment in order to illustrate it.

Consider the examples below, where A's utterance is followed by a response from B that attempts to question the reasons behind A's utterance.[14] Note that in all cases, *that* refers to the proposition that Jack knows where the body is.

(235) A NEW TEST FOR COMMITMENT

a. A: Jack knows where the body is. *Decl*

   B: Why do you think that (Jack knows where the body is)?

b. A: Does Jack know where the body is? *PPQ*

   B: #Why do you think that (Jack knows where the body is)?

c. A: If Jack knows where the body is, we should ask him about that. *Conditional*

   B: #Why do you think that (Jack knows where the body is)?

In these examples, B's utterance presupposes that A holds a particular belief about Jack. The response is felicitous after an assertion (235a), but not after an information-seeking question (235b) or if the relevant proposition occurs in environments like the antecedent of a conditional (235c). Given the difference between asserting and questioning as discussed in §4.2, these examples therefore reveal that the phrase "Why do you think that?" is only acceptable in response to utterances that entail an actual commitment on the part of the response's addressee to the proposition that serves as the antecedent of *that*.

The judgements for (235a–b) above are backed up by an online poll of 110 self-selected respondents. 86% of respondents accepted B's response to the assertion in (235a) as natural, while only 6% found it natural after the polar question in (235b).

Altering the test slightly by replacing *do* with *would*, a different pattern emerges:

[13]The reader may note that took a different approach in the previous chapter when discussing the relationship between the Japanese sentence-final particle *yo* and intonation. This is because unlike the intonation on tag questions, the intonation on *yo* is clearly distinct from normal phrasal intonation.

[14]More specifically, it targets the commitments that stem from the utterance. A similar test with *say* would target the appropriateness of the actual speech act.

(236) A NEW TEST FOR COMMITMENT, PART 2

a. A: Jack knows where the body is! *Decl*

B: ? Why would you think that (Jack knows where the body is)?

b. A: Does Jack know where the body is? *PPQ*

B: Why would you think that (Jack knows where the body is)?

c. A: If Jack knows where the body is, we should ask him about it. *Conditional*

B: #Why would you think that (Jack knows where the body is)?

Comparing (236a–b), it seems *Why would you think that?* is marginally more acceptable after a polar question than after an assertion; of the same 110 respondents, only 51% judged (236a) as natural, while 66% accepted (236b). The felicity of a *would*-response is therefore a less reliable test than the *do*-response, but the trend remains that *would*-responses are best in the absence of strong commitment.

Applying these tests to the rising and falling tag questions above, a slight contrast appears:

(237) APPLYING THE TEST TO TAG QUESTIONS

a. A: Jack knows where the body is, doesn't he?

B: a. Why do you think that? 51% ok

b. Why would you think that? 70% ok

b. A: Jack knows where the body is, doesn't he.

B: a. Why do you think that? 61% ok

b. Why would you think that? 73% ok

It appears that the *do*-response, while judged licit by a majority of speakers with both rising and falling tag questions, is slightly more acceptable with the falling tag than the rising one. This corroborates the intuition that falling tag questions involve a stronger commitment to the truth of the anchor than rising tag questions. At the same time, the trend is weak; whatever the formal difference between these questions types, the effect on strength of commitment will have to be small in order to capture the data. [15]

[15]Note also that the acceptance rate for the *would*-response is constant at around 70% regardless of the tag

4.5.1.4 Closeness of anchor and tag

Another fact first recognized by Ladd (1981) is that tag questions occur with two different intonational contours over the junction between the anchor and the tag as shown in (238). Ladd terms these 'postnuclear' and 'nuclear' intonation, based on whether the tag has its own nuclear accent, or if it shares (and so comes after) the nuclear accent on the anchor. Note that postnuclear phrasing is incompatible with falling intonation, as shown in (239), a fact corroborated by Bolinger (1989).

(238) a. (What happy news! I must share it right away.) Amanda isn't home=is she? *postnuclear*

b. (What do you mean you can't find her?) Amanda isn't home / is she? *nuclear*

(239) a. # Amanda isn't home=is she. *postnuclear*

b. Amanda isn't home / is she. *nuclear*

Reese and Asher (2009) find fault with Ladd's characterization, but generally do agree that there are two kinds of phrasing available, one where the anchor and tag are close, and

---

question's intonation. This contrasts with acceptance rates for HNPQs, shown below:

(i) APPLYING THE TEST TO HNPQS

A: Doesn't Jack knows where the body is?

B: a. ?? Why do you think that? 32% ok

b. Why would you think that? 86% ok

The relative acceptability of the *would*-response follows readily from the analysis presented in the previous section; A is holding on to a commitment based on an older state of the context, so the counterfactual flavor of the *would*-response is most appropriate.

A summary of the results from the judgement task are included below.

(ii) ACCEPTANCE RATES FOR *Why do/would you think that?* (n=110)

Responses	**Decl**	**PPQ**	**HNPQ**	**Rising TQ**	**Falling TQ**
do only	44%	4%	9%	25%	25%
would only	13%	64%	63%	44%	37%
both natural	42%	2%	23%	26%	36%
both unnatural	2%	31%	5%	5%	2%

another where the phrases are more separated. Intuitively, they claim, closer phrasing indicates a stronger connection between the anchor and tag.

Reese and Asher (2009) follow this intuition to a view where tag questions with nuclear phrasing and postnuclear phrasing have radically different semantics and pragmatics. They claim nuclear tag questions involve an assertion followed by a question, the same way HNPQs do in their view, with largely conventionalized pragmatics sorting out the combined effects. Postnuclear tag questions are more closely integrated and so correspond to just one combined act, with a distinct semantics of a non-default (i.e., biased) question.

I will not review their account here, except to say that there's reason to think this view needs more nuancing; as shown previously with the rising versus falling tag data in (237), nuclear tag questions do not seem to involve the effects of a full assertion, only some kind of weak commitment.[16] Assuming assertion requires strong, actual commitment—arguably its defining feature—Reese and Asher's (2009) view only goes through if the pragmatic portion of the calculation can 'undo' strong commitments. More generally, we should be suspicious of treating nuclear and postnuclear tags questions as radically different, because experimental evidence shows that speakers can't reliably tell them apart (Lyn Walker, p.c.).

4.5.1.5 Polarity of anchor

I turn now to the issue of polarity. As the examples below show, the polarity of the anchor can be either positive or negative, as predicted by the fact that there are few restrictions on the form of the anchor, beyond that it must be declarative.[17]

(240) a. Thomas stole the wine / didn't he?

b. Thomas didn't steal the wine / did he?

The interpretive effects here are straightforward except for interactions with negation, as noted by Reese and Asher (2009). When a perfect storm of tag question dimensions collide, we find certain tags can express either the usual epistemic bias for the truth of the anchor,

[16] It has been established that asserting p requires both committing to p and the proposing to add p to the common ground, as Krifka (2012) and Farkas and Roelofsen (2012) show. Regardless of the intonational phrasing, tag questions seem to raise the issue $\{p, \neg p\}$, rather than the singleton $\{p\}$.

[17] Anchors for invariant tags like 'ok?' are not constrained in this way. For instance, it is possible to tag an imperative, to various effects: "Pick that up and put it in the trash, ok?"

or *no bias at all*. This alternation occurs only with a negative anchor and a postnuclear (and therefore rising) tag.

(241) a. Jack doesn't know=does he? (He was oblivious earlier.) *biased*

b. Jack doesn't know=does he? (I want to be the first to tell him!) *not biased*

Similar to HNPQs, it has been claimed that the two 'readings' above can be further disambiguated by including NPIs or PPIs:

(242) a. Jack doesn't know either=does he? *biased*

b. Jack doesn't know too=does he? *not biased*

The case that is biased for the truth of the anchor ($\neg p$) admits only NPIs, like Ladd's (1981) Inner-NPQs. The unbiased reading admits PPIs, suggesting as before that the anchor's negation is somehow not strictly part of the anchor.

As suggested by the context following (241b), the unbiased TQ is best when it conveys eagerness on the part of the speaker. If this eagerness is removed, the unbiased reading disappears:

(241b′) *Jolene and Billy are discussing the fact that Jack just won an award. She says:* (# I don't really care that much, but) Jack doesn't know=does he? (I'll go tell him for you.)

This suggests that while the speaker does not convey the sort of evidence-based bias that is the focus of this chapter, she still indicates a preference for the anchor's truth when she uses this kind of TQ in the epistemically-unbiased way.[18] As will be shown in §4.5.2, tag questions in general *do* express robust, evidence-based bias, so the use explored above represents a narrow exception extension of the tag question's normal use.

[18]The reasons for this residual bias are unsurprising if we recall van Rooy and Safarova's (2003) utility-theoretic approach to the pragmatics of questions. The view of utility they present takes into account both beliefs and desires, meaning that even if an utterance does not express bias of the sort captured with the evidential base so far under the current analysis, it is plausible that there are other avenues for this preference to come across to the speaker's interlocutors.

4.5.1.6 Polarity of tag (same vs. reverse)

To end this parade of tag question variation, note that one of the most intriguing ways that tag questions can vary is in the relative polarity of the tag. While tag questions that reverse the polarity of the anchor are the most canonical in terms of what comes to mind when we think of tag questions, it is also possible to for the tag to have the same polarity as the anchor, with different effects.

(243) a. Anna loves Mr. Bates / doesn't she?

b. Anna loves Mr. Bates / does she?

(244) a. Anna doesn't love Mr. Bates / does she?

b. Anna doesn't love Mr. Bates / doesn't she?

The final case in (244b), a negative anchor with a negative tag, requires some justification; these tags are quite restricted, occurring most easily in contexts where the speaker is skeptical:

(245) A: Anna doesn't love Mr. Bates. That's crazy!

B: (Oh, so) Anna doesn't love Mr. Bates / doesn't she? Then where'd this love letter come from?

Positive same-TQs like (243b), on the other hand, occur somewhat more frequently, with a diverse range of effects. Crucially for present purposes, there is a great deal of variation in how much commitment and speaker bias these questions express. Compare (246–247) below:

(246) A: I don't know whether they still do dancing there, but the maypoles's still there.

B: It's still there, is it?

(247) A: No, I'll sit in my normal spot

B: Oh, okay.

A: Oh no, you're taping us, are you?

. (Kimps 2007, citing COBUILD Corpus)

In (246), B's same TQ functions as a rather elaborate acceptance move; the speaker expresses neither incredulity nor independent knowledge about the maypole. The situation in (247), on the other hand, evokes many of the examples of speaker bias already discussed; A seems to have concluded that it is likely he is being taped, based on B's behavior. Finally, note that these questions can also indicate that the speaker is skeptical, similar to (245). The two colorful examples below are again due to Kimps (2007).

(248) A: I promise you.

B: Are you man enough to threaten me personally or

A: Oh sir I'm more than man enough, but I don't need to do those sort of things myself.

B: Oh you have other people do you?

A: Er yes.

(249) 'It's Iraq, not Maoist China' I snapped. Without blinking she said, '[W]hy do I have to understand the newspaper? Why can't I just read it for enjoyment?' **Turks read about genocide for enjoyment, do they?**' I shrieked. To abate the feeling of utter pointlessness that swept over me, I devoted myself to my marmalade [. . .]
. (Kimps 2007, citing COBUILD Corpus)

In short, same TQs can express the full gamut of bias: they can suggest the speaker favors the anchor, disbelieves it, or merely accepts it. The modest conclusion that can be drawn from this is that the same TQ is not inherently biased the way that other TQs and especially HNPQs are. I therefore will not pursue an analysis of these questions further in this section, focusing instead on the reverse TQs (to be characterized in the next subsection).

Beyond the fact that the status of same-TQs as biased questions is in doubt, it is worth noting that any discussion of same-TQs is complicated by the fact that their frequency and use differ between British and American English (Kimps, 2007). Corpus work by Tottie (2006) suggests that speakers of British English employ same-TQs twice as often as American English speakers (8% versus 4% of TQs). Across all the categories discussed above, British English speakers use TQs about equally in *confirmatory* and *facilitating* functions. Confirmation situations, according to Tottie (2006), are those where the speaker holds a (weak) belief and wishes to strengthen it, while facilitation situations are those where the

speaker seems more or less convinced and merely seeks to "involve the addressee." American English speakers, on the other hand, are much more likely to use TQs in order to facilitate. Based on these discrepancies, it is very likely that the functions of various TQs differ between American and British speakers. The discussion that follows is based on American English usage (which extends as far as capturing the trend toward facilitation).

4.5.2 Tag questions are biased questions

Reverse TQs seem to be biased questions, but how do they differ from HNPQs? To investigate this issue, I propose borrowing the battery of tests employed by Sudo (2013) for HNPQs. While speaker beliefs (SB) and contextual evidence (CE) do not have independent status as invokable evidential bases in my analysis thanks to the findings in §4.3.2, the licit configurations still reveal the relevant generalizations. Restricting attention to rising versus falling TQs, it seems that rising TQs are felicitous whenever there is positive SB for the anchor, while falling TQs require positive CE for it.

4.5.2.1 Testing speaker belief and contextual evidence

The examples employed in this section are variations on a theme, so I will briefly sketch the relevant details. Across all of the examples, Marcine has begun preparing dinner when Gertrude gets home, and Gerturde asks a tag question to determine whether there will be a guest over for dinner. Gertrude's speaker beliefs and the contextual evidence Marcine provides differ across the examples. Gertrude's prior speaker beliefs are those she brought home with her: In some cases, she thinks Teddy, a mutual friend, was planning to join them, in some she explicitly thinks he was not planning to join them, and in other cases, she has no beliefs one way or the other. Contextual evidence is supplied by the state of Marcine's preparations. Positive evidence for a guest comes from an extra place setting at the table, while negative evidence comes from the normal number of place settings. Neutral contextual evidence is supplied when Marcine is just barely getting started preparing and so has not yet had a chance to 'bias the context' either way. This contextual evidence is assisted by background assumptions about the purpose of place settings, as well as one crucial bit of conditional knowledge: Teddy is the only person who ever comes over for dinner. Therefore, if there is evidence that *someone* is coming over, it is evidence that *Teddy*

is coming over.

The examples are presented in blocks for easier comparison. In the first set, Gertrude's positive speaker belief is held constant; she always has prior evidence that Teddy is coming over. Each example is also labeled to indicate what the default ("current" here, for clarity) and prior evidential bases support—p, $\neg p$, or neither.

(250) *Gertrude notices an extra place setting.* current: p
Independently, she believes Teddy was planning to come over for dinner. prior: p
G: a. Teddy's coming to dinner, isn't he?
b. Teddy's coming to dinner, isn't he.

(251) *Marcine has just started pulling out dishes for dinner.* current: –
Independently, she believes Teddy was planning to come over for dinner. prior: p
G: a. Teddy's coming to dinner, isn't he?
b. # Teddy's coming to dinner, isn't he.

(252) *Gertrude notices the normal number of place settings.* current: $\neg p$
Independently, she believes Teddy was planning to come over for dinner. prior: p
G: a. Teddy's coming to dinner, isn't he?
b. # Teddy's coming to dinner, isn't he.

These examples reveal that a rising TQ is licensed whenever the asker has positive speaker beliefs; the value of contextual evidence makes no difference. A falling TQ, on the other hand, is only licit if there is also prior evidence available.

(253) EMPIRICAL GENERALIZATION FOR TQS, #2: **Prior evidence conditions**
a. Contexts where the speaker has prior evidence in favor of the anchor admit **rising TQs**.
b. **Falling TQs** are only admitted if the speaker also has current evidence in favor of the anchor.

To test this generalization and continue discovering what generalizations can be made about current evidence conditions, we must take away the positive prior evidence. In the next set of examples, Gertrude lacks any particular beliefs about Teddy's dinner plans.

(254) *Gertrude notices an extra place setting.* current: p
When there's a guest, it's often Teddy, but Gertrude hadn't heard anything about Teddy's plans. prior: –

G: a. # Teddy's coming to dinner, isn't he?

b. Teddy's coming to dinner, isn't he.

(255) *Marcine has just started pulling out dishes for dinner.* current: –
When there's a guest, it's often Teddy, but Gertrude hadn't heard anything about Teddy's plans. prior: –

G: a. # Teddy's coming to dinner, isn't he?

b. # Teddy's coming to dinner, isn't he.

(256) *Gertrude notices the normal number of place settings.* current: $\neg p$
When there's a guest, it's often Teddy, but Gertrude hadn't heard anything about Teddy's plans. prior: –

G: a. # Teddy's coming to dinner, isn't he?

b. # Teddy's coming to dinner, isn't he.

It appears that the rising TQ is never licensed when positive prior evidence is lacking, meaning that neutral speaker belief is not sufficient grounds for such a question. The falling tag, however, is licit when the asker lacks pre-existing speaker beliefs, but is provided positive contextual evidence. Taken together with the first set of judgements, this suggests that current evidence in favor of the anchor may be sufficient to license a falling TQ, but to be certain, we must also consider cases where the speaker's prior evidence runs counter to the anchor's proposition.

Predictably, a similar pattern emerges with negative speaker belief:

(257) *Gertrude notices an extra place setting.* current: p
When there's a guest, it's Teddy, but Gertrude thought Teddy wasn't coming. prior: $\neg p$

G: a. # Teddy's coming to dinner, isn't he?

b. Teddy's coming to dinner, isn't he.

(258) *Marcine has just started pulling out dishes for dinner.* current: –
When there's a guest, it's Teddy, but Gertrude thought Teddy wasn't coming. prior: $\neg p$

G: a. # Teddy's coming to dinner, isn't he?

b. # Teddy's coming to dinner, isn't he.

(259) *Gertrude notices the normal number of place settings.* current: $\neg p$
When there's a guest, it's Teddy, but Gertrude thought Teddy wasn't coming. prior: $\neg p$

G: a. # Teddy's coming to dinner, isn't he?

b. # Teddy's coming to dinner, isn't he.

Once again, the rising TQ is never lict with negative speaker beliefs, but the falling TQ is, so long as there is positive contextual evidence. In short, the generalization suggested above seems to be accurate:

(260) EMPIRICAL GENERALIZATION FOR TQS, #3: **Current evidence conditions**

a. Contexts where the speaker has current evidence in favor of the anchor admit **falling TQs**.

b. **Rising TQs** are admitted only if the speaker also has prior evidence for the anchor.

4.5.2.2 Summary of bias conditions so far

The charts below summarize the data presented above, in contrast to the facts for HNPQ from the previous section. Previously, it was shown that HNPQs are generally licensed by positive speaker belief and negative contextual evidence.[19] The analysis then generalized from speaker belief and contextual evidence to prior and default current evidence. More support for this move appears shortly.

[19]Recall that counter to Roelofsen et al. (2012), I claim that there must always be current evidence against p, even if that means accommodating it. Cases which seem to feature neutral current evidence are only superficially this way. See §4.3.1.1.

(261) FELICITY FOR HNPQS

		current		
		p	$-$	$\neg p$
prior	p	#	#	✓
	$-$	#	#	#
	$\neg p$	#	#	#

This section has shown that the behavior of TQs is much simpler; rising TQs are licensed whenever the speaker holds positive speaker beliefs, and falling TQs are licensed whenever there is sufficient contextual evidence to warrant an inference:

(262) FELICITY FOR RISING TQS

		current		
		p	$-$	$\neg p$
prior	p	✓	✓	✓
	$-$	#	#	#
	$\neg p$	#	#	#

(263) FELICITY FOR FALLING TQS

		current		
		p	$-$	$\neg p$
prior	p	✓	#	#
	$-$	✓	#	#
	$\neg p$	✓	#	#

Rising and falling TQs are therefore similar to HNPQs in that they express some kind of positive bias on the basis of evidence, but unlike HNPQs, they do not need a conflict between two evidence sources.

4.5.3 The discourse effect of tag questions

With the data generalizations introduced, this section provides my formal analysis for TQs. I argue that at their core, tag questions are in fact questions. This runs counter to Reese and Asher's (2009) argument, whereby TQs (as well as HNPQs) involve both assertive and questioning acts, but in fact the spirit of the analysis is quite similar; it maintains the idea that biased questions fall into an intermediate position between default questions and default assertions, but without weakening the notion of assertion itself to account for bias.

Under the discourse model introduced in Chapter 2 and nuanced in Chapter 3, default assertions and default questions share a basic template. Both kinds of acts first place the denotation of the utterance's content on the Table, a register of issues under consideration to become common ground (c.f. Farkas and Bruce 2010; Farkas and Roelofsen 2012). An

assertion places a proposition on the Table, while a question places the set of its answers. Next, in both cases the speaker registers a commitment to the utterance's content. In the case of an assertion, this straightforwardly means the speaker commits to the proposition he proposes to add to the common ground. In the case of a polar question, the speaker makes a trivial commitment to $[p \vee \neg p]$, which has the effect of making the speaker responsible for any presuppositions of his question, and also implicates openness to each answer. I will term this the *openness commitment* of a question (c.f. Farkas and Roelofsen 2012). Crucially, these commitments share a core feature: in the absence of markers that indicate otherwise, they are made relative to the default evidential base. This means both the commitment resulting from a default assertion and the openness commitment of a question are strong, independent commitments underwritten by the speaker's current default base, which itself is comprised of the speaker's private beliefs and any shared contextual evidence. This is because the base is *not* characterized as weak, dependent, or prior. Simply by specifying the literal meaning of an expression and requiring that speakers commit to the totality of what they place on the Table, the discourse effects of both default assertions and questions fall out automatically.

This template offers a few ways that either one of the above acts could be modified in order to move the overall discourse effect somewhere between the two, which we have already seen. A rising declarative, for example, is associated with an assertion-like act which raises a single propositional issue but only weakly commits the speaker to that proposition as a way of attenuating its overall effect. The default question, on the other hand, can be bolstered by adding additional commitments to reveal that the speaker is not equally open to both alternatives.

At this point it is again worth stressing that the dividing line between non-default assertions and non-default questions is rather unclear (Farkas and Roelofsen, 2013). For example, even though we might expect that assertions (default or otherwise) can be used to inform, given that the issue they place on the Table is a single proposition, the weakness of the speaker's commitment can make them inappropriate for that use. Similarly, we might expect that because default questions can never be used to inform, non-default questions cannot either. Even this is difficult to show, however, because of the ways that non-default question forms can be used rhetorically. See §4.6 for a thorough discussion of these facts.

For the time being, however, the idea that TQs are questions to their core remains a hypothesis, albeit one that returns theoretical dividends over the course of this section.

In short, it will be proposed here that both kinds of TQs discussed here carry the default effects of a polar question: they raise the issue $\{p, \neg p\}$ and commit the speaker to the (trivial) disjunction of the two answers, which signals openness to their resolution as describe above, but they *also* carry weak commitments to signal that the speaker's willingness to commit to both alternatives is not equal.

Because TQs do not require a bias conflict, their non-default effects can be modeled with a single contingent commitment, rather than the two needed for HNPQs. This means there is a three-way distinction among default positive polar questions, HNPQs, and TQs with regard to the 'secondary' questions that their forms hint at. Default PPQs in fact have no secondary question; they simply ask whether p holds. The secondary question lurking behind HNPQs derives from their conflicting contingent commitments; an informal paraphrase would be that the speaker wishes to know if she must really give up her predilection for p in the face of recent anti-p evidence. The secondary question behind tag questions is more straightforward. Instead of clarification, it is a matter of confirmation; the positive answer is strongly favored.

Rising and falling TQs seem to differ in the reasons that underwrite weak commitment to the positive answer. Based on the above discussion, it is once again tempting to couch this in terms of commitments given the proposed bases E_{SB} and E_{CE} directly, but as previously argued, this is a dubious prospect. Instead, the commitments should be framed in terms of the default (current) base E_{DEFAULT}, and the prior base E_{PRIOR}.

Perhaps the more interesting case under the evolving view is the falling TQ, so we begin there. These questions appear to be licensed by positive contextual evidence, which means recent additions to the context must favor the TQ's anchor. This matches the intuition that falling TQs generally express a desire to confirm an inference. To indicate that an inference has occurred on the basis of recent evidence, the speaker must indicate that at a previous state, the inference wasn't available. This is possible by shifting the reasons for the otherwise-trivial commitment to $[p \vee \neg p]$ to be based on a previous state of speaker's default base, as shown in (ii) below:

(264) EFFECTS OF UTTERING A FALLING TQ QUESTIONING p:

i. Speaker raises the issue $\{p, \neg p\}$

ii. Speaker makes a commitment to $[p \vee \neg p]$ given evidential base E_{PRIOR}

iii. Speaker makes a commitment to p given the evidential base E_{WEAK}

The above provides the complete discourse effects of a falling TQ. By virtue of being a question, the speaker raises the issue of whether p and commits to the union of the question's answers. Unlike earlier examples, however, this commitment is made relative to the speaker's prior base, indicating that in the recent past, the speaker was open to both options, but leaving ambiguous whether both options are still entertained. Then, the speaker makes an additional commitment to p given his weak base, i.e., his current default base, with an indication that commitment is highly mutable.

Taken together, these two commitments reveal that the speaker rather strongly biased toward p based on inference over contextual evidence. The utterance is in fact barely a question, because the speaker does not overtly signal that he is open to both answers—although because both answers were put on the Table, some degree of openness is implied.[20]

Under this view, falling TQs are a kind of mirror of HNPQs. Rather than expressing an unwillingness to give up an old pro-p view, they express an eagerness to keep a new pro-p view.

Rising TQs are a simpler case. They rest on the availability of prior evidence in favor of the anchor, and are oblivious to recent changes in (contextual) evidence. This can be captured simply by requiring that the evidential base that underwrites the trivial commitment matches the base underwriting a commitment to p:

(265) EFFECTS OF UTTERING A RISING TQ QUESTIONING p:

i. Speaker raises the issue $\{p, \neg p\}$

ii. Speaker makes a commitment to $[p \vee \neg p]$ given the default evidential base E_{DEFAULT}

[20] It was previously noted that in American English, TQs are often used when the speaker is all but convinced that p, but seeks to 'facilitate' conversation by means of the biased question. If the openness commitment is underwritten by the prior base, this behavior makes sense; the speaker *was* open to both options, but now favors p quite strongly.

iii. Speaker makes a commitment to p given the evidential base E_{WEAK}

The only change here from the falling TQ is that in (ii), the openness commitment is conditioned on the speaker's current default base, rather than the prior base. This says, rather straightforwardly, that the speaker is open to either answer, but puts more credence in p.[21]

Note that while with falling TQs, employing the default base implied a change from earlier states, this was in fact only the case because another commitment was being made given the prior base; the default base is by its very nature unmarked, and so implies no changes or contrasts on its own. By failing to make reference the prior base, the implication is that the difference between the previous state of the speaker's default base and its current state does not in fact matter.

4.5.3.1 Novel data and solutions

§4.5.1–4.5.2 Provided the following three data generalizations for rising and falling TQs, all of which are explained under the current account. In reviewing these generalizations, some novel data is brought to bear to show that once again, the movement from working in terms of speaker belief and contextual evidence to the speaker's total evidence over time is the correct approach.

The first empirical generalization for TQs is that they are self-consistent:

(233) EMPIRICAL GENERALIZATION FOR TQS, #1: **Speaker consistency**
(Variant) tag questions are not understood as self-contradictory, even though their forms suggest they would be.

This is because unlike the sequences questions and answers the resemble, TQs are single, biased utterance, not a composite of two:

(232) a. Milo owns a VCR, doesn't he?

[21] An alternative formalization of the rising TQ would make both of its commitments relative to the prior base. On its face this seems right—it transparently amounts to ignoring recent contextual evidence—but it is not a configuration that should be permissible. Regardless of which evidential bases are employed, the speech act itself takes place relative to the current context, so to not make reference to that context would render the question an odd sort of counterfactual.

b. # Milo owns a VCR, doesn't Milo own a VCR?

c. # Milo owns a VCR, doesn't he own a VCR?

The fact that TQs are not contradictory follows directly from the account; in (232a) above, the speaker makes a trivial commitment to Milo either owning or not owning a VCR, as well as a weak commitment to the anchor answer. The weak base underwriting the latter commitment becomes inviable if more support for the proposition does not become available in the discourse, meaning that the speaker presents himself as open to either answer, albeit with a strong preference toward replacing his weak commitment with a default strong one. In (232b–c), by contrast, the speaker performs two utterances in sequence, and the effects of the two are not compatible. After his assertion, the speaker has placed the proposition that Milo owns a VCR on the Table and committed to it with his default base. Then, he immediately places both the proposition and its negation on the Table, makes the openness commitment given his default base. While it is consistent for the speaker to be committed to p and $[p \vee \neg p]$ given the same base, it is not consistent with the purpose of such an act, namely to receive an answer to dependently commit to.[22] The utterances are therefore self-contradictory.

In terms of the evidence conditions for rising and falling TQs, the earlier discussion culminated in the following:

(253) EMPIRICAL GENERALIZATION FOR TQS, #2: **Prior evidence conditions**

a. Contexts where the speaker has prior evidence in favor of the anchor admit **rising TQs**.

b. **Falling TQs** are only admitted if the speaker also has current evidence in favor of the anchor.

(266) (260) EMPIRICAL GENERALIZATION FOR TQS, #3: **Current evidence conditions**

a. Contexts where the speaker has current evidence in favor of the anchor admit **falling TQs**.

b. **Rising TQs** are admitted only if the speaker also has prior evidence for the anchor.

[22] In default information-seeking cases.

To prove that these generalizations are correct to be phrased in terms of current and prior evidence, rather than speaker beliefs and contextual evidence, consider the example repeated below:

(207) *Happy hour approaches; Zabi and Kazuko need a drink. Zabi can't easily recall where the brewery is.*

Z: Do you remember where the brewery is? Never mind; **it's on Swift, isn't it**.

In this example, the falling TQ is licit even though there is no contextual evidence at play at all—Zabi never gives Kazuko a chance. Instead, what is contrasted here is the state of Zabi's default base prior to recalling the information he needed, and after. This follows from the account: Here, Zabi first asks a default question, placing the issue of whether Kazuko remembers where the brewery is on the table, and making the openness commitment to the possibility of both answers. He then immediately asks his falling TQ about Swift Street, making the openness commitment to the possibility of both answers given his *prior* base, and a weak commitment to the anchor possibility given his current base. Precisely because no contextual information has been able to change Zabi's default base between the utterance, we understand that the change must have had to do with Zabi's speaker beliefs; he remembered something. Covering this example is only possible if the change in evidence doesn't have to be contextual evidence, because no satisfactory definition of contextual evidence is going to include "things the speaker just remembered."

The remaining sections of this chapter provide a more general discussion of what the differences between HNPQs and TQs, and what these differences reveal about the system of evidential bases that underwrites commitments. In short, the space of variation for these bases is more constrained than the preceding discussion at first suggests—a welcome result on the path toward building an explanatory model of discourse.

4.6 Rhetorical uses

So far, this chapter has assumed that all polar questions are genuine *information-seeking* question, but this is of course not the whole story. Semantic questions can be used to achieve a variety of discourse effects. One of the most intriguing uses is the family of *rhetorical questions*. Under these uses, the speaker does not in fact seek the answer to his

question. Rather, he seeks to show that the answer is already known to his interlocutors, or is otherwise trivial (Lee-Goldman 2006).

In recent work it is generally thought that there is no such thing as a unary rhetorical speech act. Rather, Rohde (2006), Caponigro and Sprouse (2007), and others argue that rhetorical questions share the syntax and semantics of ordinary questions. Recall that under the analysis presented here, the default discourse effects of a question (or any utterance) result from an interaction between the form of the utterance and the context. This requires, then, that the discourse effects of rhetorical questions build on those of default questions as well. Specifically, I assume that rhetorical questions perform their normal discourse effects, but are immediately resolved when at the conclusion of the speech act it is clear to the interlocutors that no answer is necessary.

The implied answer to a default PPQ can be positive or negative:

(267) Is the pope Catholic? *yes*

(268) Are you going to eat that? *yes/no*

(269) Do I look like a doctor? *no*

In all cases, the question proceeds as normal: The speaker raises the issue (e.g., {the Pope is Catholic, the Pope is not Catholic}) and commits to the trivial union of the answers. In fact, the common ground already includes the information needed to settle the issue, so no further action is necessary. The question self-settles and the overarching pragmatic effect involves drawing an analogy to the obviousness of the question's resolution.

Biased questions face further restrictions when used rhetorically. This section presents the data surrounding the use of HNPQs and rising and falling TQs rhetorically, and shows that these deviations from the default fall out from the analysis presented here without any modifications.

4.6.1 Rhetorical HNPQs

When HNPQs are used rhetorically, the implied answer must be the positive one:

(270) Isn't the pope Catholic? *yes*

(271) Aren't you going to eat that? *yes/#no*

(272) Don't I look like a doctor? *#no*

Happily, this falls out from the account presented in this chapter. Consider the example below, in a slightly richer context.

(273) *Obama issues an order to shut down PRISM.*

A: He can't just order us around like that!

B: Sure he can! Isn't he the president?

In this example, A and B share knowledge of the President's identity, which means the context entails the answer to B's question. This renders the question rhetorical. At the same time, however, the full discourse effects of the question must come into effect. B raises the issue of whether Obama is the president (ignoring the fact that the issue is self-resolving). She commits to the trivial proposition that Obama either is or isn't the President. She also makes the commitments that are indicated by the HNPQ form: A weak commitment to Obama's being president ($= p$), given her prior base, and a weak commitment to Obama's not being President ($= \neg p$), given the current state of her default base. The issue then resolves itself (in favor of p), which entails dismissing the conflicting evidence that underwrites B's commitment to $\neg p$. Crucially, however, the conflicted evidence must have been present, just like is required for information-seeking uses of HNPQs.

The examples below illustrate this familiar conflict requirement.

(274) *Mom steps into Todd's messy bedroom and gives a stern look.*

M: a. Are you going to clean your room today?

b. Aren't you going to clean your room today?

(275) *Mom and Todd are having breakfast when she shoots him a stern look.*

M: a. Are you going to clean your room today?

b. # Aren't you going to clean your room today?

in (274), Todd's mother can use both the PPQ and the HNPQ rhetorically—the first because it is generally available, and the second because the mess in Todd's room conflicts with her

expectations.[23] In (275), the rhetorical HNPQ is infelicitous unless we enrich the context to build in a conflict.

4.6.2 Rhetorical TQs

Similarly, rhetorical uses of rising TQs also must be positive:

(276) The pope is Catholic, isn't he? *yes*

(277) You're going to eat that, aren't you? *yes/#no*

(278) I look like a doctor, don't I? *#no*

Just as in the case of rhetorical HNPQs, this is because the conditions for felicitous use of the question form must be met, even if the question will self-resolve. Rising TQs require grounds for the speaker to be biased toward p, so using one in a context that will be able to automatically resolve for $\neg p$ is impossible.

Falling TQs, in contrast, are never rhetorical:

(279) # The pope is Catholic, isn't he. *#yes*

(280) # You're going to eat that, aren't you. *#yes/#no*

(281) # I look like a doctor, don't I. *#no*

Falling TQs are felicitous whenever the speaker was previously open to both possible answers, but has since concluded one answer is more likely. In rhetorical situations, the first condition cannot be met; if the speaker was uncertain until recently of whether p, the situation cannot be one where the rhetorical use is licit. Additionally, falling TQs already have a kind of rhetorical flavor, because they only signal a limited openness to the negative answer anyway. It's therefore arguably redundant to try to use one rhetorically.

More generally, using a rhetorical question involves temporarily ignoring the fact that the issue being raised will self-resolve in the face of the wider context, but non-trivial commitments do not tolerate make-believe. Raised issues merely interact with the broader context, while commitments are part of its fabric. These facts come for free with the analysis presented here.

[23] Again, we see the importance of separating a speaker's private expectations from those she displays publicly. Todd's mother need not actually have expected Todd to have cleaned up in order to present herself as believing so.

4.6.3 Answering rhetorical questions

Some of the most puzzling rhetorical questions are like (282) below, which are clearly biased but not conflicted:[24]

(282) Isn't this the best sandwich you've ever tasted?[25]

Under the analysis provided here, this question should only be licit if the interlocutors can locate or accommodate an appropriate conflict between two states of the speaker's default base. I argue that the licensing conditions for an example like this are analogous to those for the flooding example in (194); it is possible for a speaker's prior base and his default one to differ in the necessary way if the speaker expected (or pretended to expect) something to happen and it did not. In (282), the 'missing' evidence is that no one has yet remarked on the sandwich's superior taste and/or mouthfeel.

This approach may seem farfetched at first, until we consider an interesting fact about rhetorical HNPQs like this one: they require answers.

(282′) *Layla and Maddrox are sharing a picnic at the park.*

L: Isn't this the best sandwich you've ever tasted?

M: It's *so* good. / The avocado is perfectly ripe. / #Pass me the chips. / *# (No answer)*

This follows directly from what I've claimed about these questions. If the asker is treating the fact that his addressee hasn't commented on the sandwich's quality as evidence against its superb nature—even in momentary jest—it triggers a weak commitment to the notion that the current context is one where perhaps the sandwich is not the best that B has ever tasted. If the addressee wishes to avoid leaving the speaker out on a limb, he'd best respond. Notice also that irrelevant responses like *Pass the chips* above are easily construed as avoidance, and give the asker grounds for actual doubt.

[24] Examples of this sort appear in Anderbois (2011), using examples like *Isn't he the cutest baby you have ever seen!?* These are problematic, however, because any discussion of babies requires an unusual level of enthusiasm and conversational engagement. This obscures the facts below, so a more mundane example is used.

[25] Interestingly, some speakers prefer "Isn't this like the best sandwich you've ever tasted?," regardless of how frequently they report using the interjective *like*.

The only cases where these rhetorical questions do not require answers is when the addressee has in fact already supplied the appropriate contextual counter-evidence prior to the question itself. For example, in (282″), A's exclamative clearly conveys that he does think the sandwich is good, but B's rhetorical question is still licit.

(282″) *Layla and Maddrox are sharing a picnic at the park.*

L: This sandwich is *so* good.

M: Isn't this the best sandwich you've ever tasted?

A needn't reply in this case. Just like all rhetorical question situations, in this context B must ignore the fact that her question is already settled in the context until just after the utterance is complete. In this case, this requires temporarily ignoring A's exclamative, which means the conditions are right to pretend that A hasn't made his position on the quality of the sandwich clear yet. When B's utterance is complete and she again takes stock of the whole context, not only is the question answered, but the false evidential conflict is resolved too.[26]

4.7 Constraining expressions of bias

The system of bias elaborated in this chapter has introduced a number of ways that commitments employing the prior base can vary. This final content section discusses the predictions of the system to show that the freedom proposed in this chapter does not over-generate. Rather, when combined with a few reasonable pragmatic restrictions, the system predicts exactly the set of bias configurations for polar questions that are attested in English.

4.7.1 Dimensions of variation in commitments

Farkas and Roelofsen's (2012) discourse model offers just two dimensions for variation in commitment: conditionality and sourcehood. The conditionality of a commitment has three levels; a speaker can make an actual commitment, a conditional one, or no commitment at all. The sourcehood of a commitment is binary; a speaker can commit as source or dependent. Given that it is not meaningful to make 'no commitment' as source or dependent, this

[26] As a final note on this point, the same facts as for (282‴) hold if Layla's initial utterance is itself a rhetorical question, such as *How good is this sandwich?*.

means there are five types of commitments a speaker can make. When a speaker raises an issue, he can in principle commit to any of the alternative propositions within it, using any of those five types.

In my model, conditionality and sourcehood are absorbed into a more general component that registers the characterization of the evidential base that underwrites the commitment. The base, by default, is comprised of speaker's beliefs and contextual evidence, but it can be winnowed down and by invoking non-default characterizations. These include characterizations in terms of evidential source, as well as discourse-specific notions like dependence (Ch 2) and relative authority (Ch 3). The base even captures the idea that some commitments are more mutable than others by means of the weak base characterization. In this chapter, one further non-default base was proposed, the prior base. This base is particularly powerful because it allows the speaker to reach back and reference his default base at a previous point in the discourse. Is this a parsimonious addition?

One point to keep in mind before tackling this question is that even if the addition were not parsimonious in terms of its predictions about bias configurations, the account presented is maximally parsimonious in other ways; it builds a deep connection between bias and evidentially, and derives bias simply by generalizing the notion of the evidential base and commitment relative to it. That said, the prior base is, independently, unproblematic.

Recall from Chapter 2 that certain bases, namely counterfactual ones like a base characterized as containing hearsay known to be false (repeated below):

(62) UNATTESTED: Percy wrote a book [false-rumors-ev]

The reason this base is unattested, and the metric by which to judge the prior base, relates to the deeper purpose of evidentiality in communicative terms: marking evidence is only relevant if its source affects the reliability of the at-issue information in a consistent way. A base is appropriate to the extent that it advances this goal, which the prior base, in successfully helping to model biased questions, clearly does. On the basis of facts about the felicitous use of (esp. falling) TQs, it is suggested here that speakers can in effect opt to ignore a portion of the default base in certain circumstances. This is a move is marked, but it is extremely useful as a way of asking about how recent changes in the discourse affect a speaker's ability to commit. But how many ways can this prior base be invoked? Are they

all attested? With these questions in mind, I turn now to the space of variation I predict in polar questions. I conclude that the approach more parsimonious than it first appears, allowing only those configurations of prior bias that are attested in the English system.

4.7.2 Variation in polar questions

Polar questions offer three possibilities for commitment: p, $\neg p$, and the trivial $[p \vee \neg p]$. In addition, each of these commitments could in principle be 1) absent, 2) conditioned based on the speaker's default (current) base, or 3) conditioned on the speaker's prior base.[27] This is already a lot of variation, but it can be paired down almost immediately. First, it is a reasonable assumption that a polar question will always involve its openness commitment to $[p \vee \neg p]$, and that this commitment will never be weak. Second, any commitments to p or to $\neg p$ *must* invoke the weak base, otherwise they will conflict with the openness commitment.

More generally, I assume that at the level of discourse representation, all questions 'look' information-seeking. Rhetorical questions and quiz questions, as discussed above and in Chapter 3, involve the speaker temporarily presenting himself as though his request for information is genuine—even to the point where the speaker might use hedges or certainty-reducing devices to enhance the effect. This was shown in Japanese with the tentative copular form *daroo/desyoo*, which is preferred over the regular form *da/desu* in quiz questions.[28] This leaves eighteen possible commitment configurations for polar questions, as

[27] I exclude polar questions with evidential markers, etc.

[28] In a related note, biased questions can sometimes be used as quiz questions, too. In these cases, the emulated bias the question expresses can either be used to guide the 'student' (283) or challenge him (284).

(283) T: What's the next smallest prime number after 11?
S: Um. . . 5.
T: Isn't 7 prime?

(284) T: Shouldn't this be 'and' instead of 'or'?
S: No, we distributed the negation, so we flip the sign.
T: Yes, that's right.

Note that this is not an example of Ladd's (1981) ambiguity. Rather, what these data suggest once more is that rhetorical questions involve a kind of discourse-level play acting.

shown in Fig. 4.1.

	$[p \vee \neg p]$	p	$\neg p$	**Attested?**	**Details**
a.			–	**Yes**	**Default Q**
b.		–	current	No	Avoid reversing (4.7.2.2)
c.			prior	No	Superfluous prior (4.7.2.3)
d.			–	**Yes**	**Rising TQ**
e.	current	current	current	No	Inconsistent
f.			prior	No	Superfluous prior (4.7.2.3)
g.			–	No	Superfluous prior (4.7.2.3)
h.		prior	current	**Yes**	**HNPQ**
i.			prior	No	Inconsistent
j.			–	No	Lacks current commitment (4.7.2.1)
k.		–	current	No	Avoid reversing (4.7.2.2)
l.			prior	No	Lacks current commitment (4.7.2.1)
m.			–	**Yes**	**Falling TQ**
n.	prior	current	current	No	Inconsistent
o.			prior	No	Superfluous prior (4.7.2.3)
p.			–	No	Lacks current commitment (4.7.2.1)
q.		prior	current	No	Superfluous prior (4.7.2.3)
r.			prior	No	Inconsistent

Figure 4.1: Possible commitment configurations for polar questions

The bolded entries above show the attested configurations. The rest of this section shows how the other possible configurations are ruled out. For example, it is impossible to commit to both p and $\neg p$ given the same evidential base. This immediately rules out the (e,i,n,r) configurations. Ruling out the others requires going more in depth into what it means to ask a mannerly question and what the purpose of expressing bias is.

4.7.2.1 Commit in the present

Configurations (j,l,p) each lack any commitment based on the current, default base.

(285) EVIDENCE CONFIGURATIONS LACKING A CURRENT COMMITMENT

Question form	$[p \vee \neg p]$	p	$\neg p$
j.	prior	-	-
l.	prior	-	prior
p.	prior	prior	-

Questions induce choices and drive conversation forward by narrowing in on mutually exclusive answers. They are inherently forward-looking in this regard. Commitments that are made holding on to an early state of one's default base are therefore only useful insofar as they inform future decisions about the partitioning of the discourse. Without a default commitment to contrast with, these configurations would convey useless information.

4.7.2.2 Avoid reversing

The configurations (b,k) are notable because they mirror the rising TQ (d) and falling TQ (m) configurations.

(286) ACTUAL AND ALTERNATIVE CONFIGURATIONS FOR TAG QUESTIONS

Question form	$[p \vee \neg p]$	p	$\neg p$
d. Rising TQ	current	current	-
b. Rising TQ′	current	-	current
m. Falling TQ	prior	current	-
k. Falling TQ′	prior	-	current

The intuitive problem with these alternatives is that they are attempting to do the work of polarity; an equivalent configuration can be expressed simply by using a negated form of the question, as shown below.

(287) This isn't Bates' fault, is it? *(negative) rising TQ*
$p =$ It is not Bates' fault.
Conveys: Sp is mutably committed to p; *yes* $= p$

(288) * Is this Bate's fault [my inkling is *no*]? *'equivalent' rising TQ′*
$p =$ It is Bates' fault.
Conveys: Sp is mutably committed to $\neg p$; *yes* $= p$ (mismatch!)

Further, note that in the unattested versions, there is a mismatch between the answer that the asker is weakly committed to on one hand, and the answer that the a *yes* response would pick out on the other. After (287), if the addressee subsequently responds with *yes*, it confirms the notion that Bates is *not* at fault, which conforms to the bias that the asker expresses. After (288), however, a *yes* response confirms that Bates *is* at fault, even though the asker still presents bias that Bates is not at fault. Farkas and Roelofsen (2012; F&R) discuss

these sorts of *reversing responses*, where the answer favored by the asker differs from the response provided. For F&R, the favored answer, or prejacent, is *highlighted*. This notion, from the Inquisitive Semantics literature, corresponds to the idea that a single alternative can receive extra semantic and pragmatic prominence beyond the other alternatives that comprise an issue (if any) by virtue of how the issue is raised. The discussion here does not make use of highlighting, so I instead suggest that one source of semantic and pragmatic prominence is bias marking, i.e., weak commitment.

Because they go against the expectations of the asker, reversing responses are highly marked. Assuming a parallel between semantic and pragmatic markedness on one hand, and morphological markedness on the other, support for the special status of reversing responses comes from that fact that many languages feature a third polarity particle specifically to mark reversing (Pope, 1976; Farkas, 2011). Particles that specifically mark a reversing move include *ba* in Romanian and *de* in Hungarian. French *si* and German *doch* also qualify as special reverse-marking particles, although they are only used for a subset of reversing moves, in cases where the reversing move is toward the positive answer, like in (288) above.

The markedness of reversing moves leads to a strong, cross-linguistic preference for asking questions that will be answered affirmatively. F&R offer the following summary, phrased in terms of a minor conversational maxim:

(289) AVOID [REVERSE]
Other things being equal, a cooperative speaker formulates her initiative in such a way as to minimize the chance of eliciting a [REVERSE] response.

In the case of biased polar questions, the attested rising and falling TQs accomplish the same communicative goal as the unattested 'mirror' forms, but do so while also minimizing the likelihood of a reversing response. There is therefore no reason for the mirror forms' arrangements of commitments to be employed.

4.7.2.3 Superfluous commitments

There is another 'mirror' form among the configurations as well, this time corresponding to the HNPQ:

(290) ACTUAL AND ALTERNATIVE CONFIGURATIONS FOR HNPQ

Question form	$(p \vee \neg p)$	p	$\neg p$
h. HNPQ	current	prior	current
f. HNPQ′	current	current	prior

It is tempting to apply the above reasoning to the HNPQ/HNPQ′ split as well, but this approach leads to problems. This is because, due to the conflicted nature of HNPQs and the fact that semantic and pragmatic prominence are tied to weak commitment under this view, *both* answers to a HNPQ are in some sense 'expected'.[29]

(291) *Ness, who Paula knows was previously out exploring the wilderness with Jeff, returns seemingly alone.*

Paula: Isn't Jeff with you?

p = Jeff is with you

Conveys: Paula thought p but evidence suggests $\neg p$; *yes* = p

Ness: a. Yes... *desired answer, but strictly less likely*

b. No... *more likely based on current context*

It seems that the clearest way to ensure a *yes* response, given the speaker's current commitments, is to embrace the possibility that Ness is alone and phrase the question as a HNPQ′. Instead, the *yes* answer is tied to the response that is strictly less likely given the Paula's default base. In this way, Paula expresses doubt that the new contextual evidence that has been foisted upon her is reliable.

In more detail, Paula signals prior bias for Jeff's presence as a way of indicating that she prefers a resolution where she has maintained her beliefs, rather than one where she is forced to integrate the new contextual information. In other words, Paula ties the *yes* answer to what she considers the better resolution (returning to a default base more like her prior base), rather than just the most likely proposition (which would saddle her with an interpretation of contextual information she does not trust). This discrepancy relates to

[29]If F&R's notion of highlighting were incorporated, this difficulty might evaporate, because then there would again be a clear notion that one of the alternatives is more pragmatically prominent than the other. Note also that if we wish to capture the inner/outer-HNPQ ambiguity, highlighting therefore offers a way forward in that regard as well.

what Ladusaw (2003) terms the difference between favoring a "cell in the partition" (i.e., the likelihood of an answer), and favoring a "resolution" of the issue (i.e., the quality of the context once the answer is integrated). In the current model, this amounts to the difference between simply contrasting the odds of two propositions, and comparing the broader effects that accepting either proposition as true would have on the discourse in terms of how it affects the evidence that the speakers use to underwrite further commitments.

This distinction mirrors the division between epistemic modality and true illocutionary modification that runs throughout this dissertation. Commitment—including the weakened form of it that encodes bias—is ultimately concerned with the legitimacy of the evidence underwriting the commitment. Through this focus on evidence, agreements to act as though propositions are true can be reliably established. Epistemic modality, on the other hand, is concerned foremost with the truth of a proposition.

That said, it is still unclear why the HNPQ, clinging to an unlikely outcome, should trump the maxim to avoid reversing moves. An example of what this form might look like is provided below, and paves the way to a solution.

(292) *Jeff told Paula he was staying home, but Ness has just bought 3 backpacks for their journey instead of 2.*

a. Paula: Isn't Jeff not coming with us? *(negative) HNPQ*
   p = Jeff is not coming with us
   Conveys: Paula thought p but evidence suggests $\neg p$; *yes* $= p$ (mismatch?)

b. Paula: Is Jeff coming with us [my inkling was *no* but now it's *yes*]? *'equivalent' HNPQ′*
   p = Jeff is coming with us
   Conveys: Paula thought $\neg p$ but evidence suggests p; *yes* $= p$

To solve this puzzle, we need only remember that beyond any calculations the speaker makes in determining whether to publicize bias, the 'purpose' of a polar question in terms of its basic discourse effect is to work toward resolving whether p. A speaker's current default base an obvious role in this, but as discussed previously, prior evidence and any bias it underwrites bears on the issue of whether p only indirectly; it is only relevant to the extent that it aids or hinder future moves by the interlocutors. This is in part why

commitments given prior evidence are a marked move; that prior evidence is called up because it also bears on the speaker's willingness to commit given current evidence.

In the unattested (4.7.2.3b), the desired *yes* answer refers to the same proposition that Paula is weakly committed to. Resolving in favor of p is therefore unproblematic, and the prior $\neg p$ bias is rendered completely superfluous. In (4.7.2.3a), by contrast, the prior commitment is in fact made relevant *because* it is positioned as preferred answer.

It seems that in conversation we take for granted that our interlocutors are aiming for an answer that matches the polarity of the question. This means that the act of creating a mismatch between the preferred answer and the most likely one is what makes the prior commitment important–unlike in all the other cases discussed, where it ends up having no effect on the question of whether p. By producing this mismatch, the speaker prompts her interlocutors to contrast her previous and current evidential bases, because it bears on the likelihood of the addressee's eventual responding move being accepted. The discourse effects of HNPQs are therefore motivate by self-interest on the part of the addressee.

4.7.2.4 More superfluous commitments

A few other potential forms run into the same issue of having superfluous commitments based on the previous state of the default base. I mention just two of the more interesting cases here.

(293) MORE UNATTESTED CONFIGURATIONS

Question form	$[p \vee \neg p]$	p	$\neg p$
g.	current	prior	-
q.	prior	prior	current

the (g) configuration, if licit, would convey the idea, "Previously I thought p, but now I'm unable to commit (even contingently) to either answer." It is, therefore, a HNPQ, without the conflict. But just like the configuration in the previous section, the prior commitment is superfluous because given that the HNPQ was not used, we assume that the current base leaves the speaker unbiased. The prior commitment to p plays no role in how the question of whether p will be settled.

Configuration (q) is arguable quite useful. It conveys, "Previously I thought p, but I was open to either. Now, I think $\neg p$." And in fact, we do convey such biases rather

frequently—using a rising TQ. We can create the conditions for a question like this by taking any environment that would lead to a licit rising TQ and advancing the context toward $\neg p$:

(294) *Zelda knows that it's Saturday morning, so Link is likely on a quest. She turns toward Sara, who's seen Link more recently, and asks:*

Z: Link's on a quest, isn't he?

(295) *Zelda knows that it's Saturday morning, so Link is likely on a quest. She turns to Sara and asks:*

Z: Li–

S: *(Interrupting)* Look, Link's horse is outside the tavern!

Z: a. (So) Link's not on a quest, is he.

b. * Link's on a quest [(q)-particle] ?

The negated tag question fulfills the same function at the (q) configuration, without reference to the superfluous, prior commitment to p.

4.7.3 Conclusions on variation

After all of the discussion above, it appears that the only licit bias configurations are the four below:

(296) PREDICTED EVIDENCE CONFIGURATIONS FOR POLAR QUESTIONS

Question form	$[p \vee \neg p]$	p	$\neg p$
a. Default Q	current	-	-
h. HNPQ	current	prior	current
d. Rising TQ	current	current	-
m. Falling TQ	prior	current	-

These are, in fact, the four configurations of bias that are attested in the data presented in this chapter. While allowing commitments based on an earlier state of the context seems as though it would over-generate the possible space of variation for biased questions, in fact the results conform to expectations. Those configurations that are not attested are ruled out by robust generalizations about the pragmatics of polar questions.

4.8 Conclusions

This chapter has presented a view of bias in English polar questions as defective discourse evidentiality. These questions share a core semantics and pragmatics, variously augmented by additional commitments to the alternative propositions p and/or $\neg p$. Bias is expressed via commitments conditioned on the weak base, and are therefore a kind of discourse-evidential marking. Weak commitments further be broken down into those made relative to the default base, which is a combination of a speaker's private speaker beliefs and shared, public contextual evidence, and the prior base, which is equivalent to the default base at an earlier point in the discourse. In terms of data, some longstanding puzzles concerning high negation polar questions and tag questions were addressed along the way.

At this point, I have offered discourse-evidential analyses of a variety of phenomena, including sourcehood (Chapter 2), relative authority (Chapter 3), and now bias (Chapter 4). What conclusions can be drawn from this collection? In the next chapter, I conclude my investigation by arguing that the purpose of marking discourse-evidential notions is to facilitate a speaker's ability to coordinate his commitments and settle issues in the discourse.

Chapter 5

Conclusion

5.1 Review of the analysis

The central question I have pursued so far in this dissertation is: How do commitments vary? To that end, I have proposed a novel discourse model that draws formal connections between the domain of illocutionary evidentiality and an array of phenomena including sourcehood marking (Chapter 2), relative authority marking (Chapter 3), and expressions of bias in polar questions (Chapter 4). In this section I briefly summarize the arguments that lead to the particular discourse model I've proposed. I then discuss the broader implications of the approach and how it improves our understanding of the pragmatics of conversation. In short, the goal here is to step back and shed light on *why* commitments vary in the ways that they do. Below, I argue that what all of the phenomena mentioned above have in common with each other (and with illocutionary evidentiality) is that they facilitate *management of commitments* and *issue-raising*—the most important goals that speakers face under a commitment-based view of discourse.

The essential theoretical machinery introduced in Chapter 2 is an elaborated notion of what a discourse commitment is. Authors of previous commitment-based discourse models note that the Stalnakerian common ground is insufficient for modeling the discourse effects of interlocutors' utterances in a satisfactory way. For example, Farkas and Bruce (2010) argues that the phenomenon of "agree[ing] to disagree" is strong evidence that speakers make real distinctions between the contributions that they've introduced to the discourse, and those that others have. This reflects a more general conceptual problem with relying

too directly on a Stalnakerian common ground to model how speakers converse, because speakers cannot in fact build a shared repository of knowledge directly. Rather, what is manipulated in discourse must be an inherently dialogic structure, which *reflects* the state of the common ground, but *models* each interlocutor's state independently.

Gunlogson (2003) calls each interlocutor's component of the discourse structure their *discourse commitments* (*DC*s). A speaker's discourse commitments is a set of those propositions that the speaker has publicly agreed to act as though he believes.[1] This separation allows for easy comparison between what two different speakers are committed to. Where their *DC*s overlap, there is common ground. Where the *DC*s do not overlap, there are meaningful differences in how they are expected to act in discourse. Note that employing these commitment structures does not replace Stalnaker's common ground; rather, it enriches it. The elaboration of commitment proposed in this dissertation is a similar—and similarly necessary—elaboration.

Many authors have noted that not all commitments that a speaker makes are the same. The model presented here therefore extends the idea of commitment by linking the propositional content of every commitment to an *evidential base*, on analogy with Kratzer's (1981) modal base.

(297) ANATOMY OF A COMMITMENT $\langle p, E \rangle$

a. Propositional content p, restricts worlds that Sp may act as though he considers viable[2]

b. Evidential base E, characterizes the evidence that licensed Sp to commit to p

The evidential base is a set of propositions that records the details about the reasons that allow the speaker to make his commitment. characterizations of this base can invoke different evidence sources, requirements on speaker's relative authority over the evidence, or even some information about how fragile the commitment that invokes the base will be in light of changes or challenges to the base. These elaborations are reviewed in more detail

[1] Gunlogson suggests they might also be called *public beliefs*, although I avoid this term; it tempts one to forget that holding a public belief that p does not entail believing that p.

[2] Recall from Ch 3 that in the case of commitments to truth, p-worlds are candidates for the actual world, i.e., worlds where p is true. In the case of commitments to action, a speaker committed to p should act in order to to bring p about (meaning that it needn't be true as utterance time).

below. The evidential base allows a great deal of flexibility in the information that a commitment conveys. The limiting factor on this variation is that any distinction in commitment that is reflected in the discourse structure must have predictable effects on speakers' ability to manage their own commitments and get their interlocutors to make useful commitments of their own. This crucial component of discourse behavior is what is captured better under the current discourse model than under previous ones.

5.1.1 Sourcehood

The evidential base that underwrites a commitment might be characterized as involving evidence from a particular sense or inference process, in which case the base is a straightforward encoding of illocutionary evidentiality, in the vein of Faller (2002) or Aikhenvald (2004) (but c.f. Murray 2010). Crucial for current purposes, the base can also explain how interlocutors' commitments hang together. As discussed in Chapter 2, most commitments 'stand alone'. The speaker commits, based on a combination of his private speaker beliefs and any shared, contextual evidence, but crucially, whether the speaker's interlocutors commit as well or maintain prior commitments will not affect the speaker's commitment. Because the speaker's commitment is independent of the other interlocutors' discourse moves, the speaker is said to be the *source* for his commitment (Gunlogson, 2008). In contrast to this, speakers can make discourse moves where they commit to a proposition on the basis of an interlocutor's prior commitments, or the expectation of future ones. This kind of *dependent* commitment is a pragmatically marked move, because it jeopardizes the speaker's ability to maintain control of his discourse commitment set. This is especially clear under the discourse model proposed here: A dependent commitment is underwritten by the fact that the addressee is committed to the same proposition. If the addressee reneges, the basis for the dependent commitment falls away too. This puts a speaker who makes a dependent commitment in a precarious position; he not only relies on his interlocutor to vouch for that specific commitment, but in doing so, he gives up a measure of control over his ability to keep his commitment set consistent. It is therefore unsurprising that, as Farkas and Roelofsen (2012) remark, pragmatic markedness of this sort is reflected in formal markedness (e.g., syntactic, morphological, or intonational). It is important that all of the participants in the discourse are aware of the potential commitment instability that

the speaker is introducing.

5.1.2 Relative authority

More generally, the major function of the evidential base behind a commitment is to highlight a particular *trigger of mutability* for that commitment. This could mean highlighting potential dangers for the discourse, as detailed above, but it can just as easily involve providing information that is useful in the process of negotiating new commitments. Relative authority markers in Japanese, as discussed in Chapter 3, function in this capacity. Recall the definitions for the final particles *ne* and *yo*, from Chapter 3, where $AUTH_X(\varphi)$ stands for the measure of authority that discourse participant X has for φ, given X's default base.

(135) $[\![ne(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base $E_{\text{MIN}} = \{q | \forall X \in D : AUTH_X(q) \geq AUTH_{Sp}(q)\}$.

(133) $[\![yo(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base E_{MAX} s.t.: $E_{\text{MAX}} = \{q | \forall X \in D : AUTH_X(q) \leq AUTH_{Sp}(q)\}$.

Under this analysis, marking an utterance with *ne* indicates that the speaker regards her evidence as shared. Regardless of any other facts about the commitment, such as whether it is strong or weak (see below), the speaker's interlocutors are expected to have at least equal access to the base that underwrites it, meaning that the commitment itself ought to be uncontroversial as well. This facilitates the adoption of the commitment by the other interlocutors, because the speaker is telling them that they ought to be able to commit independently. In a similar vein, using *ne* indicates that the speaker does not anticipate a battle over how to settle the issue the utterance raises.

Marking an utterance with *yo* also aids the progression of the discourse, even though it highlights disparities in evidence rather than commonalities. When a speaker uses *yo*, she indicates that her interlocutors have no better evidence than she does, and by implicature,

that her evidence is in fact the best.[3] Again, this facilitates commitment for the speaker's interlocutors; if their evidence for the proposition in question pales in comparison to the speaker's, their best recourse will often be to take the speaker's word for it and commit dependently. In terms of settling issues, when the speaker invokes *yo*'s evidential base, she seeks to forestall argument. If the speaker really is the best authority, there is little reason to question her designs for the discourse.

In sum, the evidential bases brought to bear by *yo* and *ne* do more than merely indicate the speaker's presuppositions about his addressee's access to evidence. Rather, they provide a backchannel to ease negotiation over commitments and issues. The fact that it is this general function, rather than relative authority itself that is most important is made especially clear when we consider the distinction between *ne* above and the Singlish particle *lah*, which I have argued should receive the following analysis:

(178) $[\![lah(\varphi)]\!] =$

a. *At-issue:* φ

b. *Not-at-issue:* Any commitment to φ is conditioned on a base E^{LAH}, and $\exists q \in E_{\text{LAH}}$ s.t.:

   i. The speaker could not commit to φ given a base $E^{\text{LAH}'}$ where $E_{\text{LAH}'} = E_{\text{LAH}} - q$, and

   ii. The speaker and addressee are taken to be committed to q

Unlike *ne*, which treats the entire evidential base as 'shared' information, *lah* indicates that one crucial proposition as shared, thereby calling attention to it. By specifically publicizing that there is common ground between the speaker and addressee, using *lah* performs exactly the sort of function the present model predicts should be available for grammaticalization.

5.1.3 Bias

The evidential base is a tool for indicating specific ways that a commitment is open to influence in later discourse. For example, invoking the dependent evidential base highlights that the instability of the speaker's commitment is connected to a very specific trigger:

[3] This implicature arises from competition with the possibility of using both *yo* and *ne*.

the addressee's commitments. Invoking bases related to relative authority using *yo* or *ne* similarly highlights the speaker's expectations about how the addressee will be able to interact with her evidence. The *weak base*, by contrast, is a more general tool that speakers have for publicizing potential instability in their discourse commitments that originates from their own uncertainty. A strong commitment, the default state, indicates that the speaker does not expect that the commitment will change. In other words, the speaker declares that, independent of any specific triggers of mutability, the evidential base as a whole is one that he trusts. These notions may seem at odds, but they are not. Consider again the effect of *oh* in the following dialogue based on Gunlogson (2003, 2008):

(15) A: The server's down.

B: Oh, (we should call someone.)

In this example, when B responds with *oh*, he indicates that he is dependent on his addressee, but he is also taken to be fully committed to the proposition. His commitment has a trigger of mutability in the form of A's commitment to the server being down, but barring that, the commitment is reliable.

Weak commitment, on the other hand, takes for granted that some unspecified future evidence will come available to bolster the commitment. A weak evidential base contains a promissory note to this effect, meaning that the commitment will not be maintainable unless such evidence is found. The commitment is therefore volatile; rather than identifying a single, identified point of mutability like the speaker's commitment, invoking the weak base indicates that the continued viability of the commitment is in some sense at the whims of the future discourse, because the speaker's own evidence is too tenuous to stand alone.

Weak commitments are employed in non-default polar questions in order to indicate that the speaker does not regard both options equally. The base underwriting a weak commitment to one of the question's answers constitutes mounting evidence in favor of that answer. Weak commitment is therefore a roadblock for issue-settling, because it indicates that the addressee potentially faces an uphill battle if she wants to push for an answer counter the asker's weak commitment. Tag questions, which involve a single dependent commitment, exemplify this.

High negation polar questions (HNPQs), in contrast, feature incompatible weak commitments to *both* possible answers, one conditioned on the asker's default base, and one

conditioned on a prior version of the base that lacks recent contextual evidence. The specifics of the roadblock a weak commitment puts up depend on the evidential base that underwrites it, analogous to how the base changes points of mutability for strong commitment. This means that in order to overcome the bias expressed by a HNPQ, the speaker must address the anomalous contextual evidence head-on. This explains why (215c) is the best answer to this kind of question.

(215) *Gertrude notices one fewer place settings than expected, given that Teddy was supposed to be visiting tonight, and says to Maxine:*

G: Isn't Teddy coming over?

M: a. ? Yes, (he is). *No target*

b. OK Yes, but he won't get here until after dinner. *Targets CE*

c. # Yes, we were both there when he accepted the invitation. *Targets SB*

Crucially, the minimal answer in (215a) is odd because it seemingly ignores the chance to incorporate a strategy based on the backchannel negotiations that Gertrude has begun by asking a biased question.

In terms of managing commitments, being able to express a conflict in bias this way is extremely useful for a conflicted speaker who wishes to keep his commitments consistent. By framing bias as a roadblock for the progression of the discourse, the asker forces the addressee to help resolve what is otherwise simply an internal conflict of the asker's. Resolving the conflict benefits all the interlocutors.

5.1.4 Common thread: discourse motives

The summary above argues that the phenomena of sourcehood marking, relative authority, and bias must all make reference to negotiations over evidence, but why should this information be represented *as part of the discourse structure*? My response is purely pragmatic, in the popular sense. Time and again, we have seen that performing discourse-evidential utterances does not simply involve airing secondary meta-beliefs alongside the primary ones that drive the discourse forward. Rather, choosing to mark a discourse-evidential feature like relative authority has immediate and uncancelable effects on the felicity of current and

future utterances. For example, recall from Chapter 2 that when a speaker uses a rising declarative, it signals a dependent commitment to the utterance's content:

(20) *Max points at something at fruit stand where Laura often shops.*

M: That's a persimmon?

L: Yes, (it is)./#Oh.

When Max commits as dependent, he indicates that his commitment is underwritten by a (still pending) commitment of Laura's. If Laura responds *oh*, claiming dependence herself, the effect is not to highlight that Max is wrong to have expected an answer from her. Rather, her response is simply not felicitous. The discourse-evidential properties of an utterance are not secondary effects; they change the trajectory of the discourse, and so those effects ought to be represented directly in the discourse model, rather than relegated to informal pragmatic reasoning behind the formal pragmatics.

In Chapter 1, the central argument for formalizing discourse in terms of commitments was that even if commitments are not an explicit part of the discourse model, it is still necessary to emulate them in order to arrive at a satisfactory view of discourse effects. In the body of this dissertation, I have argued that the same reasoning applies to evidential bases: evidential bases are a driving component of discourse effects, even if they aren't explicitly formalized—so it's best to just formalize them.

5.1.5 Issues for the approach

This section wraps up the review by investigating two remaining issues that arise from connecting the various phenomena I have discussed with evidentiality: the lack of 'interrogative flip' with discourse-evidential markers, and the relationship between evidentiality and epistemic modality. I argue that in light of the broad conclusions above, neither of these issues is problematic. Rather, both can be explained as natural results of an approach where discourse-evidentiality is viewed as a strategy for managing commitments and facilitating issue-settling.

5.1.5.1 Interrogative flip

In many languages, standard evidentials in questions optionally participate in *interrogative flip*, where the evidential component can variably specify either the reasons behind the asker's question, or the kind of evidence the asker would like the addressee to base her answer on. Faller (2007) reports that the Cuzco Quechua reportative evidential can be used both ways:

(298) May-manta-s chay runa ka-n-man.
where-ABL-REP this man be-3-COND
'Where could this man be from?' (Faller, 2006)

(299) Imayna-s ka-sha-nki.
how-REP be-PROG-2
'(She says) How are you?' (Faller, 2007)

In (298), the reportative *-si* (*-s* here) requests that the addressee answer with reportative evidence, but it says nothing special about the grounds for the asker's question. In (299), the reportative evidentiality is attributed to the asker, who is asking a question on behalf of someone else. Note that this particular question is not ambiguous, because it would be pragmatically odd to request that the addressee use reportative evidence to answer the question, "How are you?"[4]

In light of this, a potential objection to the enterprise of discourse-evidentiality is that none of the discourse markers discussed in this dissertation participate in interrogative flip. For example, recall that a rising tag question under the current view weakly commits the speaker to the anchor possibility. There is no reading, however, whereby the *addressee* is asked to make a weak commitment to one of the answers:

[4]This suggests that under some circumstances, the characterization of an evidential base introduced by an expression of evidentiality can associate with *future* commitments, in particular, those that result not from the speaker's utterance, but from later discourse moves that resolve the issue that the speaker's utterance puts on the Table. In (298), this would mean that a reportative base is somehow associated with the set of possible answers to the speaker's question. Whether this means that the discourse model really needs to keep track of future states of the discourse is a Pandora's Box which I choose not to open. See Malamud and Stephenson (2011) and Krifka (2011, 2012) for proposals that move in this direction. Given the idiosyncratic nature of interrogative flip (as discussed below), I am not convinced such radical measures are warranted by these data alone.

(300) M: Teddy's coming to dinner, isn't he?
≠ 'Do you have an inkling that Teddy's coming to dinner, (or do you have an inkling that he doesn't)?'

Note that the addressee can of course respond to a tag question like this with a weakened commitment. The argument here is that unlike the reportative, the question does not anticipate or request a weak commitment.

Before retracting the whole dissertation because of this problem, though, I offer two related arguments for why this discrepancy is natural in light of the previous discussion. First, it is important to note that interrogative flip is itself a highly idiosyncratic phenomenon. Not all illocutionary evidentials participate, and even within the same language, different evidentials can show different patterns.[5] Returning to Cuzco Quechua, Faller (2007) claims that although the reportative optionally flips in questions, the conjectural *chá always* does. In other languages, however, a given evidential marker might *never* flip. This is the situation in Yukaghir, as reported by Aikhenvald (2004), citing Maslova (2003):

(301) Godo ti:-t kebej-nu-l'el-ŋi.
how here-ABL go-IMP-NONFIRSTHAND-3-PL.INTR
'How do people go away from here?' (the speaker did not see them)

In this example, the non-firsthand marker indicates the speaker's basis for asking the question. It cannot function as a request for the addressee to respond with non-firsthand evidence—although again, such a response is not ruled out.

Given how variable interrogative flip is in general, the weak answer to the criticism that discourse evidentiality seems to lack the property is that this conclusion is simply a hasty generalization; the fact that the examples I've discussed do not flip is consisted with our current understanding of interrogative flip as an idiosyncratic phenomenon. This leaves open the possibility that in the future, discourse evidentials that do flip could be discovered, just like any future discovery of a more traditional evidential could either flip or not.

[5]Further, the inventory of evidentials that a language uses needn't be consistent across clause types; some evidentials are disallowed in questions altogether (Aikhenvald, 2004, p. 242–9). For example, Chapter 2 mentioned that Tariana (Arawak) features a 5-way contrast in evidential markers, but the full set is only available with declarative sentences. With interrogatives, only three options are licit, and with imperatives, only two. This discrepancy cannot be traced back to the character of the evidentials themselves, as evidenced by the fact that a restricted evidential can have an analog in another language that is not restricted.

This is an unsatisfying conclusion, however. There is a sense in which it does not seems that most discourse evidentials *should* be able to flip. At the very least, there ought to be some way to predict whether a given evidential will tend to flip or not. A solution along these lines is within grasp, if we consider again the generalizations in the previous discussion about the utility of discourse evidentiality.

This section has argued that discourse evidentiality is marked in order to facilitate commitment management and issue-settling, by providing useful, backgrounded context for interlocutors' negotiations. But why are these negotiations so critical? The terse answer—paraphrasing from Chapter 1—is that this is how speakers maintain and grow their common ground. Any discourse moves that require a speaker to revise his commitments put him in jeopardy of undoing the common ground status of some information. Maintaining good commitments, then, goes beyond a simple requirement that one's commitments are consistent. Rather, the commitments should be as strong and well-supported as possible. Speakers, thanks to their inherent fallibility as humans, generally have good evidence for only a subset of the things they privately believe or expect. Maintaining ones' commitments therefore is an exercise in *avoiding overcommitment*, because commitments on tenuous grounds are more likely to require revision. This amounts to a restatement of Grice's Quality maxim in terms of commitment and evidence.

Opposing the idea that speakers should minimize their own commitments is the fact that the common ground can only grow if one's interlocutors participate in settling issues and make commitments of their own. It is therefore in a speaker's self-interest to be proactive and put his interlocutors into positions where they must make commitments. This evokes Grice's Quantity maxim, albeit with the focus shifted onto directing the actions of others, rather than shaping ones own discourse moves. These two requirements can themselves be phrased as maxims that guide a speaker's actions relative to a commitment and evidence-based discourse as follows:

(302) MAXIMS OF COMMITMENT NEGOTIATION

a. Do not overcommit.

b. Require your addressees to commit maximally.

A speaker who follows these maxims is acting cooperatively, because he will build the most reliable discourse commitment set that he can, while also pushing his interlocutors to com-

mit more liberally. These forces pull against one another, leading controlled growth of the common ground.

Returning to the issue of interrogative flip, it is immediately clear under this view why flip is sometimes grammaticalized for standard evidentials but never for discourse evidentials. If a speaker wishes to urge his addressee to commit maximally, it can be highly profitable to request, e.g., a commitment explicitly based on hearsay, because it gives the addressee license to make a less rigorous commitment without penalty. In other cases, the fact that the answer is hearsay might also be integral to the direction of the conversation; if A and B are discussing office rumors, requesting reportative, conjectural, or even specifically direct-visual evidence could be highly relevant to the issues at hand. Discourse-evidential notions, however, will never as useful in growing the common ground. There is little utility in specifically requesting a weak (i.e., highly mutable) commitment using a 'flipped' tag question; the expected answers, conditioned on the weak base, would not be able to stand on their own without further evidence, which the question asker would likely not be able to provide. There is even less utility in seeking out a dependent commitment; the addressee's answer, which would invoke the dependent base, would crucially contain the fact of the asker's own (independent) commitment, and if the asker could offer such a commitment, where would have been no need to ask in the first place. The canonical result of requesting that a question be answered dependently would therefore result in a discourse structure on the verge of partial collapse. Grammaticalizing interrogative flip for a given evidential is only available when requesting a modified commitment of that sort does not conflict with the goals in (302). This is a claim which should be tested in future research.

5.1.5.2 Evidentiality vs epistemic modality

Finally, I conclude this review of the analysis with perhaps the biggest open question: What is the nature of the relationship between evidentiality and epistemic modality? This dissertation has taken the view that there is a difference between these categories that must be respected, although the nature of that difference remains mostly intuitive. I have argued, for example, that making a strong commitment to a modalized proposition (e.g., to assert *might-p*) is a distinct discourse move from making a weak commitment to an unmodalized proposition (e.g., to show bias for *p*):

(303) *Lucca is investigating Marle's pendant, a family heirloom of mysterious origin.*

L: It's hard to make out, but the inscription seems to say "Schala, Princess of the Kingdom of Zeal."

M: a. (So) this might be Schala's pendant. (I have no idea.)

b. (So) this is Schala's pendant? (# I have no idea.)

Marle's possible utterances are not entirely parallel, given that the rising declarative employed in (303b) involves dependent commitment while (303a) does not. Even so, the relationship that Marle bears toward the proposition that the pendant is Schala's is different in each case. In (303a), Marle is most naturally remarking on a logical possibility, based on what it means for a piece of jewelry to be engraved, without committing to that interpretation. In (303b), Marle does seem to be committed, however weakly, as shown by the infelicity of the continuation above.

The intuition that these sorts of discourse moves are distinct remains controversial; Krifka (2004) remarks, "to assert an epistemically weakened proposition is conversationally equivalent to performing a downtoned assertion." At the very least, it is clear that utterances like (303a–b) often implicate one another, which makes separating their immediate discourse effects and their latter pragmatic extensions a difficult task.

As a preliminary suggestion based on the analysis presented here, I would argue that the narrow purposes of illocutionary evidentiality and epistemic modality are distinct. Recall from Chapter 2 that evidentiality is concerned primarily with evidence source, and that the strength of a commitment is a secondary notion (see Aikhenvald 2004). Either way, however, modulating a commitment involves adjusting one's public accountability for the truth of the commitment's content, as per the 'maxims' introduced in (302). Modulating epistemic modality, on the other hand, is more narrowly focused on the 'search for truth'.[6] A fuller exploration of the difference between the two should begin with contexts where these goals diverge, but I will refrain from committing further.

[6]This description evokes Verum focus. See Chapter 4.

5.2 Future directions

To conclude, this section presents a few major avenues that the project of this dissertation points to for future work.

First, there is still more to be said about what the proper relationship is between evidentiality and epistemic modality. By introducing the new category of discourse evidentiality as both a subcategory and cousin of more standard illocutionary evidentiality, I have narrowed an empirical and theoretical divide that already resembled the Strait of Messina. More work is needed on the question of how discourse-evidentiality and epistemic modality interact.

Second, I follow the majority of authors who discuss the idea of evidence without every moving beyond an intuitive notion of what evidence is, or specifically, what makes a particular proposition 'count' as evidence in the first place. This project is off to a promising start thanks to McCready (2011), and a complete answer will likely involve a long period of collaboration between linguists and cognitive scientists. Even preliminary results from these investigations will have major repercussions for any theory like this one, that skips over the question of where evidence comes from.

Finally, in a related point, the model proposed here has used a single mechanism, the evidential base, for what some might argue are *two* points of variation for commitments: the content of the evidence itself (Harnish's "reasons"), and the 'strength' of that commitment. The latter of these is captured here via the 'weak' base, invoked to indicate that the speaker's own evidence is insufficient. One natural question that arises is whether the strength of a commitment is unpredictable, or if it can be more-or-less determined based on the relationship between the commitment's propositional content and the other factors marked by the evidential base. It is clear from the cross-linguistic differences in the behavior of, e.g., hearsay evidentials, that there must be *some* relationship between evidence and strength, but where it can be regularized depends heavily on what evidence is taken to be in the first place, as discussed above. Moving to a model that features deterministic commitment strength is a natural progression for a model like this, which already evokes the idea that commitment is 'conditioned' on evidence, in the Bayesian sense. As linguistic pragmatics becomes increasing computational, it is inevitable that a speaker's general reasoning processes about probability and the quality of evidence will to leak into our representations of discourse. In pursuing this avenue of research we must take care not the lose the insights

that formal pragmatics has provided about the differences among commitment, belief, and expectation.

It is my hope that the ideas presented in this dissertation will find their way into the evidential bases underwriting the claims of others in the field. The reader should keep in mind, however, that the commitments I've made in this dissertation are themselves mutable in light of future evidence.

Bibliography

Aikhenvald, Alexandra Yu. 2004. *Evidentiality*. Oxford University Press.

Anand, Pranav, and Valentine Hacquard. 2013. Epistemics and attitudes. *Semantics and Pragmatics* 6:1–59.

Anderbois, Scott. 2011. Issues and alternatives. Doctoral Dissertation, UC Santa Cruz.

AnderBois, Scott, Adrian Brasoveanu, and Robert Henderson. 2010. Crossing the appositive / at-issue meaning boundary. In *Proceedings of SALT 20*, ed. Nan Li and Daniel Lutz, 328–346.

Asher, Nicholas, and Alex Lascarides. 2003. *Logics of conversation*. Studies in Natural Language Processing. Cambridge University Press.

Austin, J. L. 1962. *How to do things with words*. Clarendon Press.

Bell, R.T., and L.P.Q. Ser. 1993. 'Today *la*?' 'Tomorrow *la!*': The *la* particle in Singapore English. *RELC Journal* 14:1–18.

Besemeres, Mary, and Anna Wierzbicka. 2003. The meaning of the particle *lah* in Singapore English. *Pragmatics and Cognition* 11:3–38.

Beyssade, Claire, and Jean-Marie Marandin. 2006. The speech act assignment problem revisited: Disentangling speaker's commitment from speaker's call on addressee. In *Empirical Issues in Syntax and Semantics 8*, ed. O. Bonami and P. Cabredo Hofherr, 37–68.

Bolinger, Dwight Le Merton. 1989. *Intonation and its uses*. Stanford University Press.

Büring, Daniel, and Christine Gunlogson. 2000. Aren't positive and negative polar questions the same? UCSC/UCLA.

Caponigro, Ivano, and Jon Sprouse. 2007. Rhetorical questions as questions. In *Proceedings of Sinn und Bedeutung 11*, ed. E. Puig-Waldmüller.

Carlson, Lauri. 1983. *Dialogue games*. Reidel Dordrecht, The Netherlands.

Chernilovskaya, Anna, Cleo Condoravdi, and Sven Lauer. to appear. On the discourse effects of *wh*-exclamatives. In *Proceedings of the 30th West Coast Conference on Formal Linguistics (WCCFL 30)*, ed. Nathan Arnett and Ryan Bennett. Cascadilla Press.

Ciardelli, Ivano, and Floris Roelofsen. 2011. Inquisitive logic. *Journal of Philosophical Logic* 40:55–94.

Condoravdi, Cleo, and Sven Lauer. 2010. Imperatives and public commitments. Handout of talk presented at the 11th Annual Semantics Fest, Stanford University.

Condoravdi, Cleo, and Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. In *Empirical Issues in Syntax and Semantics 8*.

Cook, Haruko Minegishi. 1990. The sentence-final particle *Ne* as a tool for cooperation in Japanese conversation. In *Japanese/Korean linguistics*, ed. Hajime Hoji, 29–43. CSLI.

Davis, Christopher. 2009. Decisions, dynamics and the Japanese particle *yo*. *Journal of Semantics* 26:329–366.

Davis, Christopher. 2011. Constraining interpretation: Sentence final particles in Japanese. Doctoral Dissertation, University of Massachusetts, Amherst.

Davis, Christopher, Christopher Potts, and Peggy Speas. 2007. The pragmatic values of evidential sentences. In *Proceedings of SALT 17*, ed. M. Gibson and T. Friedman.

De Brabanter, Philippe, and Patrick Dendale. 2008. Commitment: The term and the notions. *Belgian Journal of Linguistics* 22:1–14.

Deterding, David. 2007. *Singapore English*. Dialects of English. Edinburgh University Press.

Ettinger, Allyson, and Sophia Malamud. 2013. Mandarin utterance-final particle *ba* in the conversational scoreboard. In *Proceedings of the Linguistic Society of America Annual Meeting*.

Faller, Martina T. 2002. Semantics and pragmatics of evidentials in Cuzco Quechua. Doctoral Dissertation, Stanford University.

Faller, Martina T. 2006. Evidentiality below and above speech acts. In *Functions of language*. John Benjamins Publishing Company.

Faller, Martina T. 2007. The Cuzco Quechua conjectural: an illocutionary analysis of an epistemic evidential. In *Proceedings of SULA 4*, ed. Amy Rose Deal. Amherst, MA: GLSA Publication.

Faller, Martina T. 2011. A possible worlds semantics for Cuzco Quechua evidentials. In *Proceedings of SALT 20*, ed. Nan Li and Daniel Lutz, 660–683.

Faller, Martina T. 2012. Evidential scalar implicatures. *Linguistics and Philosophy* 35:285–312.

Farkas, Donka. 2010. The grammar of polarity particles in Romanian. In *Interface properties: Edges, heads, and projections*, ed. Anna Maria Di Sciullo and Virginia Hill, 87–124. John Benjamins.

Farkas, Donka. 2011. Polarity particles in English and Romanian. In *Romance Linguistics 2010: Selected papers from the 40th Linguistic Symposium on Romance Linguistics (LSRL)*, ed. Julia Herschensohn, 303–328. John Benjamins.

Farkas, Donka, and Kim Bruce. 2010. On Reacting to Assertions and Polar Questions. *Journal of Semantics* 27:81–118.

Farkas, Donka, and Floris Roelofsen. 2012. Polar initiatives and polarity particles in an inquisitive discourse model. UCSC.

Farkas, Donka, and Floris Roelofsen. 2013. Assertions, questions, and hybrids in an inquisitive discourse model. Course at the 25th European Summer School on Logic, Language, and Information. August 5–16, Düsseldorf.

Grice, H. P. 1975. Logic and conversation. In *Syntax and semantics volume 3: Speech acts*, ed. Peter Cole and Jerry L. Morgan, 41–58. New York: Academic Press.

Grice, H. P. 1989. *Studies in the way of words*. Harvard University Press.

Groenendijk, Jeroen, and Floris Roelofsen. 2009. Inquisitive semantics and pragmatics. Handout of talk presented at the Workshop on Lauguage, Communication, and Rational Agency at Stanford, May 2009.

Grosz, Patrick. 2010. German *doch*: An element that triggers a contrast presupposition. In *Proceedings of the Chicago Linguistic Society 46*. University of Chicago Press.

Gunlogson, Christine. 2003. *True to form: Rising and falling declaratives as questions in English*. New York: Routledge.

Gunlogson, Christine. 2008. A question of commitment. *Belgian Journal of Linguistics* 22:101–136.

Gupta, Anthea Fraser. 1992. The pragmatic particles of Singapore Colloquial English. *Journal of Pragmatics* 18:31–57.

Gupta, Anthea Fraser. 2006. Epistemic modalities and the discourse particles of Singlish. In *Approaches to discourse particles*. Elsevier.

Hamblin, C. I. 1958. Questions. *Australasian Journal of Philosophy* 36:159–168.

Hamblin, C. I. 1970. *Fallacies*. Vale.

Hamblin, C. I. 1971. Mathematical models of dialogue. *Theoria* 37:130–55.

Han, Chung-Hye. 2002. Interpreting interrogatives as rhetorical questions. *Lingua* 112:201–229.

Hara, Yurie. 2010. Non-propositional modal meaning. In *Proceedings of Sinn und Bedeutung 10*.

Harnish, Robert M. 2005. Commitments and speech acts. *Philosophica* 75:11–41.

Hartung, S. 2007. Forms of negation in polar questions. In *Proceedings of the Workshop on Negation and Polarity at the University of Tubingen*, ed. H. Zeijlstra and J.-P. Soehn, 58–63.

Höhle, Tilman. 1992. Über Verum Fokus im Deutschen. In *Informationsstruktur und Grammatik*.

Holmes, Janet. 1984. Modifying illocutionary force. *Journal of Pragmatics* 8:345–365.

Horn, Laurence R. 1989. *A natural history of negation*. University of Chicago Press, Chicago.

Huddleston, Rodney, and Geoffrey K. Pullum. 2002. *The Cambridge Grammar of the English language*. Cambridge University Press.

Itani, Reiko. 1992. Japanese sentence-final particle *ne*: A relevance-theoretic approach. In *UCLWPL 4*, ed. Hans van de Koot.

Itani, Reiko. 1995. Semantics and pragmatics of hedges in English and Japanese. Doctoral Dissertation, University College London.

Itani, Reiko. 1998. A relevance-based analysis of hearsay particles: With special reference to Japanese sentence-final particle *tte*. In *Relevance theory: Applications and implications*, ed. Robyn Carston, Nam-Sun Song, and Seiji Uchida. John Benjamins.

Izvorski, Roumyana. 1997. The present perfect as an epistemic modal. In *Proceedings of SALT 7*, ed. Aaron Lawson, 222–239. Cornell University Press.

Jucker, A.H., and S.W. Smith. 1998. And people just you know like 'wow': Discourse markers as negotiation strategies. In *Discourse markers: Descriptions and theory*, ed. A.H. Jucker and S.W. Smith, 177–201. John Benjamins.

Kamio, Akio. 1994. The theory of territory of information: The case of Japanese. *Journal of Pragmatics* 21:67–100.

Kaufmann, Stefan. 1997. What's at stack in discourse. In *Proceedings of the 11th Amsterdam Colloquium*, 187–192.

Kaufmann, Stefan. 2000. Dynamic context management. In *Formalizing the dynamics of information*, ed. Martina T. Faller, Stefan Kaufmann, and Marc Pauly, 171–188. CSLI Publications.

Kimps, Ditte. 2007. Declarative constant polarity tag questions: A data-driven analysis of their form, meaning and attitudinal uses. *Journal of Pragmatics* 39:270–291.

King, Joseph. 2011. Post-syntactic attachment: A derivational analysis of tag questions. Ms., UCSC.

Kratzer, Angelika. 1981. The Notional Category of Modality. In *Words, worlds, and contexts*, ed. Hans-Jürgen Eikmeyer and Hannes Rieser, 38–74. Walter de Gruyter.

Kratzer, Angelika. 1987. Modality. In *Handbook of semantics*, ed. Arnim von Stechow and Deiter Wunderlich, 639–650. Walter de Gruyter.

Krawczyk, Elizabeth. 2009. Do you have evidence for that evidential? In *Theoretical approaches to understudied languages*, ed. Corinne Hutchinson and Elizabeth Krawczyk, volume 7 of *Georgetown Working Papers in Theoretical Linguistics*. Georgetown University.

Krifka, Manfred. 2004. Semantics below and above speech acts. Handout of talk presented at Stanford University.

Krifka, Manfred. 2011. Questions and question acts. Handout of a talk given at the 18th Amsterdam Colloquium.

Krifka, Manfred. 2012. Negated polarity questions as denegations of assertions. Humboldt-Universität Berlin.

Kwan Terry, A. 1978. The meaning and the source of the 'la' and 'what' particles in Singapore English. *RELC Journal* 9:22–35.

Kwan Terry, A. 1992. Towards a dictionary of Singapore English: Issues relating to making entries for particles in Singapore English. In *Words in a cultural context*, ed. A. Pakir, 62–72. Unipress.

Ladd, Robert. 1981. A first look at the semantics and pragmatics of negative questions and tag questions. In *Proceedings of Chicago Linguistics Society (CLS) 17*, 164–171.

Ladusaw, William A. 2003. Framing the issues: The biasing effect of polarity items in questions. Ms.

Leibniz, G.W. 1949. *New essays concerning human understanding*. La Salle, IL: The Open Court Publishing Company. [1704].

Ler, Soon Lay. 2005. An in-depth study of discourse particles in Singapore English. Doctoral Dissertation, National University of Singapore.

Ler, Soon Lay. 2006. A relevance-theoretic approach to discourse particles in Singapore English. In *Approaches to discourse particles*, ed. Kerstin Fischer. Elsevier.

Lock, Kit-Ken, and Mei-Yin Low. 1988. A proposed descriptive framework for the pragmatic meanings of the particle *la* in colloquial Singapore English. In *Asian Pacific Papers. Regional papers presented at the 8th World Congress of Applied Linguistics*, ed. B. McCarthy, 150–161. Applied Linguistics Association of Australia.

Luke, K.K. 1990. *Utterance particles in Cantonese conversation*. John Benjamins.

Malamud, Sophia, and Tamina Stephenson. 2011. Three ways to avoid commitments: Declarative force modifiers in the conversational scoreboard. In *Proceedings of the 15th Workshop on the Semantics and Pragmatics of Dialogue*, 74–83.

Maslova, E. 2003. Evidentiality in Yukaghir. In *Studies in evidentiality*, ed. Alexandra Yu Aikhenvald and R.M.W. Dixon, 219–35. John Benjamins.

Matthewson, Lisa, Hotze Rullmann, and Henry Davis. 2007. Evidentials as epistemic modals: Evidence from St'át'imcets. In *Linguistic Variation Yearbook 2007*, ed. Jeroen Van Craenenbroeck. John Benjamins.

McCready, Eric. 2008a. Evidentials, knowledge and belief. Ms.

McCready, Eric. 2008b. Particles: Dynamics vs. utility. In *Japanese/Korean Linguistics 16*, ed. Mutsuko Endo Hudson, Peter Sells, and Sun-Ah Jun. Center for the Study of Language and Information.

McCready, Eric. 2011. What is evidence in natural language? Ms.

McCready, Eric, and Norry Ogata. 2007. Evidentiality, modality and probability. *Linguistics and Philosophy* 30:147–206.

Murray, Sarah E. 2010. Evidentials and the structure of speech acts. Doctoral Dissertation, Rutgers University.

Murray, Sarah E. 2011. A Hamblin semantics for evidentials. In *Proceedings from Semantics and Linguistic Theory XIX*.

Murray, Sarah E. 2012. Varieties of update. Ms., Cornell.

Murray, Sarah E., and Jessica Rett. to appear. A semantic account of mirative evidentials. In *Proceedings from Semantics and Linguistic Theory 23*.

Papafragou, Anna. 2006. Epistemic modality and truth conditions. *Lingua* 116:1688–1702.

Platt, J., and Mian Lian Ho. 1989. Discourse particles in Singapore English: Substratum influences and universals. *World Englishes* 8:215–221.

Pope, Emily. 1976. *Questions and answers in english*. Walter de Gruyter.

Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In *Proceedings of SALT 14*, ed. Kazuha Watanabe and Robert B. Young.

Portner, Paul. 2007. Imperatives and modals. *Natural Language Semantics* 15:351–383.

Poschmann, Claudia. 2008. All declarative questions are attributive? *Belgian Journal of Linguistics* 22:247–269.

Potts, Christopher. 2007. The expressive dimension. *Theoretical Linguistics* 33:165–197.

Reese, Brian. 2007. Bias in questions. Doctoral Dissertation, University of Texas at Austin.

Reese, Brian, and Nicholas Asher. 2009. *Information structure*, chapter Biased questions, intonation, and discourse. Oxford University Press.

Repp, Sophie. 2012. Common ground management: Modal particles, illocutionary negation, and verum. In *Explorations of conventional non-truth-conditional meaning*, ed. Daniel Gutzmann and Hans-Martin Gaertner. Oxford University Press.

Rett, Jessica. 2011. Exclamatives, degrees and speech acts. *Linguistics and Philosophy* 34:411–442.

Richards, J.C., and M.W.J. Tay. 1977. The *la* particle in Singapore English. In *The English Language in Singapore*, ed. W. Crewe, 141–156. Singapore, Kuala Lumpur, Hong Kong: Eastern Universities Press.

Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. In *Ohio State University Working Papers in Linguistics 49*, ed. Jae Hak Yoon and Andreas Kathol, 91–136. The Ohio State University.

Roelofsen, Floris, Noortje Venhuizen, and Galit Weidman Sassoon. 2012. Positive and negative polar questions in discourse. In *Proceedings of Sinn und Bedeutung 17*, ed. E. Chemla, V. Homer, and G. Winterstein.

Rohde, Hannah. 2006. Rhetorical questions as redundant interrogatives. *San Diego Linguistics Papers* 2:134–168.

Romero, Maribel. 2005. Two approaches to biased *yes/no* questions. In *Proceedings of WCCFL 24*, 352–360.

Romero, Maribel, and Chung-Hye Han. 2004. On negative yes/no questions. *Linguistics and Philosophy* 27:609–658.

Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1.

van Rooy, Robert. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26:727–763.

van Rooy, Robert, and Marie Safarova. 2003. On polar questions. In *Proceedings of SALT 13*, ed. Robert B. Young and Yuping Zhao, 292–309.

Saigo, Hideki. 2011. *The Japanese sentence-final particles in talk-in-interaction*. John Benjamins.

Sailor, Craig. 2009. Tagged for deletion: A typological approach to vp ellipsis in tag questions. Master's thesis, UCLA.

Sailor, Craig. 2013. Questionable negation. Presented at the Linguistic Society of America (LSA) 2013 Annual Meeting, Boston, MA.

Schlenker, Philippe. 2006. 'Maximize Presupposition' and Gricean reasoning. Ms., UCLA and IJN.

Schmerling, Susan. 1982. How imperatives are special, and how they aren't. In *Papers from the parasession on nondeclaratives*, 202–218. Chicago Linguistic Society.

Searle, John R. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge University Press.

Simons, Mandy. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117:1034–1056.

Speas, Peggy. 2004. Evidentiality, logophoricity and tha syntactic representation of pragmatic features. *Lingua* 114:255–276.

Sperber, Dan, and Deirdre Wilson. 1986. *Relevance: Communication and cognition*. Blackwell Publishers.

Stalnaker, Robert. 1978. Assertion. In *Syntax and semantics volume 9: Pragmatics*, ed. Peter Cole, 315–332. Academic Press.

Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 701–721.

Sudo, Yasutada. 2013. Biased polar questions in English and Japanese. In *Beyond Expressives: Explorations in Use-Conditional Meaning*, ed. Daniel Gutzmann and Hans-Martin Gaertner, Current Research in the Semantics/Pragmatics Interface (CRiSPI) 28. Brill: Leiden.

Tottie, Gunnel. 2006. Tag questions in British and American English. *Journal of English Linguistics* 34:283–311.

Uyeno, Tazuko. 1971. A study of Japanese modality: A performative analysis of sentence particles. Doctoral Dissertation, University of Michigan.

Vanderveken, Daniel. 1997. Formal pragmatics of non literal meaning. In *Linguistische Berichte Sonderheft 8*, ed. E. Rolf, 324–341. Opladen: Westdeuscher Verlag.

Walton, Douglas. 1993. Commitment, types of dialogue, and fallacies. *Informal Logic* 14:93–103.

Willett, Thomas. 1988. A cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language* 12:51–97.

Zimmermann, Malte. 2004. Discourse particles in the left periphery. In *Proceedings of the Dislocated Elements Workshop*, volume 35 of *ZAS Papers in Linguistics*.