Title
Rhetorical questions as redundant interrogatives

Permalink
https://escholarship.org/uc/item/4xd7t5ww

Author
Rohde, Hannah

Publication Date
2006-05-01

RHETORICAL QUESTIONS AS REDUNDANT INTERROGATIVES[*]

Hannah Rohde
University of California, San Diego

Standard analyses associate rhetorical questions with negative singleton answers, but I introduce new data (including corpus data from Switchboard) with rhetorical questions that have a wider range of answers: (1) *Negative answer* Who lifted a finger to help? (2) *Positive answer* Has the educational system been so watered down that anybody who's above average is now gifted? (3) *Non-null answer* Who always shows up late to class? (4) *Multiple answers* What's going to happen to these kids when they grow up? I propose a new condition for felicitous use of rhetorical questions: the Speaker and Addressee must share prior commitments to similar and obvious answers. I frame the analysis of shared answers within Gunlogson's (2001) model of Common Ground, and I define sufficiently similar answers as those that either have an identical value or share an extreme position on a contextually relevant scale. Obviousness is measured using van Rooy's (2003) implementation of entropy for calculating the predictability of an answer set. Because they invoke an answer set, rhetorical questions resemble interrogatives, but the obviousness of a particular answer implies the bias of an assertion. As such, they are assertive, yet uninformative: instead of informing any discourse participant, rhetorical questions are redundant and serve to synchronize Speaker and Addressee beliefs.

1 Introduction

Rhetorical questions are examples of utterances whose form does not match their function. They have the structure of a question but the force of an assertion and so are generally defined as questions that neither seek information nor elicit an answer (Borkin 1971, Sadock 1971, Banuazizi 1999, and others). This makes them unique within semantic and pragmatic analyses since most utterances are assumed to be informative or at least information-seeking.

    Standard analyses typically associate rhetorical questions with single negative answers (Krifka 1995, Han 1998). Some researchers define them as biased assertions (Sadock 1971),

while others treat them as constrained questions (van Rooy 2003). I introduce new data (including naturally occurring data from the Switchboard corpus[1]) for which no prior account exists. I propose an analysis to account for rhetorical questions with a wider range of answers: positive/negative, null/non-null, single/multiple.

(1) *Negative answer*: Who lifted a finger to help?

(2) *Positive answer*: Has the educational system been so watered down that anybody who's above average is now gifted?   [sw_0393_2407]

(3) *Non-null answer*: Who always shows up late to class?

(4) *Multiple answers* What's going to happen to these kids when they grow up?
                                                                          [sw_0255_4548]

From the new data, a more general characteristic of rhetorical questions emerges: to be felicitous, rhetorical questions require that discourse participants share a prior commitment to similar, obvious, and often extreme answers. As such, rhetorical questions are biased, yet at the same time uninformative. Their effect is to synchronize discourse participants' commitments, confirming their shared beliefs about the world. The new data shows that rhetorical questions are neither information-seeking like regular questions nor information-providing like statements. I classify them as redundant interrogatives.

I consider the task of determining whether a rhetorical question succeeds based on the state of discourse participants' commitment sets in the Common Ground. I take the Common Ground to be the union of all participants' public beliefs (Stalnaker 1978). For every question, regular or rhetorical, each discourse participant invokes a set of possible answers. To measure the felicity of a rhetorical question, I look at whether the participants share an answer and whether that answer is obvious. I formalize this as three felicity conditions that must hold between the answer sets of the Speaker and Addressee: (i) the presence of an **obvious** answer, (ii) the **uninformativity** of the answer, and (iii) the **sufficient similarity** of the Speaker and Addressee's answers.

For condition (i), obviousness is measured with the mathematical tool of entropy which calculates the predictability of the Speaker and Addressee's answer sets. The presence of an obvious answer renders the answer set as a whole highly predictable, or minimally entropic.

(5) a. A: So, who's your favorite team?
     b. B: Who do you think? The Dallas Cowboys!   [sw_0389_2157]

To succeed rhetorically, (5b) requires that both participants know or be able to derive the obvious answer to (5a). *The Dallas Cowboys* are the highly predictable answer. I use van

Rooy's (2003) implementation of entropy for calculating the predictability of an answer set. However, unlike van Rooy who equates regular and rhetorical questions, I use the same entropy model to contrast the unpredictability of regular-question answers with the predictability of rhetorical-question answers. The intuition is that answers to rhetorical questions are so shared and predictable that they need not be uttered.

Informativity and answer similarity are both determined by examining the Speaker and Addressee's commitment sets. I frame the analysis of shared answers in terms of Gunlogson's (2001) model of Common Ground in which the participants' commitment sets are held separate. This separation makes it possible to isolate and then compare the Speaker and Addressee's obvious answers. Condition (ii) requires that a rhetorical question induce no updates to any participant's commitments—each participant knows the answer and they each know that they both know the answer, rendering the utterance uninformative.

(6)   How high are taxes going to be when my kids are my age?   [sw_0729_4840]

If all participants are already committed to high tax rates of the future, no commitment updates are induced and the rhetorical question succeeds. Beyond being uninformative, the answers must also be similar, as established in condition (iii). I define sufficiently similar answers as those that either have an identical value or those that share an extreme position on a contextually relevant scale. The similarity of singleton answers can be established through direct comparison—for example, *the Dallas Cowboys* in (5). For multiple answers and vague answers, the probability distribution must highlight answers at the same end of a contextually relevant scale – good/bad adulthoods in (4) or tax rates in (6).

Unlike previous analyses, I make no predictions about what the answer to a rhetorical question will be – I capture instead the Common Ground conditions that permit the felicitous use of a rhetorical question. The constraints I propose for sufficient similarity between Speaker and Addressee answer sets provide an account for new data from the Switchboard corpus in which rhetorical questions are associated with multiple answers or answers whose actual value is vague or otherwise underspecified.

Though I include many more types of rhetorical questions under my analysis than previous accounts, I still exclude from consideration rhetorical questions that only simulate a dialog. The Switchboard annotators label questions like *you know what?* as rhetorical, but the questions appear only to superimpose the structure of a dialog onto the discourse without leading the Addressee to any conclusion. These questions also fail to pass rhetorical question tests (ex. felicity of following *yet* clause, Sadock (1971)).

This paper aims to provide an account of rhetorical questions as redundant interrogatives. Section 2 reviews previous approaches to rhetorical questions. Section 3 introduces the corpus data and the variable question/answer patterns. Section 4 describes the theoretical frameworks that allow for a reanalysis, and Section 5 defines conditions for the required predictability and similarity of discourse participants' answer sets.

2 Background

In this section I describe the main conclusions from previous work. For analyses of rhetorical questions as assertions, I note the problems that arise from limiting the domain of rhetorical questions to those with single, negative answers. For the treatment of these questions as regular interrogatives, I contrast their supposed information-seeking status with the reality of their redundancy.

2.1 Rhetorical questions as assertions

Rhetorical questions have been described as selecting and asserting a single, negative or null answer from the answer set of the original non-rhetorical question (Borkin 1971, Sadock 1971, Krifka 1995, Lee 1994, Han 1998, Banuazizi & Creswell 1999, Ladusaw 2003).

(7) *yes/no* Question:
    a. Did John ever help?
    b. *Implies*: John never helped.

(8) *wh-* Question:
    a. What has John ever done to help?
    b. *Implies*: John has done nothing to help.

The rhetorical questions in (7a) and (8a) target single negative answers, either an assertion of opposite polarity (7b) or an assertion about the empty set (8b). Previous analyses account for the behavior of these rhetorical questions by appealing to negative operators (Han 1998, Lee 1994) or the influence of negative polarity items (Fauconnier 1975, Horn 1989, Progovac 1993, Guerzoni 2002, and others).

    Rhetorical questions that invoke non-negative answers are generally relegated to footnotes. They make assertions of positive polarity or specify unique non-null answers.

(9)   a. A: Is Clinton a liberal?
      b. B: Is the Pope Catholic?   (Han 1998, footnote 3, p. 8)

In uttering the rhetorical tag question in (9b), Speaker B takes advantage of shared knowledge that the Pope is indeed Catholic, using the strongly biased rhetorical question in (9b) to indicate that the answer to (9a) is similarly obvious. Question (9b) is rhetorical in that it does not elicit an answer from the Addressee, but contrary to predictions about negative singleton answers, there is nothing negative about the implied answer. Question (10) is also not negative and does not imply the predicted null answer *nobody*. Instead, it highlights a salient individual from the answer set, namely *mom*.

Existing theories incorrectly predict that the questions in (9b) and (10) correspond to an assertion of opposite polarity (*the Pope is not Catholic*) and to an assertion about the empty set (*no one has fed you and given you a proper education*), respectively. Han admits that these examples must be treated exceptionally.

(10)   Who has fed you and given you a proper education?
       CONTEXT: a mother to her son   (Han 1998, footnote 4, p.9)

Another set of questions that are routinely dismissed are rhetorical adjunct questions. They have been claimed to simply not exist (Gutierrez-Rexach 1996) or to have limited usage as exclamations of surprise for events that are presupposed to have already occurred (Lee 1994, Lawler 1971). Either way, previous work on rhetorical questions neither anticipates nor accounts for the wide array of questions seen in the Switchboard corpus where we find adjunct rhetorical questions as well as rhetorical questions associated with non-negative and non-null answers. The adjunct questions deserve further treatment because they appear to invoke a scale for ordering the possible answers (ex. time or quantity) and their obvious answers are tied to a position of extremity on that scale.

Though the rhetorical questions in the literature account for only a subset of the new corpus data, a common theme threads the variable data together: a recurrent pattern of shared, obvious answers known to both the Speaker and Addressee. Banuazizi & Creswell (1999) capture this in the Switchboard annotation guidelines for rhetorical questions:

> The answers to rhetorical questions are, in the speaker's mind, either perfectly obvious or perfectly obviously unknowable – in neither case would [the speaker] require or expect an answer. (page 10)

Despite this generalization, some analyses treat rhetorical questions as information-seeking.

2.2   Rhetorical questions as information-seeking questions

Ladusaw (1980) and van Rooy (2003) suggest that rhetorical questions function like regular questions and invoke a set of possible answers. The answer set is constrained by the presence of negative polarity items and by requirements of maximum entropy.

(11)   a. *Rhetorical question*: Did John lift finger to help?
       b. *Invoked answer set*: $\left[\begin{array}{l} \text{(i) Yes, John lifted a finger to help.} \\ \text{(ii) No, John did not lift a finger to help.} \end{array}\right]$

Ladusaw's constraint on question well-formedness requires that a speaker pose a question only when the answer is expressible without major revision to the original question. (11a) conveys the Speaker's belief that no one lifted a finger because (11bii) is the only answer that the speaker can rightfully expect given the constraint on minimal modification. Since (11bi) contains an unlicensed negative polarity item (i.e. *lift a finger* appears without overt negation in the clause), a *yes-* answer to (11a) would require significant modifications (for example, changing (11bi) to *actually, yes, John helped a lot*).

Van Rooy (2003) also treats rhetorical questions as regular questions, attributing their rhetorical effect to the negativity of all answers in the answer set and to the implicature

that only the bottom elements of a scale are relevant. Under van Rooy's information-theoretic approach, Speakers should only pose questions when the answer is unknown and the answer set is maximally entropic. Under this approach, (11a) is only appropriate if it is unresolved whether John lifted a finger. The negative polarity item invokes a scale of helpful contributions, and the question carries the implication that the Speaker knows what John contributed for all other units of help. The only remaining unknown is whether John contributed the bare minimum of *lifting a finger*. The rhetorical effect is achieved since both equally probable answers speak poorly of John (either he contributed next to nothing by lifting a finger or he really did absolutely nothing at all).

Though it is difficult to reconcile van Rooy's theory with the redundancy of an answer to a rhetorical question, his approach captures the useful notion that questions can be evaluated by the informativity of their answers. He establishes the connection between entropy and natural language questions, incorporating probabilities into the answer set and measuring bias over the set of all answers. His conclusion, that answers to rhetorical questions are informative and cannot be predicted, contrasts sharply with previous analyses which assume that the answer is obvious (Han 1998, and others) or else derivable from the structure of the question (Ladusaw 1980).

Though the responses that follow rhetorical questions suggest that they are not information-seeking and the lack of consistently negative responses suggests that the answer is not derivable from the structure of the question, I nonetheless adopt two insights from the previous research. In keeping with the classification of rhetorical questions as assertions, I will show that the answers to rhetorical questions are obvious and available to both Speaker and Addressee. In keeping with the information-seeking approach to questions, I will claim that answers are selected by the discourse participants themselves and are not calculable from the structure of the question.

3   New Data

The Switchboard data challenges restrictions on the kinds of questions that have rhetorical readings and confirms the generalization that rhetorical questions do not elicit an answer. Section 3.1 describes the range of questions seen in the corpus, confirming their status as rhetorical. Section 3.2 compares regular and rhetorical questions, showing that rhetorical questions pattern with assertions in the types of responses they generate.

3.1   Range of rhetorical questions

The annotated Switchboard corpus from which I draw my examples contains 1155 discourse-annotated conversations comprising 205,000 utterances and 1.4 million words.[2] Of the 583 questions in the corpus that are tagged as rhetorical, many go beyond the

---

[2]Inter-annotator agreement was 0.80 using the Kappa statistic, though no information is provided about the inter-annotator agreement of individual discourse tags, such as rhetorical questions.

traditional class with negative singleton answers, indicating the importance of modeling 'exceptions' like rhetorical questions with non-null, non-negative, and multiple answers.

Traditional data from the literature associates *yes/no* and *wh-* rhetorical questions with a negative singleton answer. In the Switchboard corpus, one finds examples of traditional polarizing *yes/no* questions as well as *wh-* questions with empty set answers.

(12)   I mean, can you force somebody to be a good productive citizen?   → *no*

[sw_0168_4765]

(13)   Who would steal a newspaper?   → *no one*

[sw_0141_2060]

The felicity of (12) and (13) depends on the shared sense of absurdity associated with stealing a newspaper or forcing someone to be a good citizen. Both questions point to a negative answer. However, as noted in the literature, some rhetorical questions do pick out a single non-null, non-negative answer from the answer set. The examples from the literature include tag questions like *Is the Pope Catholic?* and *wh-* questions like *Who has fed you and given you a proper education?* (Han 1998). In the Switchboard corpus, there are similar non-negative questions and non-null *wh-* questions that function rhetorically, as in (14) and (15b).

(14)   Has the educational system been watered down so that
       anybody who's above average is now gifted?   →*yes*   [sw_0393_2407]
       CONTEXT: public schools in America

(15)   a. A: So, who's your favorite team?
       b. B: Who do you think? The Dallas Cowboys!   [sw_0389_2157]

The rhetorical question in (14) assumes that all participants are aware of the deterioration of the public schools. Similarly, (15b) signals that the answer to the information-seeking question in (15a) is known or should be easily deduced given the context. Even though B provides an answer, the obvious "duh" answer informs no one because both participants already know or can derive the answer from their knowledge about B's public commitments (i.e. B's silver and blue jacket). The obvious answer is a salient non-null answer, *the Dallas Cowboys.*

The Switchboard corpus also contains examples of non-argument *wh-* rhetorical questions. Like the *yes/no* and argument *wh-* rhetorical questions above, non-argument *wh-* rhetorical questions can invoke negative or extreme answers.

Examples (16-19) refute the claim that rhetorical questions are limited to only *yes/no* and argument *wh-* questions, and they confirm my generalization that rhetorical question felicity is dependent on both participants' shared commitment to an obvious answer. Adjunct questions also invoke new kinds of answer scales (times, locations) whose positive and negative endpoints are made salient by the rhetorical question.

(16)   How can you justify raising anybody's salary if you have to lay people off?

[sw_0302_3309]

(17)   When is the sun going to come out?   [sw_1013_4822]
CONTEXT: life in outer space

(18)   Where's the money supposed to come from?   [sw_0458_2316]
CONTEXT: medical bills

(19)   Why do you do this to yourself?   [sw_0461_2436]
CONTEXT: actors who act on network TV

Question (20) conjures up a scale of normal-to-extreme outcomes of adulthood, ranging potentially from 2.5 kids and a dog to prison and terrorism. Though Speaker and Addressee may not share the same ranking, the question accesses a shared extreme region on the scale.

(20)   What's going to happen to these kids when they grow up?   [sw_0255_4548]
CONTEXT: juvenile delinquents

Though the Switchboard annotators tag these questions as rhetorical, one still must ask whether they are rhetorical by the standards of previous researchers. Several tests exist for determining the function of a question (Sadock 1971, Ladusaw 1980), including the presence of negative polarity items like *lift a finger* as well as the use of *yet-* clauses. Examples (21) and (22) demonstrate that the felicity of *yet-* clauses immediately following a question is restricted to rhetorical questions.

(21)   *Rhetorical question followed felicitously by yet clause*
    a. Did John lift a finger to help?
    b. Yet he didn't get punished.

(22)   *Regular question followed infelicitously by yet clause*
    a. What time is it?
    b. #Yet I have to leave.

A rhetorical question from the corpus patterns with previous rhetorical data.
Having described the variety of questions in the corpus and refuted the claim that rhetorical questions are restricted to negative singleton answers, I turn now to the kinds of responses that follow regular and rhetorical questions in the corpus. The response patterns provide support for claims that rhetorical questions do not function like regular questions and that they align more closely with statements.

(23) *Corpus question appears felicitously with yet- clause $\rightarrow$ rhetorical*
    a. What will become of these kids when the grow up?
    b. Yet the schools are doing nothing to help.   [sw_0255_4548]

3.2 Variation in responses to regular and rhetorical questions

Van Rooy's analysis predicts that rhetorical questions permit informative responses since all answers in the answer set have equal probability. My corpus study shows that responses to rhetorical questions are predominantly confirmations and backchannels, suggesting that rhetorical questions are uninformative and generate very little surprise. Discourse participants use backchannels (*uh huh, right, mmhmm*) to provide discourse-structuring feedback to the speaker (Jurafsky *et al.* 1998b).

With the exception of researchers like van Rooy (2003), there is general agreement that rhetorical questions do not elicit an answer but instead directly assert a biased position. The Addressee, however, may actually respond in certain cases: (i) if the rhetorical intent was misunderstood, (ii) if the question was an "exam-question", or (iii) if the response confirms or acknowledges the implication of the rhetorical question.

This last case, the confirmation, is the most common response following rhetorical questions in the Switchboard corpus. Regular questions, in contrast, consistently are followed with statements of opinion, statements of fact, or *yes/no* responses from the Addressee.
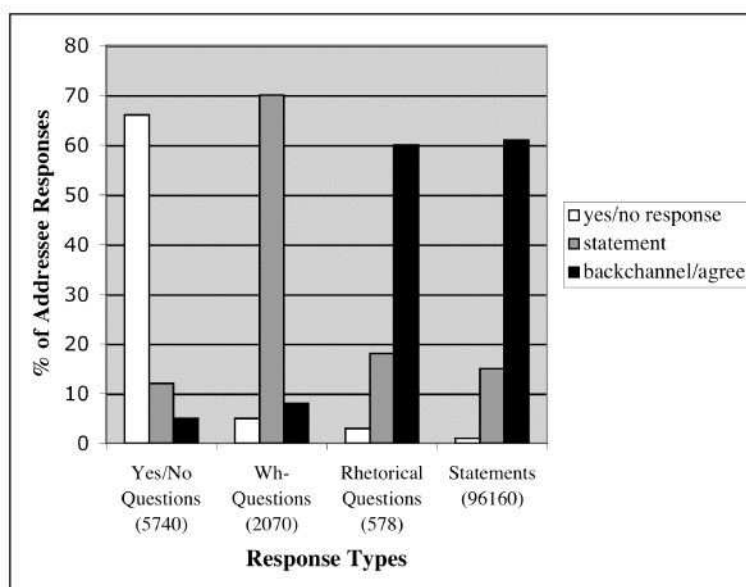

Figure 1. Addressee Responses to Regular and Rhetorical Questions

Figure (1) shows the distribution of responses to regular and rhetorical questions.

The most common responses to regular questions (*yes/no* answers and statements) are notably underrepresented in the responses to rhetorical questions. The most common responses to rhetorical questions (confirmations and backchannels), on the other hand, are the answers one typically finds following statements. Rhetorical questions and statements pattern together in response type and differ significantly from regular questions ($X^2(3) = 87.1940; p < .0001$). This suggests that rhetorical questions do not function in discourse like regular questions but instead express an opinion, to which the Addressee may provide confirmation or agreement.

In order to account for this new data I need to choose a framework for analyzing question form and function. The data suggests the use of a system for comparing Speaker and Addressee commitments and a metric to assess the obviousness of those commitments.

4   Frameworks

This section summarizes tools that will be important in the rest of the discussion. I build on aspects of Gunlogson's (2001) discourse model of Common Ground, and I rely on Groenendijk & Stokhof's (1997) semantic analysis of questions to construct partitions over possible worlds. I use entropy to establish a relationship between question-answer uncertainty and informativity, following van Rooy's (2003) adoption of Shannon (1948).

4.1   Common Ground

Felicitous use of a rhetorical question depends on shared commitments in the Common Ground. Gunlogson (2001) presents a new formalization of Common Ground, separating out participants' commitments into independent commitment sets composed of worlds that are consistent with that participant's public beliefs. The Common Ground consists of the union of worlds from all commitment sets. This approach refines Stalnaker's (1978) conception of Common Ground which collapses all participants' beliefs.

Gunlogson captures key properties of informativity and bias that are relevant to rhetorical questions, but her analysis is restricted to the behavior of *yes/no* questions which do not exhibit the full range of question intonation and structure seen in the rhetorical data. My analysis requires additional constraints beyond those imposed by Gunlogson's *yes/no* context updates, but I adopt her formalization of the Common Ground in order to gain access to the separate commitments held by two discourse participants. I also compare the characteristics of rhetorical questions with the properties she associates with falling intonation and interrogative structure.

In her contextual bias condition for rising declaratives, Gunlogson asserts that existing beliefs in the Common Ground can influence the felicity of an utterance. Her Contextual Bias Condition states that *Rising declaratives can only be used as questions in contexts where the Addressee is already publicly committed to the proposition expressed.* Rhetorical questions also impose restrictions on the existing commitment sets, requiring shared bias toward a sufficiently similar obvious answer.

Using a context-update model of discourse, Gunlogson establishes an important connection between the form/intonation of an utterance and the attribution of commitment. Declarative structure is linked to commitment update and bias while interrogative structure is connected to commitment preservation or neutrality. Under Gunlogson's model, falling declaratives commit the Speaker to the propositional content of an utterance, while rising declaratives commit the Addressee. Interrogatives commit no one and their context update is an identity function. In (24), a context C is updated with the falling declarative $\downarrow S_{decl}$.

(24)   Context Update of falling declarative (*It is raining.*)
  $C+ \downarrow S_{decl} = C'$ such that:
  a. $cs_{speaker}(C') = cs_{speaker}(C) + S_{decl}$
  b. $cs_{addressee}(C') = cs_{addressee}(C)$ (from Gunlogson ex. 78, p. 52)

After the commitment update for the falling declarative, the Speaker's commitment set $cs_{speaker}$ includes the content of the sentence. In general, declaratives induce an update to some participant's commitment set, and falling intonation assigns the commitment to the Speaker. In (25), the rising intonation associates commitment with the Addressee.

(25)   Context Update of rising declarative (*It's raining?*)
  $C+ \uparrow S_{decl} = C'$ such that:
  a. $cs_{speaker}(C') = cs_{speaker}(C)$
  b. $cs_{addressee}(C') = cs_{addressee}(C) + S_{decl}$ (from Gunlogson ex. 77, p. 52)

In contrast to declaratives, interrogatives act as identity functions, preserving the commitment sets of both the Speaker and the Addressee.

(26)   Context update of a rising interrogative (*Is it raining?*)
  $C+ \uparrow S_{interr} = C'$ such that:
    $cs_X(C') = cs_X(C)$

From (26), we see that a rising interrogative $\uparrow S_{interr}$ commits neither the Speaker nor the Addressee to its content and their commitment sets cs_X remain unchanged.[3] I show that rhetorical questions are also uninformative, not because they commit neither participant,

---

[3]The full context update is:
  (i) $cs_{speaker}(C') = cs_{speaker}(C)$
  (ii) $cs_{addressee}(C') = cs_{addressee}(C) + S_{interr}$ (from Gunlogson ex. 88, p. 56)

Given that interrogatives preserve commitments, it is unclear why the $cs_{addressee}$ contains the content of the interrogative (unless the update includes some kind of presupposition accommodation).

but because they require that both participants already be committed to a joint mutual belief. Joint commitments require all participants to not only believe a proposition P but also be aware that other participants believe P as well. Shared knowledge about others' beliefs and about the real world is precisely what licenses rhetorical questions.

4.2   Meaning of Questions

Gunlogson specifically restricts her analysis to *yes/no* questions, but much of my data includes *wh-* questions. Following Groenendijk & Stokhof (1997), I take the meaning of a question to be a partition over possible worlds. Each cell of the partition contains worlds in which some proposition P is true, where P counts as a full answer to the question $?xPx$. Two worlds occupy the same cell if they share the same answer to the question.

nobody is such that $\lambda x.Px$
a1 is the one that $\lambda x.Px$
a2 is the one that $\lambda x.Px$
a1 and a2 are the ones that $\lambda x.Px$
everybody such that $\lambda x.Px$

Figure 2. Set of partitions for a *wh-* question

Alternately, one can construct an answer set by referring to the entities that satisfy the predicate $\lambda x.[P(x)]$ using a structured meaning approach to questions (Krifka 2001). Either way, a set is generated consisting of possible answers. The advantage of a questions semantics using partitions is that the negative and null-set responses to questions are included in the answer set.

Though it remains unresolved in the literature, I include in a question's denotation answers to adjunct *wh-* questions like *when/where/why/how*. These questions are not represented in the form $?xPx$ because variables over times/places/reasons/manners are not constituent arguments of the predicate P. Instead they describe temporal, locative, reason, and manner information about an event. I assume a semantic analysis of adjunct questions which partitions the set of possible worlds into subsets indexed for non-argument variables like time and place. The negative answers are included as well: *no time, no place.*

For rhetorical questions with positive/negative, null/non-null, single/multiple answers, both participants' partitions must contain a shared obvious answer, and I turn now to the information-theoretic tool of entropy which allows me to measure uncertainty over the answer partition. The more obvious an answer in the answer set, the lower the uncertainty or entropy over that set.

4.3   Mathematical Interlude: Entropy

Entropy is often described using metaphors of surprise or uncertainty, and it is relevant here because it provides a metric for measuring the surprise associated with an answer

to a question. For regular questions, there is a high degree of uncertainty because the Speaker has asked a question for which he lacks an answer. For rhetorical questions, there is no uncertainty about the answer, and, in fact, the answer would be redundant.

Van Rooy makes the connection between entropy and natural-language question informativity. He defines a question well-formedness constraint, asserting that a good question is one with a maximally entropic answer set. Though I disagree with van Rooy's claim that rhetorical questions adhere to this constraint, I do follow his insight that an answer set can be evaluated in terms of its entropy.

Like van Rooy, I calculate entropy using Shannon's (1948) formula for the uncertainty of a system, measured in bits. Entropy calculates the predictability of a random variable over multiple trials. It estimates how surprised one will be, on average, to learn the value of the random variable. Consider a scenario in which three individuals, John, Mary, and Bob, are invited to a party. Question (27) is asked to determine who attended.

(27)   Who came to the party?

Entropy is calculated over the set of answer-probability pairs, shown in Figure 3.

$$
\begin{bmatrix}
\text{John} & \frac{1}{8} & \text{John \&Bob} & \frac{1}{8} \\
\text{Mary} & \frac{1}{8} & \text{Mary \&Bob} & \frac{1}{8} \\
\text{Bob} & \frac{1}{8} & \text{John \&Mary \&Bob} & \frac{1}{8} \\
\text{John \& Mary} & \frac{1}{8} & \oslash & \frac{1}{8}
\end{bmatrix}
$$

Figure 3. Uniform probability distribution over answers to (27)

The formula in (28) measures the entropy of a random variable X. The formula sums over all outcomes, combining the cost of expressing that outcome (the log probability) with the likelihood that that outcome occurs (the probability co-efficient).

(28) $H(X) = -\sum_{x \in X} P(x) log_2 P(x)$ DEFINITION OF ENTROPY

For the party scenario, the random variable is party attendance. Entropy measures how surprised one will be on average to learn which subset of people attended the party given a probability assignment over the eight possibilities. If all subsets are equally probable, then one will experience a high degree of surprise upon learning who attended.

Entropy is much lower when one answer is more probable than any other. The skewed probability distribution Figure 4 is associated with low entropy, or low surprise.

With the answer probabilities in Figure 4, one is rarely surprised to learn who attended the party because it is typically only John who attends. One can correctly predict the outcome half of the time, and therefore entropy is reduced as in (30).

(29)   *All answers equally likely → high surprise*
    a. $H(X) = -\left(8 * \left[\left(\frac{1}{8}\right) log_2 \left(\frac{1}{8}\right)\right]\right)$
    b. $=$ 3

$$\begin{bmatrix} \text{John} & \frac{32}{64} & \text{John \&Bob} & \frac{1}{64} \\ \text{Mary} & \frac{16}{64} & \text{Mary \&Bob} & \frac{1}{64} \\ \text{Bob} & \frac{8}{64} & \text{John \&Mary \&Bob} & \frac{1}{64} \\ \text{John \& Mary} & \frac{4}{64} & \oslash & \frac{1}{64} \end{bmatrix}$$

Figure 4. Skewed probability assignment to answers to (27)

The important fact to note is that skewed probabilities highlight a subset of answers and decrease the entropy over the answer set. The entropy drops from 3 with the unbiased question to 2 with the question that has a skewed probability distribution over its answers.[4]

Under my analysis, the case of complete lack of surprise corresponds to rhetorical questions because the answer is predictable to both the Speaker and the Addressee. The answer is so unsurprising that it need not be uttered at all. In this way, I propose to combine elements of Gunlogson's discourse model with a measure of entropy in order to determine what answer(s) the participants might select and how surprised they would be to hear the answer. The idea is that an answer to a rhetorical question should generate significantly less surprise than an answer to a regular question because rhetorical questions are asked only when the answer set is sufficiently skewed towards an obvious answer.

5   Common Ground conditions for rhetorical questions

In this section I establish the conditions under which a rhetorical question can be uttered felicitously. My claim is that rhetorical questions require probabilistically highlighted obvious answer(s) in the existing commitments of the discourse participants. These commitments must be sufficiently similar. In some cases, this means that the answer value itself needs to be shared, while in other cases the answer(s) may occupy similarly extreme regions of a relevant scale.

In 5.1, I motivate the new analysis by noting the problems that arise from focusing only on negative endpoints as the meaning of rhetorical questions. In 5.2, I use a measure of entropy to confirm the presence of an obvious answer in both participants' answer

---

[4]The value 3 is the number of bits it takes, on average, to represent the answer using binary digits. Since there are eight possibilities for party attendance, each one can be described with a 3-bit string (ex. 001 for John, 011 for John&Mary, 111 for John&Mary&Bob, etc.). If some answers are more probable than others, it makes more sense to represent common answers with shorter strings and rare answers with longer strings (ex. 1 for John, 0 for Mary, 01 for Bob, etc.). The cost drops from 3 bits on average to 2 bits.

(30)   *Subset of answers more likely than others → lower surprise*

$$
\begin{aligned}
H(X) &= -\left[\frac{32}{64}\log_2\frac{32}{64} + \frac{16}{64}\log_2\frac{16}{64} + \frac{8}{64}\log_2\frac{8}{64} + \frac{4}{64}\log_2\frac{4}{64} + 4x\frac{1}{64}\log_2\frac{1}{64}\right] \\
&= \frac{1}{2} + \frac{1}{2} + \frac{3}{8} + \frac{1}{4} + \frac{3}{8} \\
&= 2
\end{aligned}
$$

sets. In 5.3, I compare the context-update of a rhetorical question with that of a regular interrogative. In 5.4, I use corpus data to demonstrate that the obvious answer must also be shared. Sections 5.5 presents metrics for evaluating similarity between two answers, comparing either the actual answer values or their mutual extremity on a relevant scale. In 5.6, I discuss the pragmatic uses of rhetorical questions and show how my analysis makes predictions about cases where rhetorical questions fail.

5.1   Reasons for a new analysis

Unlike previous analyses which restrict the domain of rhetorical questions to those with negative answers, my analysis appeals to general conditions on discourse participants' answer sets to determine the felicity of these questions. This new approach has broader coverage over the space of possible rhetorical questions.

Negation plays a significant role in most of the rhetorical questions discussed in the literature (Han & Siegel 1997, van Rooy 2003) and is apparent in corpus examples like (31).

(31)   Can you force somebody to be a good productive citizen?   →*no*

[sw_0168_4765]

The rhetorical effect of question (31) is accounted for differently by Han and van Rooy, but the result is similar – clear coverage only over the class of negative rhetorical questions. As described earlier, Han uses a negative operator to predict the implied answer to (31), relegating positive-answer rhetorical questions to footnotes. Van Rooy uses an information-theoretic approach, predicting that a question about *forcing* someone to be a good citizen defines a negative bottom element of a scale. It implies that all other means of achieving good citizenship are known to not work (*letting s.o., asking s.o., begging s.o.*). In that case, the rhetorical effect arises from the fact that the *yes* and *no* answers both speak poorly of mankind's good citizenship – either good citizenship can be achieved, but only by force (the positive answer), or it cannot be achieved, not even by force (the negative answer).

Han's account also captures *wh-* rhetorical questions which are associated with the null set ($\oslash$). Though van Rooy does not provide the details for an extension to *wh-* questions, his use of bottom-element implicatures could be expected to locate the minimal element on some scale (possibly the empty set on a lattice constructed of supersets from $\oslash$ to all). These analyses make the wrong prediction for rhetorical questions that use constructions which require a non-null answer.

(32)   Who was it that fed and clothed you for 25 years?   [mom]

The *wh-* cleft construction triggers a presupposition that someone fed and clothed you. As rhetorical questions, *wh-* clefts point only to real-valued answers from the answer set (as in (32)). They are infelicitous with negative polarity items that point to the bottom of a scale (33) because the answer *no one* is not available.

(33)   *Who was it that lifted a finger to help?   [implied, no one]

Since a direct link between rhetorical questions and negation is untenable, I turn to more general characteristics of rhetorical questions. A common theme of the old and new data is that the answer is already available to both discourse participants. Whether the question invokes negative or positive answers, the underlying pattern is the same – each participant knows the answer and they each know that they both know the answer. Pragmatically, the effect is to highlight a shared answer and to synchronize the beliefs of Speaker and Addressee with regard to the content of that answer, whether it be positive or negative. The questions that my analysis covers include those with positive/negative, null/non-null, and single/multiple answers. I offer no account for rhetorical questions with an unknown answer, like example (34).

(34)   What are we going to do with all the soldiers over there?   [sw_0264_2252]
        CONTEXT: troops in Eastern Europe after the cold war

Here, it is the lack of an answer that must be shared between Speaker and Addressee (the response *I don't know*). My analysis proposes conditions that require shared knowledge of an obvious answer; it does not make predictions about rhetorical questions that require shared responses.

5.2   Requirement for an obvious answer

In order to succeed rhetorically, the question must carry strong bias towards an obvious answer in the Speaker and Addressee's commitment sets. The presence of an obvious answer renders the question itself, and its answer, uninformative. The result of this redundancy is that no answer is required or expected. Regular questions, on the other hand, are information-seeking, and their answer is highly informative.

(35) Who has spare time?   [sw_0959_3735]

Question (35) can function both rhetorically and non-rhetorically. One can genuinely inquire about who among a set of equally probable individuals has free time on his hands. That information-seeking question lacks a predictable answer (though there may be some bias over the possible answers). The rhetorical question, however, has a very predictable, obvious answer to which all participants are committed (for example, *no one* or *John*).

I measure predictability using the model of entropy that van Rooy (2003) employs for his analysis. Though I do not adopt his treatment of rhetorical questions as information-seeking, I incorporate his idea of measuring entropy over answer-probability pairs. I claim that the differences between questions with predictable and unpredictable answer sets can be captured by comparing their relative entropy. Rhetorical questions succeed when the presence of an obvious answer renders the answer set minimally entropic. Van Rooy claims, however, that all questions, even rhetorical ones, have maximally entropic answer sets.

Regular (unbiased) $\rightarrow$ high entropy	**Rhetorical (skewed to $\oslash$)** $\rightarrow$ low entropy	**Rhetorical (skewed to non-$\oslash$)** $\rightarrow$ low entropy
a. $\begin{bmatrix} \text{John} & 25\% \\ \text{Mary} & 25\% \\ \text{John \& Mary} & 25\% \\ \oslash & 25\% \end{bmatrix}$	b. $\begin{bmatrix} \text{John} & 1\% \\ \text{Mary} & 1\% \\ \text{John \& Mary} & 1\% \\ \oslash & 97\% \end{bmatrix}$	c. $\begin{bmatrix} \text{John} & 97\% \\ \text{Mary} & 1\% \\ \text{John \& Mary} & 1\% \\ \oslash & 1\% \end{bmatrix}$

Figure 5. Answer sets for regular and rhetorical interpretation to (35)

Figure 5 shows sample answer sets, corresponding to regular and rhetorical questions. In answer set (a), the probabilities proposed for the regular question are uniformly distributed over the answers, resulting in high surprise or high entropy.[5] The probabilities proposed for the rhetorical question are shown in answer sets (b) and (c) where the probability mass is skewed to the null set ($\oslash$) or to the non-null answer (*John*). In (b), the answer with the highest probability is the empty set, conforming to predictions in the literature about the negativity of the answer to a rhetorical question. In (c), however, the answer is a salient non-null individual *John* known to both participants for having copious free time. *John* does not represent a negative answer.

My claim is that rhetorical questions, rather than having a uniform distribution over answers like regular questions, are characterized by a skewed distribution over the answer

---

[5]Though Figure 5 shows a uniform distribution for the regular question, question bias might indicate a normal or near-normal distribution over answers – either way, the regular question's answer set still has higher entropy than one with a sharply skewed distribution.

set. In terms of the answer partition, the cell containing worlds in which the obvious answer is true receives the most probability mass.

c_0: nobody $\lambda x.Px$	$prob(c_0)$
c_1: a1 is the one that $\lambda x.Px$	$prob(c_1)$
c_2: a2 is the one that $\lambda x.Px$	$prob(c_2)$
c_3: a1 and a2 are the ones that $\lambda x.Px$	$prob(c_3)$
$\vdots$	
c_n: everybody $\lambda x.Px$	$prob(c_n)$

Figure 6. Set of partitions for a *wh-* question

The obvious answer in Figure 6 corresponds to the cell c_i with the highest probability $prob(c_i)$ in the partition. An implicit assumption here is that all worlds participate in the answer partition and in the Common Ground. Commitment updates do not eliminate worlds; they just reduce the probability on untrue worlds to a minimum likelihood. Given that all answers are still present and are associated with a probability, the predictability of the answer set can be measured as the entropy over the entire partition.

(36)   CONDITION 1: Answers to Rhetorical Questions are **Obvious**
    for a set of worlds $A_X = \text{ANSWER}(Q_{rhet}, X)$:

$$\frac{H(A_X)}{H(A_{unbiased})} < \theta, \text{ for } X \in [Speaker, Addressee]$$

CONDITION 1 establishes a metric for evaluating the presence of an obvious answer by comparing the ratio of the entropy of the Speaker and Addressee answer sets with that of an unbiased answer set. The function ANSWER takes as input a question and a participant X. It returns X's probability assignment over the answer partition for that question. If participant X believes that the answer to (35) *Who has spare time?* is the proposition *that nobody has spare time*, then ANSWER returns a probability assignment highlighting the cell containing worlds in which no one has spare time. The answer can be either alone in the spotlight of the skewed probability distribution or it may share the spotlight with a handful of other obvious answers by dividing the probability mass.

    Either way, the answer set of a rhetorical question must be sufficiently skewed towards the obvious answer(s) to guarantee a reduction in entropy between the unbiased answer set $(A_{unbiased})$ and the answer sets of the Speaker and Addressee $(A_{speaker}, A_{addressee})$. The ratio of the entropy of $H(A_{unbiased})$ and $H(A_X)$ must be less than some threshold θ where θ is crucially less than 1. A ratio of 1 represents the non-rhetorical equivalence between the entropy of two answer sets. $A_{unbiased}$ is the answer set of a hypothetical discourse participant who lacks any answer bias and assigns a uniform distribution over the partition.

5.3 Rhetorical questions and informativity

Rhetorical questions' lack of surprise is evidence of their uninformativity. Though rhetorical questions carry strong bias, they require no update to participants' commitment sets. This preservation of commitments is precisely what Gunlogson (2001) proposes for the class of rising interrogatives. According to Gunlogson, an utterance informs a participant's commitment set if the commitment set does not already entail it. If the commitment is entailed, as it is with rhetorical questions, the utterance generates no commitment updates.

(37)   CONDITION 2: Commitments are **Preserved** in context C after Q_{rhet} is uttered
   $C + Q_{rhet} = C'$ *such that:*
   $cs_X(C') = cs_X(C)$ for $X \in [Speaker, Addressee]$

CONDITION 2 uses Gunlogson's participant-separated model of Common Ground to show how rhetorical questions crucially update no participant's commitment set. Like regular interrogatives, they invoke an answer set, but they are *redundant* interrogatives because they serve only to reiterate information already in the Common Ground. A possible objection to this lack of informativity is the use of rhetorical question as tag questions to address a previous information-seeking question in the discourse, as in (38b).

(38)   a. A: Is Clinton a liberal?
   b. B: Is the Pope Catholic?

I assert that the obvious answer to (38b) is not informative – at least not with respect to the proposition *the Pope is Catholic*. What B does in (38) is signal to A that the answer to A's question is just as obvious as the answer to B's rhetorical question. Rhetorical questions may inform participants about the answer to another question, but the answer to the rhetorical question itself must be obvious and uninformative to everyone.

Given that rhetorical questions raise the same issues of informativity and commitment that Gunlogson (2001) addresses, it is worth establishing how my conception of rhetorical questions differs from categories in her utterance typology. As mentioned earlier, her context-update model of discourse defines classes of utterances differing in two dimensions: syntax (declarative/interrogative) and intonation (falling/rising). Her model targets *yes/no* questions; her idealizations about question syntax and intonation do not map directly to rhetorical questions or to most *wh-* questions.

	Falling Intonation	Rising Intonation
Declarative Structure	*It is raining.*	*It's raining?*
Interrogative Structure		*Is it raining?*

Figure 7. Gunlogon's utterance typology

Rhetorical questions share some properties with interrogatives (syntactic structure and preservation of commitment sets) and some properties with falling declaratives (intonation contour and Speaker commitment). This confirms the intuition that properties of both questions and assertions influence the behavior of rhetorical questions. At first glance, the open category of *falling interrogatives* is a tempting classification for rhetorical questions. Unfortunately rhetorical questions use neither consistent falling intonation nor consistent interrogative structure. What follows is a breakdown of the question-like and statement-like properties of rhetorical questions – a breakdown that shows the parallels between rhetorical and regular questions/assertions but concludes that they require their own usage conditions.

There is limited evidence of consistent intonation in rhetorical questions.[6] In her review of English question intonation, Bartel (1999) notes the wide variation for rhetorical questions, refuting claims that rhetorical questions rise consistently like *yes/no* questions or fall consistently like statements. For example, the same rhetorical question can be uttered with both rising and falling intonation.[7] In their study of questions in the Switchboard corpus, Banuazizi & Creswell (1999) show that *yes/no* rhetorical questions fall more often than regular questions, but still less than half the time.

Given the differences between rhetorical questions and real falling utterances, classifying rhetorical questions as regular interrogatives seems potentially more promising. Such a classification has the advantages of ensuring commitment set preservation and of invoking a set of possible answers, but it provides no means for predicting differences in responses to regular and rhetorical questions. Furthermore, though rhetorical questions usually pattern structurally with interrogatives, there are exceptions like the rising declarative in (39).

(39)   And you think I care?[8]

The variation in syntax and intonation makes it difficult to classify rhetorical questions in an utterance typology that uses categorical distinctions of intonation and syntactic structure. In contrast, Gunlogson's existing classes occupy non-overlapping domains: rising *yes/no* interrogatives differ from rising declaratives in sentence structure, and rising and falling declaratives are distinguished by their intonation. I mention these challenges

---

[6]Some evidence for rhetorical falling intonation comes from the large number of *wh-* adjunct rhetorical questions in the corpus – a class of questions which usually receive falling intonation (Bartel 1999, page 260).

[7]The following example from Bartel shows the two pitch contours that are available for the same rhetorical
   question (ToBI tonal transcription):
      (i) I don't see why I should pay for Greg's excesses.
         Am I my brother's keeper? (= I am not my brother's keeper.)
         (a)   H*/L* H-H%
         (b)   H*/L* L-L%   (from Bartel (1999), examples (5-6), page 252)

in identifying and classifying rhetorical questions as support for my claim that shared real world knowledge is crucial to an understanding of rhetorical questions. What characterizes rhetorical questions is their dependence on shared knowledge in the Common Ground.

5.4   Condition for Shared Answer

Up to now, we have seen the importance of a known, obvious answer and the accompanying preservation of participants' commitment sets. I now further constrain the obvious answers in individuals' commitment sets to be shared across participants. For instance, for question (40) to succeed rhetorically, both the Speaker and the Addressee need to share a commitment to the desirability of writer Danielle Steele's money.

(40)   And wouldn't I like to have her money?   → *yes*   [sw_0194_3595]
      CONTEXT: Danielle Steele

Question (40) is only a single example, but it exemplifies the pattern throughout the corpus that rhetorical questions rely on a shared answer between participants. I believe the necessity of a shared answer holds even if (40) is classified as an exclamative; Portner & Zanuttini (2000) describe exclamatives as questions that presuppose their content and for which it "makes no sense to provide the information [in an answer]" (p. 201). This description seems to blur the line between exclamatives that use interrogative structure and truly rhetorical questions. Further work is needed to determine the relationship between the factivity of exclamatives and the assertiveness of rhetorical questions.

As a surface-level indication of the joint commitment to a shared answer, I point to the frequency of the filler *you know* in rhetorical and regular questions.
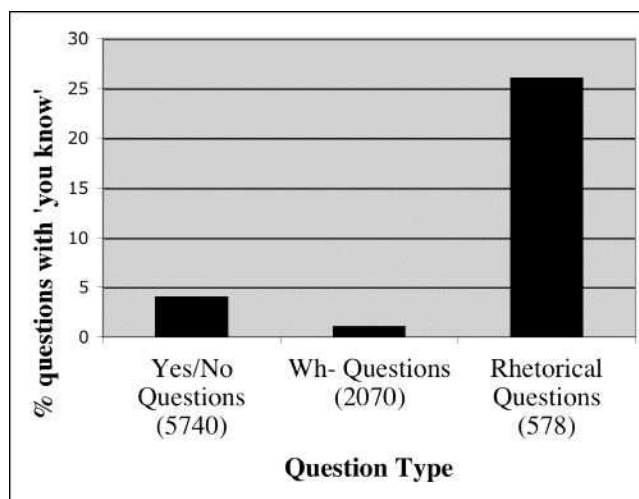

Figure 8. Frequency of *you know* across questions in Switchboard

Figure 8 shows the high frequency of the phrase *you know* in rhetorical questions (26%) and its relative scarcity in regular *yes/no* questions (4%), and *wh-* questions (1%). A test of statistical significance confirms that the three groups do not behave similarly and the exceptional group is the rhetorical group ($X^2(2) = 36.0645; p < .0001$). *You know* serves to emphasize the prior joint commitment to a shared answer inherent to rhetorical questions. An example from the corpus is shown in (41).

(41)   So it's like, what are you going to do, you know?   [sw_0440_4347]
       CONTEXT: futility of trying to buy only American-made products

Not only does the Speaker in (41) think there is no way to avoid buying foreign-made products, but he makes a gesture of synchronization, checking that the Addressee shares his belief by using the filler *you know.*

5.5   Condition for Sufficient Similarity

Shared beliefs between discourse participants can either be checked by comparing the two beliefs directly or by comparing the beliefs' positions on a relevant scale. The easiest evaluation for belief similarity is identity. However, when obvious answers to a rhetorical question cannot be directly compared, their mutual extremity on a relevant scale guarantees the presence of an obvious answer. Following the introduction of a test of answer equivalence, I discuss the kinds of scales dictated by questions, the necessity of a shared endpoint, and the predictions that arise from the presence of scales invoked by rhetorical questions.

The standard rhetorical questions described in the literature have singleton answers whose shared status across discourse participants can be directly confirmed: a shared negative answer for questions with negative polarity items, as in (42), or a shared empty set answer, as in (43).

(42)   Did John lift a finger to help?   → *no*

(43)   Who lifted a finger to help?   → *no one*

Answers are selected by the participants from their answer sets and are not usually derivable from the question – except in cases where a polarity item induces a scale and the rhetorical question highlights the bottom element. I now introduce a measure of mutual extremity to establish similarity in cases where the shared answers are located in the same region of a relevant scale but do not share identical values.

For questions that lack singleton answers, several properties must hold in order to ensure sufficient similarity: (i) the answer(s) must still be obvious, having a skewed probability distribution that reduces entropy; (ii) the distribution must be skewed towards one end of a relevant scale, highlighting a position at or near an endpoint, and (iii) the Speaker and Addressee must both locate their answers at the same endpoint of that scale.

Different forces can draw the answer to an extreme position along a scale: negative polarity items draw answers inevitably to an extreme bottom element, while other questions depend on real-world knowledge to select an extreme answer. I propose several different sources for scales: negative polarity items, *wh-* words of quantity, and lexical items within the question. I also make an appeal to a salient scale in the pragmatic context.

Throughout the literature, the emergence of scales has been linked to the presence of polarity items (Fauconnier 1975, Horn 1989, Progovac 1993, and Guerzoni 2002 and others). In (44), the answer partition can be ordered along a scale of prices, where a red cent is the bottom element. Those individuals who pay the least (or the cells containing those worlds in the answer partition) are ranked lowest on the scale.

(44)   Who paid a red cent for that old car?

The rhetorical question in (45) uses the quantity *wh-* word to dictate a numerical scale.

(45)   How many stars are there in the sky?

Though the question induces a scale with salient, easily comparable integer values, the participants' answers may not occupy identical positions along that scale. One individual might pick out a value of a *trillion* while another might consider a *googol* or a vague answer like *a whole heck of a lot*. In those cases, Speaker and Addressee answers may be different but they share sufficiently similar positions of mutual extremity along the numeric scale.

Rhetorical questions also induce scales without the help of negative polarity items or closed class *wh-* words. They may invoke a contextually relevant scale dictated by a lexical item in the utterance itself. The prevalence of rhetorical *wh-* adjunct data in the corpus provides examples of rhetorical questions whose scales are selected by the *wh-*phrase. In (46-48), an adjective in the *wh-* phrase determines the dimension of the scale.

(46)   How high will taxes be when my kids are my age?   →*very high*
       [sw_0729_4840]
       CONTEXT: taxes

(47)   How deprived could they be if they had a camcorder?   →*not very deprived*
       [sw_0410_2970]
       CONTEXT: material possessions

(48)   How soon are we going to start to get our money's worth?   →*not very soon*
       [sw_0677_3805]
       CONTEXT: government services

Dimensions of height and deprivation are made salient in these corpus examples. Without a negative polarity item to draw probability mass to the bottom element along the

scale, the answer(s) are permitted to cluster around either endpoint. The Speaker in (46) expects that taxes will be very high in the future, while the Speaker in (47) calls into question the real deprivation of someone who owns a camcorder.

In cases where there is no overt marker present to dictate the kind of scale, rhetorical questions invoke scales estimating *good/bad* along some contextually relevant dimension.

(49)   Who would steal a newspaper?   [sw_0141_2060]

The partition ordering imposed by (49) ranks worlds in which different individuals, at different times steal newspapers. Individuals who engage in newspaper stealing on a regular basis appear at the extreme endpoint of the scale. The high probability worlds are those in which no one steals newspapers and those in which lunatics do.


Figure 9. Participant's skewed distribution for (49)

Figure 9 shows the ranked list of possible newspaper thieves along the x-axis. The y-axis corresponds to probability values between 0 and 1. Since the discourse participant skews the distribution towards a single answer (the null answer), that answer is identified as obvious. Each individual on the x-axis corresponds to the partition cell containing worlds in which that individual is a newspaper thief. Here, the cell in which Grandma steals newspapers is given very low probability while the cell in which no one (∅) steals newspapers receives most of the probability mass. This distribution reflects the participants' knowledge and expectations about the real world.

Recent work confirms the importance of pragmatic scales which are determined by the content of the question (Hoeksema & Rullmann 2001, Wilkinson 1996). The scale of normal/abnormal individuals invoked by question (49) is a salient scale in much the same way as Hoekseman & Rullmann appeal to real-world context-driven rankings to invoke scales like tastiness, as in (50).

(50)   Tommy will not eat the most delicious food.
                        (Hoeksema & Rullmann (2001), example 3, page 131)

Whether they are numerical scales or deprivation scales or scales of tastiness, the scales for rhetorical questions can be characterized as one-dimensional bounded or unbounded

spectra, driven by the context and content of the question. Some scales are less obvious than tax rates and culinary taste; they require a measure of *good* and *bad* relevant to a topical domain. This scale of *good* and *bad* does not necessarily map onto more general metrics of *good* and *evil*, as shown in the scenario in (51).

(51) CONTEXT: writing a history paper on a famous person from WWII
    a. A: Who should I write about?
    b. B: Well, who was involved in the Hitler-Churchill standoff?

The discourse is structured so that the rhetorical question in (51b) uses the same domain of entities (namely, individuals) invoked by the information-seeking question (51a) and provides a ranking of *good/bad* paper topics. The obvious answers to the rhetorical question are Hitler and Churchill, two individuals who are crucially ranked differently on most scales of *good* and *evil*. Yet they are adjacent on a scale of good WWII paper topics.

    The example in (52) shows that mutual extremity also requires that the answers be skewed towards the same endpoint of the scale. The Speaker and Addressee may each hold an obvious answer to a question, but if those answers are not sufficiently similar, the synchronization fails. The participants will be forced to resort to some repair strategy in order to correct whatever inaccurate perceptions of the Common Ground led them astray.

(52) *Necessity of similarly skewed answer sets*
    a. A: How much do those linguists make?
    b. B: Yeah, they have no right to complain.   [*skew toward high salary*]
    c. B′: Yeah, they must really love their work. [*skew toward low salary*]

If B interprets (52a) as a complaint that linguists are paid too much for sitting around thinking about words all day then (52b) is an appropriate response. On the other hand, if the interpretation of (52a) is that linguists are paid too little for their noble research, (52c) is an appropriate response. Even though each participant picks out an obvious and extreme answer, the rhetorical question in (52a) can fail if the two discourse participants have different probability distributions skewed towards opposite ends of the salary spectrum. Subsequent discourse will reveal the discrepancy, indicating a misalignment between the Speaker and Addressee's perceptions of the Common Ground.

    Non-matching skewing directions are prohibited, but non-identical answer values are fine. To be sufficiently similar, multiple answers need simply to exist on the same end of the spectrum of possible answers. As mentioned above in question (45) *How many stars are there in the sky?* the two discourse participants may identify different values for the answer, but they share a commitment to a similarly extreme answer. For rhetorical

questions with multiple answers as in (53), each participant may choose a different cluster of extreme outcomes and still satisfy the condition on sufficient similarity.

(53)   What's going to happen to these kids when they grow up?   [sw_0255_4548]
       CONTEXT: juvenile delinquents

Even if the actual members of the answer cluster overlap minimally or not at all, clusters which occupy positions at or near the same endpoint of the scale provide sufficient similarity.

get rich	< 1%
visit the Grand Canyon	< 1%
⋮	
abuse kittens	3%
go to prison	91%
blow up LA	5%

Figure 10. Speaker answer partition to (53)

get rich	< 1%
visit the Grand Canyon	< 1%
⋮	
shoplift	4%
commit a felony	90%
hijack a plane	5%

Figure 11. Addressee answer partition to (53)

In Figures 10 and 11, both the Speaker and Addressee answer sets highlight a cluster of extreme answers. Though the highlighted answers are not identical, the midpoints of the two clusters are both located near the same bottom endpoint of the scale, and this ensures sufficient similarity. Figures 12 and 13 show the mutual extremity of the high probability items, even though the actual values are non-overlapping.



Figure 12. Speaker distribution for (53)

Figure 13. Addressee distribution for (53)

Figures 12 and 13 depict similar, but not identical, skewed probability distributions over similar, but not identical, answers. As the condition in (54) states, two answer sets are similar when they share a singleton answer value or, for questions like (53) with multiple answers, when they share similarly extreme regions on a scale.

   Condition (54a) formalizes answer value equivalence, checking that the most probable answers from the Speaker and Addressee answer partitions are the same. The function arg max returns participant X's obvious answer – the cell c_i with the highest probability in X's answer partition. Condition (54b) provides a metric of similarity for questions with multiple or vague answers. Condition (54bi) requires that the answer be located sufficiently far (by some distance δ) from the middle of the scale ($\frac{n}{2}$). Condition (54bii)

(54)   CONDITION 3:   Rhetorical Questions require that Speaker and Addressee
              share **sufficiently similar** answers

Given Speaker/Addressee answer partitions, at least (a) or (b) must hold:

a. *Answer equivalence*

$$\arg\max_{c_i} \; prob_{speaker}(c_i) = \arg\max_{c_j} \; prob_{addressee}(c_j)$$

for a cell c_i or c_j in the respective participant's answer partition

b. *Mutual extremity*

for $d = \frac{n}{2} - i$,

(i) $|d_X| > \delta$

(ii) and SIGN(d_X) is the same for all X

where i is the position on the scale of the most probable answer
and n is the number of elements on the scale (in cases with mul-
tiple answers, i may be the midpoint of the highlighted region
of the scale)

ensures that both participants' answers are located at the same end of the scale (positive
or negative relative to the midpoint).

    With the constraints in (54), two kinds of questions are prevented from receiving a
rhetorical reading. The first set of questions are those that locate a cluster of obvious
answers whose midpoint is a non-extreme position.[9] The second set are those that have
more than one extreme answer but the answers are not contiguous. These are cases where
the obvious answers occupy opposite ends of the scale. They fail as rhetorical questions
because the midpoint of the dispersed cluster averages to a position on the scale which
is closer to the middle of the scale than to either endpoint. These answer sets fail the
condition on mutual extremity (54bi) even though they satisfy the obviousness condition.
An example of a question with non-contiguous answers is (55).

(55)   Wow, John lost weight! What did he do?

The idea is that a discourse participant who believes that the obvious answers to (55)
include both the safe and dangerous extremes on the dieting strategy spectrum (exercise
and starvation) will deny the question a rhetorical reading. If, however, the relevant scale
is effective dieting strategies, (55) can be read rhetorically as an assertion that whatever
strategy John used, it was extremely effective, as in (56).

---

[9]Rhetorical questions with non-extreme answers are notably absent from the corpus, but evidence
regarding the requirement on extremity should come from the infelicity of rhetorical questions with
obvious midpoint answers. Most constructed cases, however, invoke some scale on which the answers
are still extreme, making it difficult to isolate a case where the obviousness condition is satisfied but
extremity is violated.

(56)   Wow, John lost weight! What did he do?! He must have lost over 50 pounds! I've never seen him look so good. Imagine losing 50 pounds in a month!

In (56), the question *what did he do?!* is not information-seeking, and the implication is that whatever John did, it was very effective.

The existence of relevant scales for rhetorical questions is predicted to have some repercussions in the surrounding discourse. The importance of scales as a metric for answer similarity licenses references to the invoked scale, even when the scale has not been overtly specified by a polarity item or a *wh-* phrase.

(57)   a.   A: What's going to happen to these kids when they grow up?   [sw_0255_4548]
             CONTEXT: juvenile delinquents
       b.   B: Oh, they're not *that* bad.
       c.   B′: Well, they'll *at least* graduate from high school.

The salience of scales in the discourse surrounding rhetorical questions is what licenses comments like (57b and 57c). B's responses indicate that the condition for shared obvious answers has not been satisfied for A's rhetorical question. The question generates a discourse repair while participants reestablish their relevant commitments. Since rhetorical questions aim to synchronize the discourse, they may reveal discrepancies in the perceptions of the Common Ground that require repair.

5.6   Rhetorical question usage and failure

Because of the properties of answer obviousness and similarity, rhetorical questions can be considered uninformative. Their effect is commitment synchronization. Highlighting and synchronizing shared views is useful for affirming linguistic and cultural community as well as to the upkeep and maintenance of the Common Ground. Previous work on synchronization comes from (Taranto 2003) in which commitment confirmation establishes a shared perception of the Common Ground between participants.

Rhetorical questions also are often used to reiterate an agreed-upon state of the world, from which other conclusions can be drawn. In example (46) about tax rates (*How high are taxes going to be when my kids are my age?*), one can hypothesize from the obviousness of an extreme answer that the Speaker believes his children will suffer or that politicians ought to fix the problem. Since rhetorical question usage is associated with strong bias, this predicts that midpoint values on a scale will be unavailable as obvious answers because they provide less ammunition for an emotional proclamation. If the obvious answer is a non-extreme value, the question is less effective because there are fewer conclusions to draw. At least in the corpus, there is an absence of rhetorical questions that highlight the status quo (though one could construct an example where the status quo is actually the extreme value on the dimension of conformist normality).

Besides capturing the synchronization effects of rhetorical questions, the proposed analysis has an advantage over other models in that it can predict ways in which rhetorical questions are misinterpreted or result in Common Ground repairs. Analyses that do not establish Common Ground conditions have difficulty explaining failed rhetorical questions.

In cases where Speakers incorrectly assume that the obvious answer they know is also obvious to the Addressee, the question either loses its rhetorical interpretation or requires amendments to the Common Ground to fix the incorrect assumption.

(58)   Who was the only heavy weight champion to retire undefeated?

Under my analysis, question (58) is felicitous as a rhetorical question if the Speaker and the Addressee both know about famous boxers. In the case where the Addressee cannot identify the obvious answer (*Rocky Marciano*), the question serves to indicate a misalignment in the Speaker and Addressee perceptions of each other's knowledge. This kind of rhetorical failure reveals the importance of separating the Common Ground, not just into Speaker and Addressee commitment sets, but into Speaker and Addressee perceptions of other participant's commitment sets. Issues of recursion aside, participant-defined perceptions of the Common Ground are valuable for identifying cases in which the Speaker has assumed something invalid about what the Addressee knows or believes.

A question may also fail rhetorically if the Speaker and the Addressee answers are not sufficiently similar. These are cases like (59a) in which an Addressee informs the Speaker that the condition of equivalent beliefs has not been met.

(59)   a. A: Who would steal a newspaper?   [sw_0141_2060]
       b. B: Well, actually, a very shy nudist might steal a newspaper.
       c. B′: Well, actually, I once stole a newspaper to read on the train.

In (59), the B correctly identifies the A's rhetorical intent but then asserts that the condition of similar answer sets has not been met by citing evidence that individuals exist who do steal newspapers and they do so with relatively high probability. Similar examples include (52) in which individuals may place the same salary on different extremes of a scale.[10]

---

[10]The different perceptions of the same salary is a problem for my analysis. Given that similarity can be established with either answer equivalence or mutual extremity, different perceptions about the same salary satisfy answer equivalence though they violate mutual extremity.

Requiring that both conditions for sufficient similarity be met is not feasible because there are obvious answers whose extremity can only be measured on scales of *expected* answers or *normal* answers. These scales then are inextricable from the probability distribution superimposed over the scale – in which case, the obvious high probability answer will always be at the endpoint on a scale of answer probability.

Another case of a failed rhetorical question is exemplified in (60a) where the rhetorical question is interpreted as a regular information-seeking question.

(60)   a.   A: How high will taxes be when my kids are my age?   [sw_0729_4840]

   b.   B: Well, that's a great question! Let me tell you, based on the current trajectory

of income tax valuation along with the growing number of Americans on social security and Greenspan's waning confidence in the dollar, I'd say taxes are likely to increase drastically over the next thirty to thirty-five years.

The rhetorical question fails because the Addressee does not recognize the rhetorical intent of the falling interrogative and provides an (overly) informative response. This means that the Addressee has misclassified the Speaker's question and is using the context update of a rising interrogative, mistakenly treating the question as information-seeking.

These failures indicate that rhetorical question felicity is sensitive to the conditions I have proposed: answer obviousness and metrics for sufficient similarity. The failed questions also point to the importance of separating Speaker and Addressee commitments and defining Speaker/Addressee perceptions of the Common Ground.

6   Conclusion

In this paper, I have introduced naturally-occurring rhetorical questions from the Switchboard corpus that differ from those described in the literature. Previous analyses limit the meaning of rhetorical questions to negative, singleton answers; my examples include *yes/no*, *wh-*, and non-argument *wh-* rhetorical questions that have the full range of positive/negative, null/non-null, single/multiple answers. I claim that the variation seen in the corpus data merits a new analysis to define discourse conditions that permit the felicitous use of a rhetorical question. I characterize rhetorical questions using notions of bias and informativity from Gunlogson's (2001) work as well as measures of answer set predictability from van Rooy (2003).

My claim is that participants in a discourse understand rhetorical questions when they share sufficiently obvious and similar answers. Obviousness is measured by the decrease in entropy or surprise associated with the answer set. I use the same model of entropy that van Rooy (2003) uses to equate regular and rhetorical questions, but I use entropy to distinguish the high predictability of a rhetorical question answer set from the low predictability of a regular question's. Sufficient similarity is measured either through answer equivalence or through mutual extremity on a relevant scale. The formalization

---

A possible solution is the treatment of salaries, not as numbers, but as (subjective) labels like *too much* or *too little* which correctly fail the similarity metrics of answer equivalence and mutual extremity.

of obvious, shared answers to rhetorical questions captures their inherent lack of informativity. They invoke answer sets like regular questions, but they carry strong bias like statements. Their answer is predictable by both participants: in this sense, rhetorical questions are redundant interrogatives. This new approach to rhetorical questions captures a much broader collection of data and makes predictions about cases in which rhetorical questions fail. It also raises questions concerning discourse synchronization and the perceptions of other's beliefs.

6.1 Future work

This paper has examined how speakers use rhetorical questions to communicate their personal biases and what effect these questions have on the entropy, or surprise, associated with the answer set. Further work is needed to compare the varieties of rhetorical devices used in discourse and the kinds of responses they generate. In a sociolinguistic study, Ilie (1994) presents a taxonomy of rhetorical questions as they are used in court rooms and talk shows. These naturally-occurring examples, as well as those from the Switchboard corpus, embody a common theme of synchronizing the Speaker and Addressee's bias towards a shared redundant answer. These same issues of bias and answer informativity also play a role in computational tasks designed to assess speaker bias and author opinion.

A tacit assumption through this paper is the accessibility of real world knowledge and its representation in discourse participants' commitment sets. The conditions I specify for a rhetorical question require that participants share beliefs about the answer. I do not provide a mechanism for knowledge update or insertion, but assume that participants have a Common Ground of shared real world knowledge. Because I cast commitment updates as probability adjustments, I redefine the Common Ground to contain all possible worlds with probabilistic highlighting to draw attention to an individual's commitments. A more general question is what implications the role of probabilities in the Common Ground has for categorical assessments like presupposition or contradiction.

Though questions remain about the accessibility of real world knowledge and the selection of relevant scales, the proposed analysis covers a new and broader range of rhetorical question data. It attempts to capture the underlying theme of shared, extreme, and obvious answers across a variety of naturally occurring examples. This work offers a characterization of rhetorical questions within a context-update model of discourse and it isolates discourse conditions conducive to felicitous use of these questions.

REFERENCES

Banuazizi, Atissa, & Cassandre Creswell. (1999) Is that a real question? final rises, final falls, and discourse function in yes-no question intonation. *CLS* 35, 1-14.

Bartel, Christine. (1999) *The intonation of English statements and questions: a compositional interpretation.* New York: Garland.

Borkin, Ann. (1971) Polarity items in questions. *CLS* 7 , 53-62.

Fauconnier, Gilles.(1975) Polarity and the scale principle. *CLS* 11, 188-199.

Groenendijk, J., & M. Stokhof. (1997) Questions. In van Benthem and ter Meulen (eds.),

Handbook of logic and lingusitics, Amsterdam: Elsevier Science.

Guerzoni, Elena. (2002) Even-NPIs in questions. *NELS* 32, 153-170.

Gunlogson, Christine. (2001) *True to form: Rising and Falling Declaratives as Questions in English*. PhD Dissertation, UCSC.

Gutierrez-Rexach, Javier. (1996) The semantic basis of NPI licensing in questions. In Benjamin Bruening (ed.), *Proceedings of the Eighth Student Conference in Linguistics, MIT Working Papers in Linguistics* 31, 359-376.

Han, Chung-Hye. (1998) *The structure and interpretation of imperatives: mood and force in universal grammar*. Doctoral Dissertation, University of Pennsylvania.

——, & Laura Siegel. (1997) Syntactic and semantic conditions on NPI licensing in questions. *WCCFL* 15, 177-191.

Hoeksema, Jack, & Hotze Rullmann. (2001) *Scalarity and polarity*. Amsterdam: John Benjamins.

Horn, Laurence. (1989) *A natural history of negation*. Chigago: University of Chicago Press.

Ilie, Cornelia. (1994) *What else can I tell you? A pragmatic study of English rhetorical questions as discursive and argumentative acts*. PhD dissertation, University of Stockholm.

Jurafsky, Daniel, Rebecca Bates, Noah Coccaro, Rachel Martin, Marie Meteer, Klaus Ries, Elizabeth Shriberg, Andreas Stolcke, Paul Taylor, & Carol Van Ess-Dykema. (1998a) Switchboard discourse language modeling project final report. In *Large Vocabulary Continuous Speech Recognition Summer Research Workshop Technical Reports*, Research Note No. 30, Center for Speech and Language Processing, Johns Hopkins University Baltimore, MD.

——, Elizabeth Shriberg, Barbara Fox, & Tracy Curl. (1998b) Lexical, Prosodic, and Syntactic Cues for Dialog Acts. *Proceedings of ACL/COLING-98 Workshop on Discourse Relations and Discourse Markers*, 114-120.

Krifka, Manfred. (1995) The semantics and pragmatics of polarity items. *Linguistic Analysis* 25, 1-49.

——. (2001) For a structured meaning account of questions and answers. In C. Fery & W. Sternefeld (eds.), *Audiatur Vox Sapientia. A Festschrift for Arnim von Stechow*, 287-319. Berlin: Akademie Verlag.

Ladusaw, William. (1980) *Polarity sensitivity as inherent scope relations*. PhD dissertation, University of Texas.

——. (2003) Framing the issues: The biasing effect of polarity items in questions. *Proceedings of Where Semantics Meets Pragmatics*, Michigan: East Lansing.

Lawler, John. (1971) Any questions? *CLS* 7, 163-173.

Lee, Felicia Ann. (1994) *Negative-polarity licensing in wh-questions: the case for two licensers*. Master's thesis, UCLA.

Portner, Paul, & Raffaella Zanuttini. (2000) The force of negation in wh-exclamatives and interrogatives. In L. Horn & Y. Kato (eds.), *Studies in Negation and Polarity*, 193-231. Oxford: Oxford University Press.

Progovac, L. (1993) Negative polarity: entailment and binding. *Linguistics and Philosophy* 16:2.149-180.

Rooy, Robert van. (2003) Negative polarity items in questions: Strength as relevance. *Journal of Semantics* 20, 239-273.

Sadock, Jerrold. 1971. Queclaratives. *CLS* 7, 223-231.

Shannon, Claude. (1948) The mathematical theory of communication. *Bell System Technical Journal* 27, 379-423, 623-656.

Stalnaker, Robert. (1978) Assertion. *Syntax and Semantics* 9.315-332, San Diego: Academic Press.

Taranto, Gina. (2003) *Discourse Adjectives*. PhD dissertation, UCSD.

Wilkinson, Katrina. (1996) The scope of *even*. *Natural Language Semantics* 4:3, 193-215.

Author's Address:
Department of Linguistics 0108
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0108
hannah@ling.ucsd.edu

References

Banuazizi, Atissa, & Cassandre Creswell. 1999. Is that a real question? final rises, final falls, and discourse function in yes-no question intonation. In *CLS 35*, 1–14.

Bartel, Christine. 1999. The intonation of english statements and questions: a compositional interpretation. In *Outstanding Dissertations in Linguistics*. Garland Publishing, New York.

Borkin, Ann. 1971. Polarity items in questions. In *CLS 7*, 53–62.

Fauconnier, Gilles. 1975. Polarity and the scale principle. In *CLS*, volume 11, 188–199.

Groenendijk, J., & M. Stokhof. 1997. Questions. In *Handbook of logic and lingusitics*, ed. by van Benthem & ter Meulen. Amsterdam: Elsevier Science.

Guerzoni, Elena. 2002. Even-npis in questions. In *NELS 32*, 153–170.

Gunlogson, Christine, 2001. *True to form: Rising and Falling Declaratives as Questions in English*. Unviersity of California, Santa Cruz dissertation.

Gutierrez-Rexach, Javier. 1996. The semantic basis of npi licensing in questions. *Proceedings of the Eighth Student Conference in Linguistics, MIT Working Papers in Linguistics* 31.359–376.

Han, Chung-Hye, 1998. *The structure and interpretation of imperatives: mood and force in universal grammar*. University of Pennsylvania dissertation.

——, & Laura Siegel. 1997. Syntactic and semantic conditions on npi licensing in questions. In *15th West Coast Conference on Formal Linguistics*, 177–191.

Hoeksema, Jack, & Hotze Rullmann. 2001. Scalarity and polarity. John Benjamins Publishing Company.

Horn, Laurence. 1989. *A natural history of negation*. University of Chicago Press.

Ilie, Cornelia, 1994. *What else can I tell you? A pragmatic study of English rhetorical questions as discursive and argumentative acts*. University of Stockholm dissertation.

Jurafsky, Daniel, Rebecca Bates, Noah Coccaro, Rachel Martin, Marie Meteer, Klaus Ries, Elizabeth Shriberg, Andreas Stolcke, Paul Taylor, & Carol Van Ess-Dykema. 1998a. Switchboard discourse language modeling project final report. In *Large Vocabulary Continuous Speech Recognition Summer Research Workshop Technical Reports, Research Note No. 30, Center for Speech and Language Processing, Johns Hopkins University Baltimore, MD*.

——, Elizabeth Shriberg, Barbara Fox, & Tracy Curl. 1998b. Lexical, prosodic, and syntactic cues for dialog acts. In *ACL/COLING-98 Workshop on Discourse Relations and Discourse Markers*.

Krifka, Manfred. 1995. The semantics and pragmatics of polarity items. *Linguistic Analysis* 25.1–49.

——. 2001. For a structured meaning account of questions and answers. In *Audiatur Vox Sapientia. A Festschrift for Arnim von Stechow*, ed. by C. Fery & W. Sternefeld, 287–319. Akademie Verlag, Berlin.

Ladusaw, William, 1980. *Polarity sensitivity as inherent scope relations*. Austin: University of Texas dissertation.

——. 2003. Framing the issues: The biasing effect of polarity items in questions. In *Where Semantics Meets Pragmatics*, East Lansing, Michigan.

Lawler, John. 1971. Any questions? In *CLS 7*, 163–173.

Lee, Felicia Ann, 1994. Negative-polarity licensing in wh-questions: The case for two licensers. Master's thesis, UCLA.

Portner, Paul, & Raffaella Zanuttini. 2000. The force of negation in wh exclamatives and interrogatives. In *Studies in Negation and Polarity*, ed. by L. Horn & Y. Kato, 193–231. Oxford: Oxford University Press.

Progovac, L. 1993. Negative polarity: entailment and binding. *Linguistics and Philosophy* 16:2.149–180.

Rooy, Robert van. 2003. Negative polarity items in questions: Strength as relevance. *Journal of Semantics* 20.239–273.

Sadock, Jerrold. 1971. Queclaratives. In *CLS 7*, 223–231.

Shannon, Claude. 1948. The mathematical theory of communication. *Bell System Technical Journal* 27.379–423,623–656.

Stalnaker, Robert. 1978. Assertion. *Syntax and Semantics* 9.315–332.

Taranto, Gina, 2003. *Discourse Adjectives*. San Diego: University of California dissertation.

Wilkinson, Katrina. 1996. The scope of *even*. *Natural Language Semantics* 4:3.193–215.