

Decision Theory and Discourse Particles: A Case Study from a Large Japanese Sentiment Corpus

Christopher Davis

University of Massachusetts Amherst

Abstract. The distribution and use of the Japanese particle *yo* is examined using a large annotated sentiment corpus. The data is shown to support a decision-theoretic account of *yo*'s meaning (Davis, 2009). A decision-theoretic approach to the analysis of sentiment corpora is proposed, by which empirical predictions of decision-theoretic formal analyses can be tested using large sets of naturalistic data.

Keywords: discourse particles, sentiment corpora, corpus pragmatics, decision theory, Japanese

1 Introduction

There has been a recent surge of interest in the formal semantics and pragmatics literature on the topic of *discourse particles* (Zimmermann, to appear). Discourse particles straddle the border of semantics and pragmatics, and provide a perfect empirical domain for developing and challenging formal models of linguistic meaning. Discourse particles are, as the name implies, connected to the context of an entire discourse, and force the analyst to go above the sentence level and develop a theory of discourse contexts within which sentences and their associated particles are situated and interpreted.

One problem for the development of formal theories of discourse particles is the fact that they typically make no truth-conditional contribution to the sentences in which they occur, and the contribution that they do make is typically very difficult to pin down. Most formal studies of particles rely on intuitionistic data from small sets of typically constructed examples. The ineffability and extreme context-sensitivity of discourse particles make it difficult to study them using corpora and other naturalistic data, in which the analyst is unable to control the discourse context and cannot probe the often subtle speaker intuitions that guide the use of these particles.

Recently, a number of researchers have exploited large sentiment corpora to explore empirical regularities in the use of *expressives* and other emotionally-charged language (Potts and Schwarz, 2008; Constant *et al.*, 2008; Davis and Potts, to appear). The structure of these corpora has allowed researchers to explore the use of these often ineffable items using large sets of naturalistic texts, on the basis of which empirical estimates of the expressive effects of this kind of language can be made. In this paper, I expand on this line of research by showing how sentiment corpora can be used as an empirical tool in the exploration of *decision-theoretic* analyses of the semantics and pragmatics of lexical items and constructions. I focus on a particular formal analysis of the Japanese sentence final discourse particle *yo* (Davis, 2009). This analysis builds on recent developments in decision/game-theoretic semantics and pragmatics (Parikh, 2001; van Rooy, 2003; Benz *et al.*, 2005a). By testing the formal analysis with quantitative data from naturalistic texts, I demonstrate the utility of corpus methods for lexical pragmatics.

In Section 2, I outline the decision-theoretic analysis of *yo* in terms of which the corpus data is analyzed. Section 3 introduces the sentiment corpus used in this paper, and explores the distribution of *yo* across ratings categories in this corpus. I show that *yo* occurs more frequently in more extreme reviews, and argue that this distribution falls out from the semantics presented in Section 2. The data is also consistent with other analyses, in particular ones in which *yo* contributes expressive meaning by indexing speaker emotionality. In Section 4 I present data suggesting that this

alternative approach is insufficient. The discussion outlines ways in which the “expressive profiles” of lexical items in a sentiment corpus can emerge in several ways, so that the analyst must combine corpus data with other tools to arrive at the correct explanation for the distribution of a given item. In Section 5 I present evidence showing that *yo* tends to appear late in the text in which it is found, with a noticeable bias toward text-final position. I argue that this fact falls out from the way that *yo*’s denotation depends on the state of the post-update contextual common ground, rather than the information encoded by the sentence on which it appears. Section 6 concludes.

2 Formal Semantics of *yo*

The semantics of *yo* described in this paper is motivated by examples like (1), in which *yo* is used with an utterance intended to guide the behavior of the addressee. Note that the same sentence without *yo* is perceived by native speakers as being significantly less felicitous than with *yo*.

- (1) A: *tabe-te-kara eiga-o mi ni ik-ou ka na*
 eat-INF-from movie-ACC see to go-HORT Q PRT
 “I wonder if I should eat before going to the movie?”
- B: *mou 7-ji sugi deshou? eiga-wa 8-ji kara hajimaru #(yo)*
 already 7-o’clock past right movie-TOP 8-o’clock from starts #(yo)
 “It’s already 7, right? The movie starts at 8 #(yo).”
 (Davis, 2009)

In (1), A has expressed some uncertainty about whether he should eat before going to the movie. B responds with information he expects will be sufficient to make A choose not to go eat, since there is not much time left until the movie. This requires a certain type of inference from the information expressed by his utterance and various pieces of background information (how long it generally takes to eat, how long it takes to get to the theater, etc). The utterance without *yo* is felt to be infelicitous, which tells us that the use of *yo* is somehow implicated in triggering this inference.

Davis (2009) provides an analysis of *yo* motivated by examples like this one, arguing that *yo* generates a pragmatic presupposition that the utterance it attaches to is sufficient to resolve the addressee’s *decision problem*. In this paper, I adopt the denotation in (2), which is similar to that proposed by Davis (2009), but with a difference that will be important in explaining a restriction on the repeatability of *yo* to be discussed in Section 4.

- (2) a. $\llbracket yo \rrbracket(\text{CCP})(c)$ is defined iff
 $\exists a \in \mathcal{A}(c') \neg \text{OPT}(a(\text{addr}), c) \wedge \text{OPT}(a(\text{addr}), c')$,
 where $c' = \text{CCP}(c)$
- b. where defined, $\llbracket yo \rrbracket(\text{CCP})(c) = \text{CCP}(c)$

The first argument of *yo* is a context change potential (CCP), a function from discourse contexts to discourse contexts. In this paper, I will consider only assertive sentences, which are assumed to be headed by the assertive operator defined in (3).¹

- (3) $\llbracket \text{ASSERT}(p) \rrbracket = \lambda c. \text{that context } c' \text{ that is just like } c \text{ except } \llbracket p \rrbracket \in \text{CG}(c')$

A declarative sentence with propositional content *p* headed by *ASSERT* thus denotes a context update, in which an input context *c* is mapped to an output context *c'* in which *p* has been added

¹ Davis (2009) has a more nuanced view of the role of operators like *ASSERT*, that interact with the intonation of *yo* to produce a variety of update types. These details do not affect the main points in this paper, and are ignored for the sake of space and simplicity.

to the common ground of c (Stalnaker, 1978), where the common ground of a context c , $\text{CG}(c)$, is modeled as a set of propositions.

The denotation of yo takes a context change potential (CCP) and context (c) as arguments, returning a (pragmatic) presupposition. This presupposition relies on a set of contextually salient actions, $\mathcal{A}(c')$, representing the options from which our agent(s) must choose.² Formally, the alternative actions are understood as *properties*, so that for a given world w , action a , and agent x , we have $a(x)(w) = 1$ iff x chooses a in w , and $a(x)(w) = 0$ iff x does not choose a in w . In the presupposition associated with yo , we require that there is some action in this set³ which is *optimal* in the output context c' , but is *not* optimal in the pre-update context c . Optimality is defined as follows:

(4) Definition of Optimality:

$$\text{OPT}(p, c) = 1 \text{ iff } \forall w_i, w_j \in \cap \text{CG}(c) [(p(w_i) \wedge w_i <_c w_j) \rightarrow p(w_j)]$$

A proposition p is optimal relative to a context c just in case p is true of all the most highly-ranked worlds consistent with the CG in c , where ranking is determined by the contextual ordering defined in (5).

(5) Partial Ordering of Worlds [modified from (Portner, 2007)]

For all worlds $w_i, w_j \in \cap \text{CG}(c)$, $w_i <_c w_j$ iff

$$\exists p \in g(c) [p(w_j) \ \& \ \neg p(w_i) \ \& \ \forall q \in g(c) [q(w_i) \rightarrow q(w_j)]],$$

where $g(c)$ is the *ordering source* in c . (Kratzer, 1981)

The contextual *ordering source* is a set of propositions, upon which the partial order $<_c$ is defined. A world w_i is ordered above another world w_j if there is a proposition in the ordering source that holds of w_i and not of w_j , while all other propositions in the ordering source that hold of w_j also hold of w_i .

While the above semantics may seem a bit complicated, its consequences for pragmatics are easily summarized. According to (2), the use of yo presupposes that there is some action a that is salient in the post-update context and is also optimal for the addressee in that context, but not optimal in the pre-update context. Putting the pieces together, what this means is that by attaching yo to an utterance, the speaker in effect indicates that his utterance serves to resolve some choice facing the hearer which the context prior to the utterance was not sufficient to resolve.

The semantics of yo presented in this section rather directly captures the use of yo in examples like (1). In brief, (1) presents a background context in which A is trying to decide whether to eat before going to the movies. The assertion by B places a new fact into the common ground.

² We can further articulate the context by having a potentially different set of actions for each discourse participant x , $\mathcal{A}_x(c)$.

³ Note that the set of actions is defined relative to the *output* context. In example (1), where the alternative actions have already been made salient in the prior context, this set will be the same in both the input and output contexts. But examples like the following (due to an anonymous reviewer) show that yo can be used to indicate actions that were not salient in the prior discourse context:

Context: The speaker notices that the hearer has dropped his ticket, a fact that the hearer does not realize.

a, kippu-o otoshimashita yo
ah, ticket-acc dropped yo

“Oh, you dropped your ticket yo .”

What this example shows is that the use of yo does not require for its felicity a salient decision problem in the prior discourse context, and can in fact be used to introduce such an issue into a previously neutral context. By relating the set of actions to the *output* context in the presuppositional semantics of yo , we can understand examples like this through a kind of accommodation, whereby the hearer must figure out what the salient action is on the basis of what the speaker has said, along with other factors.

Without *yo*, the sentence is infelicitous; native speakers report that it sounds as if B is “just stating facts”, without expressing any connection between what he is saying and the problem faced by A. By using *yo*, B indicates that the post-update context in which his assertion has been integrated is sufficient to resolve A’s problem. This invites an inference as to why the post-update context resolves A’s decision problem, and in what direction.

In the next section, I adduce quantitative support for the analysis presented in this section, relying on data from a large Japanese sentiment corpus. As will be seen, the decision-theoretic analysis developed on the basis of hand-crafted examples like (1) receives further support from the distribution of *yo* in the corpus.

3 *yo* and Speaker Sentiment: Evidence from Sentiment Corpora

The data in this section come from a recently expanded version of the publicly available UMass Amherst Sentiment Corpora (Constant *et al.*, 2009). The Japanese portion of this corpus contains approximately 33 million words of review text culled from reviews of various products (books, DVDs, electronics, and games) appearing on the Japanese Amazon website, Amazon.co.jp. All reviews on the site are associated with a product rating given by the reviewer, ranging from 1 to 5 stars. The ratings data provide an objective scale along which the author’s *sentiment* or *evaluation* of the target product can be estimated. 1 and 5 star reviews are extremely negative and positive, respectively, while 2 and 4 star reviews are associated with more moderate negative and positive evaluations. 3 star reviews are associated with a high degree of ambivalence or lack of a strong evaluative stance with respect to the target product.

To analyze the association between specific lexical items and associated rating scores, the *relative frequency* of an item across the five rating categories is calculated.⁴ The rating categories are transformed to a *sentiment index* such that sentiment index = star rating – 3, so that a star rating of 3 maps to a centered sentiment index of 0 on the *x*-axis of the graphs to be presented. In this way, negative numbers reflect negative evaluations (1 and 2 star reviews correspond to sentiment indices of –2 and –1), and positive numbers reflect positive evaluations (4 and 5 star reviews correspond to sentiment indices of 1 and 2).

Figure 1 shows the distribution of the English expressives *wow* and *damn* along with *yo* in the review texts of the English and Japanese Amazon corpora across the five centered ratings categories. The *y*-axis plots the *log odds* of the item in the corpora. The use of log odds allows us to fit logistic regression models to the data, in order to test for the statistical significance of certain trends in the distribution of an item across rating categories. All three items have a clear U-shaped distribution across the rating categories, an impression that is confirmed by the significance of the quadratic terms in the associated quadratic logistic regression. The U-shaped distribution indicates a tendency for these items to be used in reviews whose author has a more extreme opinion toward the item being reviewed, with a correspondingly strong recommendation, whether positive or negative.

It is conceivable that expressives like *wow* or *damn* directly index speaker emotionality, in which case its distribution in the corpus is a direct reflection of its meaning, insofar as review category serves as a proxy for emotional state. This use of the sentiment data relies on a (potentially indirect and fuzzy) mapping from emotional state to sentiment index, and vice-versa. Their distribution across sentiment indices thus supports the analysis of these item as expressing heightened speaker emotionality, and at the same time provides a means for empirically estimating the degree of heightened emotion expressed by this item by comparison with other expressive items. This perspective has been adopted for the analysis of expressive items exhibiting a U-shaped distribution in these corpora (Potts and Schwarz, 2008; Constant *et al.*, 2008).

⁴ Relative frequencies are used because there is a bias toward more positive reviews in the corpus.

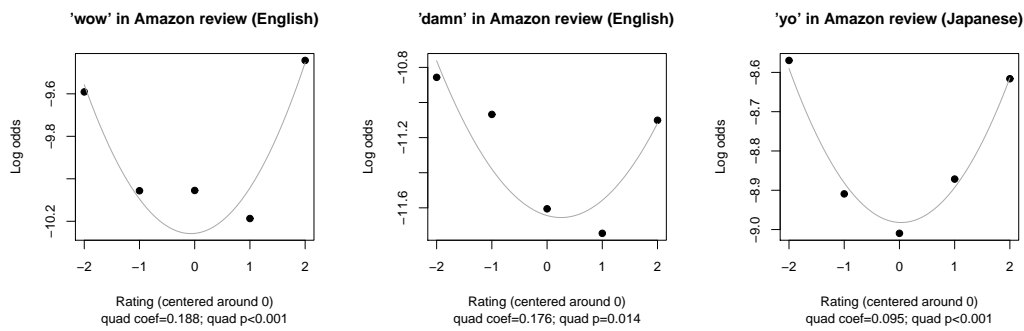


Figure 1: Distribution of *wow*, *damn*, and *yo* in the review text of the Amazon corpora.

The distribution of *yo* is also U-shaped, with more uses in the extreme rating categories than in the more moderate ones. From this we might conclude that, like *wow* and *damn*, *yo* is used to directly index speaker emotionality. I would like to suggest an alternative analysis of the data. In line with the formal semantics outlined in the previous section, I argue that the tendency for *yo* to occur in more extreme reviews is due to the fact that in such reviews the author is more confident in his recommendation that a reader either buy the product (5 star review) or not buy the product (1 star review). If *yo* conventionally indicates the existence of a contextually salient optimal action, then the bias for its occurrence in reviews with a more extreme rating follows straightforwardly.

This alternative encourages a different perspective on the sentiment indices in our corpora. In addition to correlating with author *emotionality*, these indices reflect the degree to which the author endorses the product, and hence the degree to which (s)he thinks that any given reader should purchase the product. The structure of our sentiment corpus can be mapped readily onto the formal model of decision problems and discourse contexts assumed in the analysis of *yo*. The background problem addressed by a review for item i , we can assume, is the question “Should one purchase item i ?” Formally, this is represented by a contextual action set $\mathcal{A}(c) = \{\lambda x. \text{buy}(i)(x), \lambda x. \neg \text{buy}(i)(x)\}$. With the 0-centered five star rating system, we can further assume that reviewers that give a negative (-2 or -1) rating are recommending that the reader not buy the product, while those who give a positive (1 or 2) rating are recommending that the reader buy the product. Reviewers who give a neutral (0) rating can be taken as maximally ambivalent, not making a firm recommendation either way.

Within the positive and negative review categories, we can further distinguish moderate (1, -1) and extreme (2, -2) reviews. In the context of our decision problem an extremely positive (2) review is one in which the author is fully endorsing the product, and consequently unreservedly recommending that the reader purchase the product. An extremely negative review presents the opposite extreme, in which the author strongly recommends that the reader *not* buy the product. More moderate positive and negative reviews (1 and -1) maintain a bias on the part of the author toward purchasing or not purchasing the product, but this recommendation is tempered, with potential reservations about the product not allowing for an unreserved recommendation either way.

We can now link the structure of the reviews in our sentiment corpus data to the denotation of *yo* in (2). The review is written in a context with a highly salient decision problem: whether or not to buy the product being reviewed. Reviewers in the more extreme categories have more extreme views on this issue, and are correspondingly more likely to make a categorical suggestion that the reader either purchase or not purchase the product. This fact is reflected in the relative frequency with which *yo* is used in these reviews. Informally, we can understand the U-shaped distribution of *yo* as resulting from the fact that more extreme reviews make stronger recommendations than

moderate reviews, and that *yo* tends to occur in contexts where the speaker is making a strong recommendation.

While I have shown how the distribution of *yo* in our sentiment corpus is consistent with the decision-theoretic account of *yo* presented in the last section, the data do not distinguish this account from an expressive one in which *yo* serves to index speaker emotionality. In the next section, I point out a crucial difference between the use of *yo* and that of canonical expressive items like *damn*. This difference in behavior falls out from the account presented in the last section, providing further support for modeling the U-shaped distribution of *yo* in our corpus decision-theoretically.

4 (Non-)Repeatability

While space limitations prevent me from exploring the host of ways that *yo* differs from a “pure” expressive like *damn*, the following contrast will serve to illustrate the need for a distinct analysis. Potts (2007) posits *repeatability* as one of the characteristics of expressive items:

- (6) *Repeatability*: If a speaker repeatedly uses an expressive item, the effect is generally one of strengthening the emotive content, rather than one of redundancy.

Potts illustrates this characteristic of expressive items with the following examples; as we move down the list, the repetition of *damn* serves to strengthen the sense of speaker emotionality:

- (7) a. Damn, I left my keys in the car.
 b. Damn, I left my damn keys in the car.
 c. Damn, I left my damn keys in the damn car.

Turning to *yo*, syntax prevents us from using the particle more than once in a single sentence. The principle of repeatability should apply across sentences as well, however. The following example illustrates the fact that, in general, the repetition of *yo* is *not* allowed across sentences, at least when those sentences are (in a sense to be made more explicit shortly) “about the same thing”.⁵

- (8) Context: A sushi chef is making recommendations to a customer. He makes the following two utterances, (implicitly) suggesting that the customer purchase the sea urchin.
- a. kyou-wa uni-ga oishii desu *yo*
 today-TOP sea.urchin-NOM delicious be *yo*
 “We have good sea urchin today *yo*.”
- b. kesa Hokkaido-de toreta mono desu *yo*
 this.morning Hokkaido-at caught thing be (#*yo*)
 “It was caught in Hokkaido this morning (#*yo*).”

The example illustrates the following principle: When *yo* is used with an utterance to suggest to the addressee some action, it cannot in general be used again with a subsequent utterance that is used to suggest the same action. For the example above, the action suggested to the customer by both sentences is ordering the sea urchin. It is fine to mark the first sentence with *yo*, but then the second one cannot also be so marked. The use of *yo* is thus, in an important sense, *not* repeatable, and contrasts in this respect with an expressive item like *damn*.

The decision-theoretic of *yo* can explain the restriction seen in (8). The explanation goes like this: By using *yo* with the first sentence, the speaker is suggesting that there is some salient action that is optimal for the hearer in the *post*-update context that was not optimal in the *pre*-update

⁵ I thank an anonymous reviewer for bringing this example to my attention.

context. The salient action is naturally interpreted as ordering sea urchin. The second utterance is made in the context generated, in part, by the first utterance. Using *yo* in the second utterance commits the speaker to the existence of some action that is salient and optimal in the new post-update context, but not optimal in the input context. But since the second utterance is about sea urchin, the most natural interpretation is that the utterance is still suggesting that the hearer buy sea urchin. But this was already an optimal action in the input context, due to the use of *yo* in the prior utterance. So the presupposition of *yo* is not satisfied, and the second utterance is infelicitous with *yo*.

These facts support the decision-theoretic analysis account of *yo*'s corpus distribution. While the usage profile of *yo* matches a canonical expressive like *wow* or *damn*, non-repeatability provides some reason to think that this profile is generated in a distinct way. For both types of items, we see a systematic distributional effect in our corpus, but the explanation for that effect is different. Expressive like *damn* index speaker emotionality directly, and rating category correlates (by hypothesis) with this emotional index. The particle *yo*, by contrast, serves as a guide to optimal action, and this is also reflected systematically in the rating category of the review.

In the next section of the paper, I explore one further aspect of the decision-theoretic account of *yo*: its context-dependency. I show that *yo* tends to be used later in the text of a review, and suggest how this fact might arise from the semantics given to the particle in this paper.

5 Sentence Final, Discourse Final

The particle *yo* is syntactically restricted to matrix clause-final position. Examination of the corpus data shows a tendency for *yo* to appear text-finally as well. In this subsection, I present statistical evidence from the sentiment corpus supporting this generalization. I then discuss the way in which this empirical generalization fits within the theory of *yo* outlined above.

To explore the textual position of *yo*, I extracted from the Japanese Amazon corpus every review containing one or more instances of a matrix, sentence-final use of *yo*. This excludes uses of *yo* in quotative contexts, as well as cases where *yo* is followed by another particle; such cases do not fall within the analysis presented in this paper.⁶ A total of 4,486 reviews were found containing such tokens of *yo*, containing a total of 5,283 tokens. The textual position of each token of *yo* was then calculated by counting the number of characters that preceded *yo* in the text. For a given review text, we can then get the textual position of *yo* by dividing the textual position of *yo* by the total number of characters in the text, to get a value between 0 and 1.⁷

The sentence-finality of *yo* introduces a confound in the calculation of textual position described above. To illustrate, consider a subset of reviews consisting of just two sentences of roughly equal length. Syntactically, *yo* can only occur at the end of the first sentence, or at the end of the second sentence. If it occurs at the end of the first sentence, its textual position will be approximately 0.5, or halfway through the text. If it occurs after the second sentence, its textual position will be 1. If *yo* occurs equally often on the first or second sentence in such reviews, then the average textual position will come out to 0.75. The sentence-finality of *yo* has introduced a bias towards occurring later in the text, which has nothing to do with discourse or text-level constraints on the use of *yo*.

To eliminate this confound, I calculated a corrected textual position for each occurrence of *yo* using the following procedure: I calculated the average sentence length in a review, then subtracted half of the average sentence length from the character position of each occurrence of *yo* in that review. In the example outlined above, this would give corrected textual positions of 0.25 for a

⁶ In particular, the particle sequence *yo ne* is excluded from consideration.

⁷ For technical reasons involving text processing unicode characters, the values were actually calculated in terms of bytes rather than characters. This difference does not introduce any significant differences from the idealization of the calculation given in text.

token of *yo* occurring after the first sentence, and a value of 0.75 for a token occurring after the final sentence. The corrected average textual position for a set of two-sentence reviews with an equal likelihood of *yo* after either sentence would tend toward a mean corrected position value of 0.5.

The graph in Figure 2 shows a histogram and estimated density plot of the corrected textual position of *yo* in the corpus. The mean value of the corrected textual position is 0.6, with a median of 0.67. Even with the corrected positional values, it is clear that there is bias toward later positions in the text, with a highly skewed distribution of values. This distribution can be compared with that of the question particle *ka* and the discourse particle *ne*, both of whose syntactic distribution is similar to that of *yo*, in that they must appear sentence-finally.⁸ The estimated densities for these particles across textual positions were calculated using the same procedure as described for *yo*. The mean corrected textual position of *ka* is 0.49, with a median value of 0.51. The mean value for *ne* is 0.52, with a median value of 0.55. As can be seen from the graph in Figure 2, neither particle is as biased toward text-finality as *yo*, although *ne* seems to exhibit a slight bias in the same direction, for reasons I do not know.

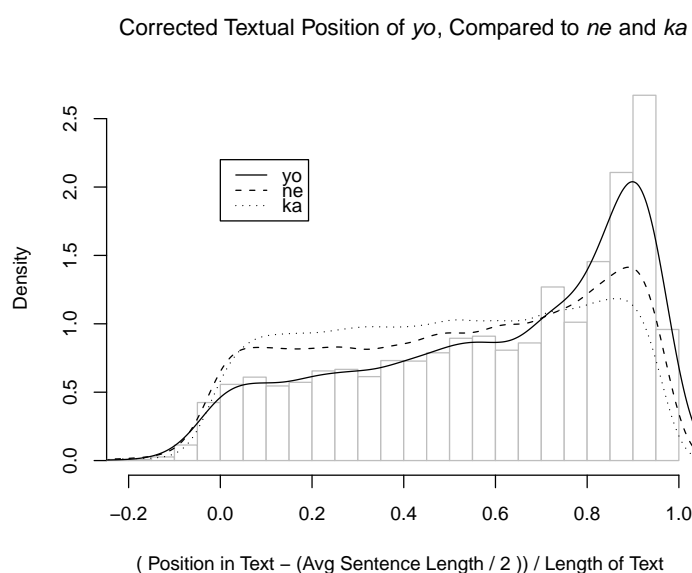


Figure 2: Histogram and density plot showing the density of the corrected position of *yo* at different points in the review text. Density estimates for two other sentence final particles are provided for comparison.

The empirical tendency for text-finality of *yo* follows from its semantics when we make a few idealizations about the structure of the review texts and the rhetorical strategies adopted by authors. In the case of extremely favorable or extremely negative reviews, we can assume that most or all of the sentences in the review will be positive or negative, respectively. In the case of a 5-star review, for example, we expect a text whose sentences are uniformly positive with respect to the product. Each sentence provides a fact or sentiment that supports the conclusion that one should buy the product. The first sentence in the review is made in a null context, and adds a single fact or sentiment relevant to the question of whether to buy the product. This adds a piece of information relevant to this decision, intended to sway the reader toward buying the product. The next sentence is made in the (positive) context created by the previous sentence. If this sentence is also positive, we now have a context with *two* pieces of information supporting the conclusion advocated by the

⁸ Like *yo*, *ne* is restricted to matrix clause-final position, while *ka* can appear in embedded clauses. In making my calculations, I considered only those instances of *ka* that appeared matrix clause-finally.

author. And so on. When the author uses *yo*, they indicate that the issue has now been settled. Rhetorically, it makes sense to save this sort of move for last.

An example will help to illustrate the principle. The text in (9) is an English translation of the review text from a 5-star review of the children's book *Hyakkai-date no ie* "The One-Hundred Story House". I have numbered the sentences to aid the discussion. The review consists of three sentences. The first sentence is neutral, describing the situation that led her to buy the book. The second sentence provides a positive sentiment. This is followed up by another sentence expressing an additional fact that supports the positive sentiment expressed by the review. This final sentence is the one that is marked by *yo*.

- (9) [1] When I asked my 1st grade child what book he would like for reading over the summer, he answered "The One-Hundred Story House", so I promptly searched for and bought it. [2] Watching my son reading it over and over, I felt glad for buying it. [3] My four year old daughter also listens to the story enthusiastically *yo*.

Why should *yo* be used to mark the last sentence in this review? I propose that the reason relies on the cumulativity of contextual update. In the context of sentence [2], we have no other positive pieces of information, while sentence [3] contains as background information the positive sentiment expressed by sentence [2]. The result is that the post-update context after sentence [3] supports the author's conclusion to a greater degree than the post-update context after sentence [2]. The presupposition of *yo* refers to the degree to which the entire post-update common ground supports a particular action, rather than the degree to which the *yo*-marked sentence itself supports the conclusion. The more positive sentences that have been asserted, the greater the degree to which the common ground supports the positive conclusion "buy the product", and thus the greater the degree to which it supports the felicitous use of *yo*.

The same holds in the case of negative reviews, where the more negative sentiments that have been expressed, the greater the degree to which the common ground supports the negative conclusion "do not buy the product". This illustrated by the 1-star review of a video game strategy guide in (10).

- (10) [1] It's just a "dictionary" in which the data from the software has been put on paper. [2] No art, no effort; I had been looking forward to it, but was really disappointed. [3] With simple data like this, it's easier just to check a wiki or something. [4] It would be nice if there were advice about weapons and armor, different ways to play, or strategies for difficult quests, but when I read this there was nothing interesting. [5] For people with an internet connection, there is absolutely no need for this book. [6] A book made like this is behind the times *yo*.

The review consists of six sentences, all of which are highly negative. The entire review builds a strong case for not buying the book, which is emphasized by the final use of *yo*. Notice that the information provided by sentence [6], on which *yo* occurs, does not seem to be a stronger strike against the book than any of the other negative sentences [1-5]. The fact that *yo* occurs with this sentence is not because it expresses a more negative or powerful argument against purchasing the book than the other sentences. Instead, its text final position follows from the fact that it is here that the argument against buying the product is strongest, since it contains all of the negative information in sentences [1-6].

The distributional data and examples provided in this section show that *yo* tends to occur late in the review text. This tendency follows from the strongly context-oriented denotation presented in Section 2. As an author builds a case for a position, the common ground becomes more supportive of that position. Since *yo* requires that the common ground be sufficient to make a particular action optimal, it tends to occur later in a text, when the context has been enriched with enough

information to favor one action over another. Looking at things from the other direction, once an author has used *yo*, he has rhetorically indicated that he takes the issue to be settled. Such an issue-settling move, I suggest, tends to be made text-finally.

6 Conclusion

In this paper, I used data from a large sentiment corpus to explore a decision-theoretic account of the Japanese discourse particle *yo*. I showed how the structure of sentiment corpora can be mapped onto a decision-theoretic model of discourse contexts, and argued that this structure is consistent with a decision-theoretic account of the particle. The distribution is, however, consistent with an expressive analysis as well, so that multiple lines of evidence are needed in order to get at the right account. Two additional pieces of evidence were adduced to this end. The non-repeatability of *yo* was explained in terms of the decision-theoretic account, as was the tendency toward text-finality in the sentiment corpus. I hope to have shown how data from sentiment corpora can be combined with other data in developing decision-theoretic models of meaning on a firm empirical foundation.

References

- Benz, A., G. Jäger, and R. van Rooy. 2005. *Game Theory and Pragmatics*. Palgrave Macmillan.
- Benz, A., G. Jäger, and R. van Rooy. 2005. *An Introduction to Game Theory for Linguists*, chapter 1. *Game Theory and Pragmatics*. Palgrave MacMillan, Houndsmills, Basingtoke, Hampshire.
- Constant, N., C. Davis, C. Potts, and F. Schwarz. 2009. The pragmatics of expressive content: Evidence from large corpora. *Sprache und Datenverarbeitung* 33(1-2).
- Constant, N., C. Davis, C. Potts, and F. Schwarz. 2009. UMass Amherst sentiment corpora.
- Davis, C. 2009. Decisions, dynamics, and the Japanese particle *yo*. *Journal of Semantics* 26(4):329-366.
- Davis, C. and C. Potts. To appear. Affective demonstratives and the division of pragmatic labor. In *Proceedings of the 17th Amsterdam Colloquium*.
- Kratzer, A. 1981. The notional category of modality. In *Words, Worlds, and Contexts. New Approaches in Word Semantics*, pages 38–74. de Gruyter, Berlin.
- Liu, B. 2010. *Sentiment Analysis and Subjectivity*. Handbook of Natural Language Processing, Second Edition.
- Parikh, P. 2001. *The Use of Language*. CSLI, Stanford.
- Portner, P. 2007. Imperatives and modals. *Natural Language Semantics*, 15(4):351–383.
- Potts, C. 2007. The expressive dimension. *Theoretical Linguistics*, 33(2):165–197.
- Potts, C. and F. Schwarz. 2008. Exclamatives and heightened emotion: Extracting pragmatic generalizations from large corpora. Ms., UMass Amherst.
- Stalnaker, R. 1978. Assertion. *Syntax and Semantics*, 9:315–332.
- van Rooy, R. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy*, 26:727–763.
- Zimmermann, M. To appear. Discourse particles. In Claudia Maienborn, Klaus von Heusinger, and Paul Portner, editors, *Semantics: An International Handbook of Natural Language Meaning*. Mouton de Gruyter, Berlin.