UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Rational use of prosody predicts projection in manner adverb utterances

Permalink

https://escholarship.org/uc/item/8p82k85g

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Stevens, Jon Scott
de Marneffe, Marie-Catherine
Speer, Shari R.
et al.

Publication Date

2017

Peer reviewed

Rational use of prosody predicts projection in manner adverb utterances

Jon Scott Stevens, Marie-Catherine de Marneffe, Shari R. Speer, Judith Tonhauser

{stevens.400,demarneffe.1,speer.21,tonhauser.1}@osu.edu

Department of Linguistics, The Ohio State University, Columbus, OH 43210 USA

Abstract

Speakers can be taken to be committed to utterance content even when that content is contributed in the scope of an entailment-canceling operator, like negation (e.g., Chierchia & McConnell-Ginet, 1990). We develop a probabilistic model of this phenomenon, called 'projection', that relies on the prosodic realization of utterances. We synthesize existing theoretical claims about prosody, information structure and projection into a model that assumes a rational speaker (Frank & Goodman, 2012) who produces utterances with prosodic melodies that can signal which utterance content she is committed to. Predictions of the probabilistic model are compared to the responses of an experiment designed to test the effect of prosody on projection in manner adverb utterances. Key behaviors of the model are borne out empirically, and the quantitative fit is surprisingly good given that the model has only one free parameter. Our findings lend support to analyses of projection that are sensitive to the information structure of utterances (e.g., Simons, Beaver, Roberts, & Tonhauser, 2017).

Keywords: Projection; prosody; information structure; probabilistic pragmatics; rational speech acts; manner adverbs

Introduction

Projective content is utterance content that the speaker may be taken to be committed to even when the content is introduced by an expression in the scope of an entailment-canceling operator (e.g., Chierchia & McConnell-Ginet, 1990). A speaker who utters (1a) is taken to be committed to the content ϕ that Sam is ill and to the content that Jo discovered ϕ. A speaker who utters the negated variant in (1b) or the polar question in (1c) is not taken to be committed to Jo having discovered ϕ (rather, Jo's discovery is negated or asked about), but the speaker may still be taken to be committed to ϕ, that Sam is ill. Hence, ϕ is projective content.

(1) a. Jo discovered that Sam is ill.
    b. Jo didn't discover that Sam is ill.
    c. Did Jo discover that Sam is ill?

Simons et al. (2017) develop a question-based analysis of the projection of the content of the complement of predicates like *discover* according to which the speaker can be taken to be committed to the content of the complement if it is entailed by the Question Under Discussion (Roberts, 2012) that the utterance is taken to address. This analysis correctly predicts that whether the speaker is taken to be committed to the content of the complement depends on the prosodic realization of the utterance (see Tonhauser, 2016, for empirical evidence). For instance, if (1c) is uttered with prosodic prominence on *discover*, as in (2a), where capital letters indicate prosodic prominence, then the speaker is more likely to be taken to be committed to the content of the complement than if (1c) is uttered with prosodic prominence on *ill*, as in (2b).

(2) a. Did Jo DISCOVER that Sam is ill?

b. Did Jo discover that Sam is ILL?

This paper provides conceptual and empirical support for the question-based analysis of projection from a novel empirical domain: utterances with manner adverbs, like *Masha didn't run quickly*. To formalize the link between prosody, information structure and projection hypothesized in Simons et al. 2017, we develop a Rational Speech Act (RSA) model (Bergen & Goodman, 2015; Frank & Goodman, 2012, and many others) that predicts the projectivity of the so-called prejacent (that Masha ran) from the prosodic realization of manner adverb utterances. The model is evaluated based on empirical observations about the projectivity of the prejacent.

Projective content in manner adverb utterances

In a manner adverb sentence, e.g., (3), a manner adverb like *quickly, beautifully* or *easily* modifies the activity-denoting verb. (3) entails that Masha ran (the prejacent). It has been long observed that the prejacent can project and that the projectivity of the prejacent depends on the prosody of the utterance (e.g., Abrusán, 2013; Simons, 2001). For instance, the speaker of (4a), with *quickly* prosodically prominent, may be taken to be committed to the prejacent even though the sentence is negated. On the other hand, the speaker of (4b), with *Masha* prosodically prominent, is not typically taken to be committed to the prejacent, i.e., the prejacent does not project. The speaker of (4b) may be committed to a different content, namely that somebody ran quickly. For reasons of space we only consider the prejacent here, though the formal proposal extends to the projectivity of other content.

(3) Masha ran quickly.
(4) a. Masha didn't run QUICKLY.
    b. MASHA didn't run quickly.

Simons et al.'s (2017) question-based analysis of projection accounts for the dependence of the projectivity of the prejacent on prosody based on two independently made empirical observations. The first observation is that information-structural focus can be prosodically indicated in American English (e.g., Eady & Cooper, 1986). Compared to non-focused expressions in the same intonational phrase, focused expressions are more likely to be realized with pitch accents, a longer duration, an expanded pitch range and/or greater intensity. Thus, the two manner adverb utterances in (4) are compatible with different expressions being focused, as shown in (5), where focus is marked by angle brackets subscripted with 'F':

(5) a. Masha didn't run [QUICKLY]$_F$.
    b. [MASHA]$_F$ didn't run quickly.

The focused expression is used to calculate the focus alternatives set of an utterance: following Rooth 1992, a focus alternatives set of (5a) is the set of propositions of the form 'Masha didn't run x', with x a modifier, and a focus alternatives set of (5b) is the set of propositions of the form 'x didn't run quickly', with x an entity. Following Beaver and Clark 2008, we assume that alternatives can also be calculated under negation: thus, another focus alternatives set of (5a) is the set of propositions of the form 'Masha ran x', with x a modifier, and another focus alternatives set of (5b) is the set of propositions of the form 'x ran quickly', with x an entity.

Importantly, the prosody of utterances merely constrains what is focused but does not determine it. Thus, although the prosodic realization of (4a) is not typically compatible with the focus marking indicated in (5b), it is compatible with a larger expression of the sentence being focused, as in (6).

(6)   Masha didn't [run QUICKLY]$_F$.

The second empirical observation is that answer utterances are congruent with the interrogative utterances they address (e.g., Paul, 1880; Rooth, 1992; von Stechow, 1990). Consider B's utterance in (7a), and assume that B's utterance is realized with the so-called rise-fall-rise contour, which consists of a rising pitch accent on *quickly* (L+H* in the ToBI annotation scheme, Beckman & Ayers, 1997) and a rising end contour (L-H% in ToBI notation). When B's utterance in (7a) is realized with the rise-fall-rise contour, it is congruent with, i.e., judged to be acceptable in response to, A's interrogative utterance in (7a). However, because of the placement of the pitch accent on *quickly*, it is not congruent with A's interrogative utterance in (7b). Likewise, when B's utterance in (7b) is realized with the rise-fall-rise contour, it is congruent with A's utterance in (7b), but not with A's utterance in (7a).

(7)   a. A: How did Masha run?
           B: Masha didn't run [QUICKLY]$_F$.
        b. A: Who ran quickly?
           B: [MASHA]$_F$ didn't run quickly.

When B's utterances are realized with the rise-fall-rise contour, B indicates that their utterance does not provide a complete answer to the question (see e.g., Wagner, 2012); instead, B's utterance only eliminates a possible true answer to the question. A prosodically motivated model of projection must take these contributions of contours into account.

In alternative semantics, question-answer congruence is accounted for by assuming that a focus alternatives set of a congruent answer includes the denotation of the question, which is a contextually restricted set of propositions. For instance, B's answer in (7a) is congruent with A's interrogative utterance since the focus alternatives set 'Masha ran x' includes the set of propositions denoted by A's utterance (a subset of propositions of the form 'Masha ran x'), but B's answer is not congruent with A's interrogative utterance in (7b) since the focus alternatives set 'Masha ran x' does not include the set of propositions denoted by A's utterance (a subset of propositions of the form 'x ran quickly').

Importantly, in naturally occurring discourse, many utterances are not made in response to an interrogative utterance (an explicit question). Simons et al. (2017) assume that such utterances address an implicit question: given question-answer congruence, the focus marking of an utterance that addresses an implicit question provides a cue to the question that the utterance addresses (e.g., Halliday, 1967; Most & Saltz, 1979; Roberts, 2012). The potentially implicit question that is addressed by an utterance is called the Question Under Discussion (QUD). It follows that the prosodic realization of a manner adverb utterance provides listeners with a cue to the focus of the utterance, which in turn provides a cue to the QUD that the speaker was intending to address. Thus, (5a) can be taken to address the QUD 'How did Masha run?' and (5b) can be taken to address the QUD 'Who ran quickly?':

(8)   a. Possible QUD of (5a): {Masha ran quickly, Masha ran slowly, Masha ran clumsily,... }
        b. Possible QUD of (5b): {Masha ran quickly, Jack ran quickly, Sue ran quickly,... }

To predict projection, Simons et al. (2017) make the following assumption:

(9)   **Projection under the question-based analysis:** Content ϕ of utterance U projects if ϕ is entailed by each alternative in the QUD addressed by U.

The question-based analysis predicts that the prejacent is more likely to project from (4a) than from (4b). If a speaker utters (4a), with prosodic prominence on the manner adverb, and the utterance is taken to have the focus shown in (5a), then she can be taken by the listener to intend her utterance to address the QUD 'How did Masha run?' in (8a). Since each alternative in the QUD has the form 'Masha ran x', each entails that Masha ran and so the prejacent is predicted by (9) to project. On the other hand, if a speaker utters (4b), with prosodic prominence on the subject, and the utterance is taken to have the focus shown in (5b), she can be taken by the listener to intend her utterance to address the QUD 'Who ran quickly?' in (8b). Since each alternative in the QUD has the form 'x ran quickly', the QUD entails that somebody ran quickly, but not the prejacent, that Masha ran. Thus, the prejacent is not predicted by (9) to project from (4b). Importantly, since prosody does not determine focus, but merely provides listeners with a cue, and since focus does not determine the QUD, but merely provides listeners with a cue, the question-based analysis does not predict categorical (non-)projection of the prejacent from (4a) and (4b), but merely that the prejacent is more likely to project from (4a) than (4b).

Modeling projection

For the purpose of modeling the link between prosody and projection, we consider an utterance to be a sentence with a melody. For the utterances considered here, a melody is the combination of a single pitch accent (L+H*, H*, L*), aligned with the stressed syllable of the accented word, and an end contour (L-H%, H-L%, L-L%). Our model considers pitch

accent positions either on the adverb, the verb, the negated auxiliary or the subject of the sentence. For example, L+H* L-H%, with the L+H* realized on the first syllable of *quickly* may be the melody of the manner adverb utterance in (4a). We develop a probabilistic model of projection based on the assumption of a 'rational speaker', who chooses a melody for a given sentence to most effectively signal the identity of the QUD that an utterance of the sentence is intended to address. Following the basic framework of the RSA model, we begin by developing a notion of speaker *utility*, which is taken to be the 'usefulness' of a melody for a given sentence as a signal of which QUD is being addressed by the speaker's utterance. More concretely, the utility of a melody M for a given sentence and a target QUD Q, is the probability that a hearer would randomly select Q, given that the sentence was uttered with melody M, and given what the hearer knows about the compatibility between M and possible QUDs. We define this utility as follows, where Q_M is the set of possible QUDs addressed by the given sentence that are *compatible* with M:

(10) $U(Q,M) = \frac{1}{|Q_M|}$ if $Q \in Q_M$, else 0

To define what it means for a melody for a given sentence to be *compatible* with a QUD, we draw on the following independently motivated sets of assumptions from the prosody-pragmatics literature:

- Assumption A – Compatibility of pitch accent with focus: A pitch accent on expression X is compatible with focus on X or a constituent that contains X and an immediately adjacent expression (e.g., Selkirk, 1996).
- Assumption B – Compatibility of focus with QUD: The QUD that is addressed by the utterance must be congruent with the focus marking of the utterance (Beaver & Clark, 2008, 45).
- Set of assumptions C – Compatibility of pitch contour with QUD:

  1. Pitch contours that lack a final fall (L-L%) and contain either an L+H* pitch accent or a continuation rise (L-H%)—H* L-H%, L+H* L-H%, L* L-H% and L+H* H-L%—are only compatible with *incomplete* answers to the QUD, i.e., answers that do not pick out a single true alternative (see e.g. Lai, 2012; Wagner, 2012, for evidence that L+H* and L-H% signal that the QUD has not been completely answered).
  2. Pitch contours that either (i) have a final fall, or (ii) have neither an L+H* pitch accent nor continuation rise to suggest incompleteness—H* L-L%, L+H* L-L%, L* L-L%, H* H-L% and L* H-L%—are only compatible with *complete* answers to the QUD.

For any melody-QUD pair $\langle M, Q \rangle$ for a given sentence, these assumptions can be used to generate the set Q_M and therefore determine the utility of using M to signal Q.

Adopting an RSA-based view, we posit a rational speaker who chooses melodies to maximize utility, i.e., maximize the chance that listeners retrieve the QUD intended by the speaker, though the maximization is *approximate*, i.e., there still remains some probability of choosing a non-rational (non-utility-maximizing) melody. This is accomplished by setting the probability P_S of the speaker producing a melody M given a QUD Q equal to a *soft max* function of $U(Q,M)$. The soft max function approximates utility maximization using a *rationality parameter*, λ, where higher values of λ result in lower probability of a non-rational melody being chosen. Thus, in cases where there is a single utility-maximizing melody, the probability of selecting that melody will approach 1 as λ increases. The formula for P_S is given below, where M' is any member of the set of possible melodies the speaker could use:

(11) $P_S(M|Q) = \frac{e^{\lambda U(Q,M)}}{\sum\limits_{M'} e^{\lambda U(Q,M')}}$

We use Bayes' rule to determine the probability P_H of the hearer deciding that the QUD is Q given that she has heard the sentence uttered with melody M:

(12) $P_H(Q|M) = \frac{P_S(M|Q) \times P(Q)}{\sum\limits_{Q'} P_S(M|Q') \times P(Q')}$

The denominator in this equation is a sum of probabilities over all possible QUDs Q', i.e., the set of QUDs that are compatible with any of the melodies we assume could have been used to utter the sentence. For instance, given the sentence *Masha didn't run quickly*, the set of all QUDs compatible with some melody for that sentence includes 'Who ran quickly?', 'Did Masha run quickly?', 'What did Masha do?', 'What did Masha do quickly?' and 'How did Masha run?', as well as the corresponding QUDs with negation ('Who didn't run quickly', etc.). For current purposes we assume a uniform prior probability distribution over QUDs.

We use P_H to calculate the probability that the prejacent of a manner adverb sentence uttered with melody M projects, i.e., that the speaker is taken to be committed to the prejacent. Recall that under assumption (9) from Simons et al. 2017, content projects if it is entailed by the QUD. For manner adverb sentences, the only QUD that entails the prejacent is the set of alternatives obtained by abstracting over the manner adverb (e.g., {Masha ran x | x is a modifier} entails that Masha ran, as discussed above). We call the prejacent ϕ, and the prejacent-entailing QUD Q_ϕ. The probability of ϕ projecting, given melody M, is the probability of the hearer assuming Q_ϕ given M:

(13) $P(\text{PROJECT}(\phi)|M) = P_H(Q_\phi|M)$

The link between Q_ϕ and projection is not probabilistic—it is a categorical consequence of the theory set forth in Simons et al. 2017 (see (9)). The probabilistic character of the model results from the fuzzy link between prosodic melodies and the implicit questions that utterances with those melodies are taken to address. The hearer must determine how likely it is that Q_ϕ is the intended QUD, based on how the speaker selects melodies to convey the QUD she intends to address.

This model operates over the possible melodies that the speaker could use and considers those possibilities when calculating the probability of the projection of the prejacent. To test the model, we examine the model's predictions for two particular melodies and compare those predictions to experimental results. The next section provides information on the two melodies and how their effect on the projection of the prejacent was assessed experimentally.

Experiment methodology

Using the method of Tonhauser 2016, participants listened to audio recordings of manner adverb utterances and judged whether the speaker was certain of the prejacent.

Participants. We recruited 100 self-reported native speakers of American English on Amazon's Mechanical Turk platform.

Stimuli. Each participant listened to 16 utterances—10 target utterances and 6 fillers. The target sentences were all of the form, "*subject* didn't *verb adverb*", where each adverb was a manner adverb, and where each subject was a proper name. The target sentences were:

(14) a. Amanda didn't clap loudly.
      b. Jennifer didn't drive carelessly.
      c. Elizabeth didn't leave silently.
      d. Linda didn't write neatly.
      e. Susan didn't sing beautifully.
      f. Jerry didn't knock frantically.
      g. Justin didn't smile cheerfully.
      h. Alexander didn't sneeze softly.
      i. Tyler didn't lie deliberately.
      j. Dennis didn't win easily.

The filler items were:

(15) a. Sandy wasn't invited to the party.
      b. Did Mario bring a chocolate cake?
      c. Who knows if Maggie is at the party?
      d. Mike forgot to bring the ketchup.
      e. Paul loves that pie!
      f. Mandy was out gardening in the yard.

Each target sentence was uttered with one of two possible melodies, L+H* L-H% with the pitch accent on the adverb (the **LH-Adverb** condition), or L+H* L-H% with the pitch accent on the proper name subject (the **LH-Name** condition). Fillers were pronounced with pitch accents on words other than the subject noun, and contained a variety of pitch accent and end tone types not used in the target utterances.

In addition to exposing participants to melodies other than the ones used for the target items, filler utterance-question pairs tested whether participants comprehended direct consequences of an action described by the speaker. For example, the utterance, *Mike forgot to bring the ketchup* was followed by the question, 'Is Debby certain that Mike brought the ketchup?'. Given that Debby's utterance implies that Mike forgot the ketchup, participants were expected to rate the speaker as "not certain", but to reliably respond in this


Figure 1: A screenshot of one experimental item.

way requires attention to the meaning of the utterance as well as its component words.

Each participant was assigned to one of two lists, where the two lists contained the same sentences but were counterbalanced for prosodic condition. The same 6 fillers occurred on both lists.

Procedure. Participants were instructed to imagine themselves at a party, where they overhear Debby, the host, utter various sentences to somebody else. For each of the 16 utterances, the participant was asked to rate on a 7-point Likert scale labeled at 4 points (1/"No, not certain", 3/"Possibly not certain", 5/"Possibly certain", 7/"Yes, certain") whether Debby was certain of some content based on what she said and how she said it. On each trial, participants were presented with a display as in Fig. 1, clicked the audio icon, heard an utterance, read the related question, clicked on the radio button that corresponded to their chosen response, and clicked the 'continue' button to proceed to the next trial. For the target item shown in Fig. 1, the utterance was *Amanda didn't clap LOUDLY*, and the participant was asked to rate the speaker's certainty about the prejacent, i.e., 'Amanda clapped'.

Data exclusion. If participants answered more than one filler incorrectly (an answer greater than 3 on the Likert scale for something that Debby would be uncertain about or an answer smaller than 5 for something that she would be certain about), their responses were excluded from analysis. We excluded 28 participants on these grounds, leaving 72 participants whose responses we analyzed. Whether these participants are excluded does not change the main effect of condition on response.

Model predictions

If participants take Debby to be committed to the truth of ϕ, we expect them to respond that Debby is certain that ϕ. If Debby is not taken to be committed to the truth of ϕ, we expect participants to respond that Debby is not certain that ϕ. However, we do not expect mean responses at the extreme ends of the 7-point Likert scale, because it is possible for participants to exhibit uncertainty about whether Debby is committed to ϕ. Participants can therefore give a response in the mid-range of the scale. To directly compare our model's predictions to the experimental results, we use the model to pre-

Figure 2: Model predictions about the extent to which listeners take the speaker to be certain of the prejacent, given the two melodies. The x-axis represents the λ-parameter, which encodes the degree to which the predictions reflect utility-maximizing reasoning. The y-axis is mapped to a 7-point scale to parallel the experimental task.

dict participants' probabilistic evaluation of whether Debby is certain that ϕ. The probability $P(\phi|M)$ is the probability that Debby is certain that ϕ, given M. This is expected to be 1 when the participant takes ϕ to project based on M—which occurs with probability $P(\text{PROJECT}(\phi)|M)$. When the participant does not take ϕ to project based on M, we expect some baseline uncertainty about whether Debby is committed to ϕ, which we encode as a prior probability $P(\phi)$. We thus define $P(\phi|M)$ as follows:

(16)
$$P(\phi|M) = P(\text{PROJECT}(\phi)|M) \times 1$$
$$+ P(\neg\text{PROJECT}(\phi)|M) \times P(\phi)$$

To account for the fact that our stimuli do not provide any prior evidence (i.e., evidence apart from the manipulated prosody) for whether the speaker is committed to ϕ, we take $P(\phi)$ to be uniform, i.e., equal to 0.5. Assuming a uniform prior over ϕ maintains a model with only one free parameter, the rationality parameter λ, and makes the model more informative by limiting the range of predictions that it can make.

Fig. 2 shows the model predictions[1] as λ increases to 10 (a relatively high value given ones used in the literature). We see that the modeled participant responses in the two prosodic conditions LH-Adverb and LH-Name diverge rather shallowly, predicting significantly higher certainty in the LH-Adverb condition, but not by a huge margin. The model predicts that projection in the LH-Adverb condition, though higher, will not be at ceiling. This is because utterances in the LH-Adverb condition are not only compatible with the prejacent-entailing QUDs but also with QUDs that do not entail the prejacent, and thus the probability of projection never exceeds fifty percent (4 on the Likert scale), even for high values of λ.[2]

---

[1]Model-generated probabilities are mapped onto a 7-point scale with the following formula: RATING $= 1 + 6 * P(\phi|M)$.

[2]The Python code used to implement the model can be downloaded at https://github.com/jonscottstevens/Prosody-Projection.


Figure 3: Responses by prosodic condition. Violin plots show frequency of participant means. Bar plots show overall means with 95% bootstrap confidence intervals.

Experiment results

The experimental results are shown in Fig. 3. Mean Likert scale response in the LH-Adverb condition was 5.7, compared to 4.8 in the LH-Name condition. This difference is in the expected direction: participants rated Debby as being less certain of the prejacent when the pitch accent was on the subject than when it was on the manner adverb. A mixed-effects ordinal regression model with random intercepts for participant and item and random slope for participant shows responses to be significantly lower in the LH-Name condition than in the LH-Adverb condition ($\beta = -1.13$, $SE = 0.24$, $z = -4.68$, $p < 0.0001$).

The model predictions in Fig. 2 are in line with the experimental results shown in Fig. 3 in three key ways:

1. The model correctly predicts a significant difference in mean responses between the two conditions, with the LH-Adverb items showing higher certainty ratings.
2. The model correctly predicts the magnitude of this difference to be rather small (within about one point on the Likert scale).
3. The model correctly predicts that even in the LH-Adverb condition, where projection is expected, the ratings are not at ceiling.

Thus, three qualitative experimental behaviors are accounted for by our RSA model, which builds on existing theoretical assumptions about the links between prosody, the QUD and projection, and incorporates those assumptions into a probabilistic pragmatic model.

The quantitative match with the model's predictions is not exact—certainty is a bit higher across the board than predicted—but as we see in Fig. 4, it is not far off, either. We would expect a more exact match if we experimentally obtained priors over the hearer's evaluation of the speaker's certainty for the various sentences used (instead of assuming a uniform prior of 0.5), a possible task for future research.

Discussion

This paper showed that the question-based analysis of projection developed in Simons et al. 2017 can be extended to manner adverb utterances and formalized in an RSA model.
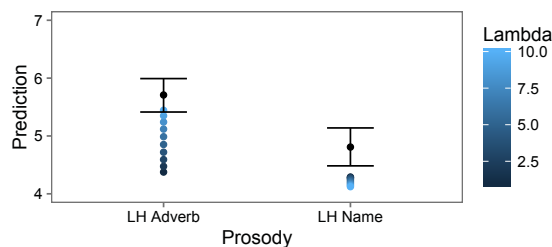
Figure 4: Model predictions by λ value (blue), along with the experimentally observed means (black).

The experimental findings empirically support the predictions of the model and, hence, the question-based analysis of projection. They also add to the growing empirical evidence that formal analyses of projection, including conventional triggering analyses (e.g., Heim, 1983; van der Sandt, 1992), need to be sensitive to information structure (e.g., Beaver, Roberts, Simons, & Tonhauser, 2017; Tonhauser, 2016). Finally, the RSA model demonstrates the feasibility of formal pragmatic analyses of projection.

Future research needs to investigate the predictions of the model for other projective contents of manner adverb utterances, other prosodic realizations of such utterances, and the projective contents of other utterances. We also observed that the influence of prosody on the projectivity of the prejacent was heterogeneous across items. This observation suggests enriching the model with information about listeners' prior expectations about the prejacent, e.g., about how likely somebody is to smile given that they didn't smile cheerfully.

Acknowledgments

References

Abrusán, M. (2013). A note on quasi-presuppositions and focus. *Journal of Semantics*, *30*, 257-265.

Beaver, D., & Clark, B. (2008). *Sense and sensitivity: How focus determines meaning*. Oxford: Wiley-Blackwell.

Beaver, D., Roberts, C., Simons, M., & Tonhauser, J. (2017). Questions Under Discussion: Where information structure meets projective content. *Annual Review of Linguistics*, *3*.

Beckman, M. E., & Ayers, G. (1997). Guidelines for ToBI labelling, version 3.0 [Computer software manual].

Bergen, L., & Goodman, N. D. (2015). The strategic use of noise in pragmatic reasoning. *Topics in Cognitive Science*, *7*(2), 336–350.

Chierchia, G., & McConnell-Ginet, S. (1990). *Meaning and grammar*. Cambridge, MA: MIT Press.

Eady, S., & Cooper, W. E. (1986). Speech intonation and focus location in matched statements and questions. *JASA*, *80*, 402-415.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998. doi: 10.1126/science.1218633

Halliday, M. A. K. (1967). Notes on transitivity and theme in English (part 2). *Journal of Linguistics*, *3*, 177-274.

Heim, I. (1983). On the projection problem for presuppositions. In M. Barlow, D. Flickinger, & M. Westcoat (Eds.), *WCCFL 2* (p. 114-125).

Lai, C. (2012). *Rises all the way up: The interpretation of prosody, discourse attitudes and dialogue structure* (Unpublished doctoral dissertation). University of Pennsylvania.

Most, R., & Saltz, E. (1979). Information structure in sentences: New information. *Language and Speech*, *22*, 89-95.

Paul, H. (1880). *Prinzipien der Sprachgeschichte*. Tübingen: Niemeyer.

Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, *5*, 1-69.

Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics*, *1*, 75-116.

Selkirk, E. O. (1996). Sentence prosody: Intonation, stress, and phrasing. In J. Goldsmith (Ed.), *The handbook of phonological theory* (p. 550-569). Blackwell.

Simons, M. (2001). On the conversational basis of some presuppositions. In *Proceedings of Semantics and Linguistics Theory (SALT) XI* (p. 431-448). Ithaca, NY: CLC Publications.

Simons, M., Beaver, D., Roberts, C., & Tonhauser, J. (2017). The Best Question: Explaining the projection behavior of factive verbs. *Discourse Processes*, *54*, 187-206.

Tonhauser, J. (2016). Prosodic cues to speaker commitment. In *Proceedings of Semantics and Linguistic Theory (SALT) XXVI*. Ithaca, NY: CLC Publications.

van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, *9*, 333-377.

von Stechow, A. (1990). Focusing and backgrounding operators. In W. Abraham (Ed.), *Discourse particles* (p. 37-84). Amsterdam: John Benjamins.

Wagner, M. (2012). Contrastive topics decomposed. *Semantics and Pragmatics*, *5*(8), 1–54.