Factivity in doubt: Clause-embedding predicates in naturally occurring discourse

Marie-Catherine de Marneffe[*], Mandy Simons[♦] and Judith Tonhauser[*]
[*]Department of Linguistics, The Ohio State University
[♦]Department of Philosophy, Carnegie Mellon University

Clause-embedding predicates, like *know* and *believe*, are a central case in the study of projection. Conventionalist approaches (e.g., Heim 1983, van der Sandt 1992) assume that this class of predicates is divided into 'factive' ones, which are conventional presupposition triggers, and 'non-factive' ones, which are not. This binary, categorical division is typically motivated by claims about projection from under entailment-canceling operators, e.g. negation, epistemic possibility modals and questions. We present results from a large-scale empirical study of projection in a new corpus of naturally occurring data that we built, the CommitmentBank. While these data support the claim that predicates give rise to projective readings to different degrees, they challenge the existence of a categorical distinction between 'factive' and 'non-factive' predicates. These findings instead support analyses of projection that do not assume conventional specification of presupposition but instead derive projectivity from a range of lexical and contextual factors (e.g., Simons et al. 2010, 2017, Abrusán 2011, Beaver et al. 2017).

CommitmentBank: The CommitmentBank is a corpus of 982 naturally occurring discourses extracted from the Wall Street Journal, the fiction section of the British National Corpus and Switchboard. Each discourse consists of a target sentence with a clause-embedding predicate embedded under an entailment-canceling operator as well as up to 2 prior context sentences, as in (1). In this corpus, 45 clause-embedding predicates are represented.

(1) *What fun to hear Artemis laugh. She's such a serious child. I didn't know she had a sense of humor.*

The projectivity of the content of the clausal complement of the target sentence was judged by at least 8 theoretically untrained native speakers of English recruited on Amazon's Mechanical Turk platform. These projectivity judgments were collected by asking Turkers to assess how certain the author or speaker of the target sentence is about the content of the clausal complement. Turkers gave ratings on a 7-point Likert scale labeled at 3 points: 3/speaker is certain that the complement is true, 0/speaker is not certain whether complement is true or false, -3/speaker is certain that complement is false. In (1), for instance, Turkers were asked to judge how certain the author is that Artemis has a sense of humor.

Using the CommitmentBank to evaluate the 'factive' / 'non-factive' distinction

Analyses that assume a categorical division between 'factive' and 'non-factive' predicates lead to three predictions. CommitmentBank data are inconsistent with all three predictions. First prediction**:** The first prediction is that discourses with 'factives' receive markedly higher projection ratings than discourses with 'non-factives'. This prediction is not borne out in the CommitmentBank data, as shown by the boxplot in Figure 1, which plots the raw ratings for all 982 discourses by predicate, with the predicates ordered from left to right by increasing mean ratings (given in blue). Although discourses with 'factives' tend to receive higher ratings than discourses with 'non-factives' (e.g., discourses with *occur, suppose* and *seem* are at the far left and discourses with *forget*, *learn* and *notice* are at the far right), we do not observe a sharp division between discourses with 'non-factives' on the left and discourses with 'factives' on the right. Although discourses with 'factives' all occur in the upper half, several 'non-factives' also occur in this set, including *tell* and *guess*; and some of these 'non-factives' receive higher mean certainty ratings than 'factives'. Overall, these findings suggest that projectivity is a gradient property of utterance content, as also suggested by the experiments in Tonhauser et al in press.

Second prediction: Conventionalist analyses derive projectivity from presuppositionality. Consequently, these analyses predict that projection itself should be binary: presupposed complements of 'factives' should either clearly project or, due to local accommodation, not project. Gradient speaker commitment to the complement is expected only for 'non-factives'. As shown in Figure 2, which plots the mean certainty ratings for the 578 discourses with the four most well-populated 'factive' predicates and with *think*, this prediction is not borne out. The projectivity of the content of the complement of 'factives' is as gradient as those of *think*.

Third prediction: Accounts assuming a 'factive'/'non-factive' distinction predict that 'factives' with complement *p* are more likely to occur in contexts which indicate speaker commitment to *p* than in contexts which do not, as the latter case requires accommodation. To test that prediction, we collected ratings of the extent to which the speaker/author is certain of the content of the clausal complement from reading only the 2 context sentences (on a subset of the corpus). These ratings, in Figure 3, show that the more *p* follows from the context, the more 'factives' with complement *p* are realized in such contexts. There is, however, no categorical division between 'factives' and 'non-factives' as 'non-factives' like *prove, signal, fear, mean*, *guess* and *suspect* also occur in the upper half of the graph.

Conclusions: Our findings are difficult to reconcile with a categorical division between 'factives' and 'non-factives', and suggest that projectivity arises from fine-grained features of lexical meaning and contextual features. Our empirically-driven approach aims to identify factors that contribute to projectivity and addresses the projective potential of *all* predicates.

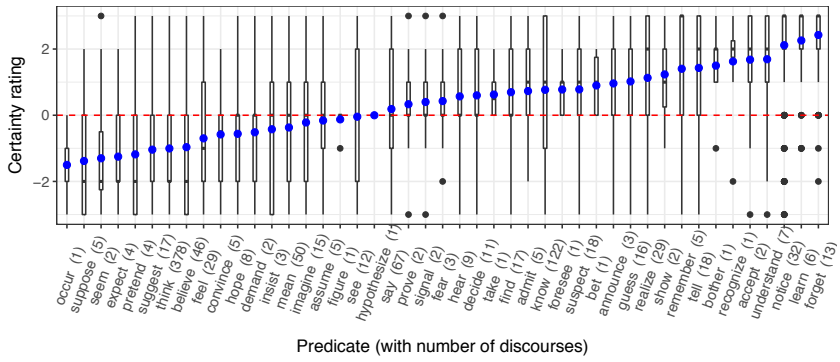

Figure 1: Boxplot of certainty ratings for discourses by predicate (mean ratings in blue).

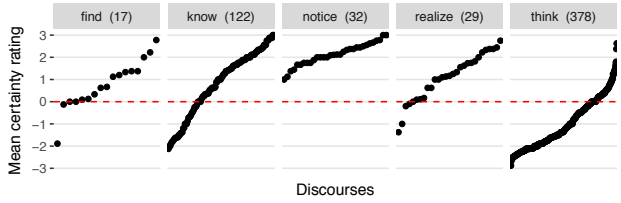

Figure 2: Mean certainty ratings for discourses with the 4 most well-populated 'factive' predicates and *think*.

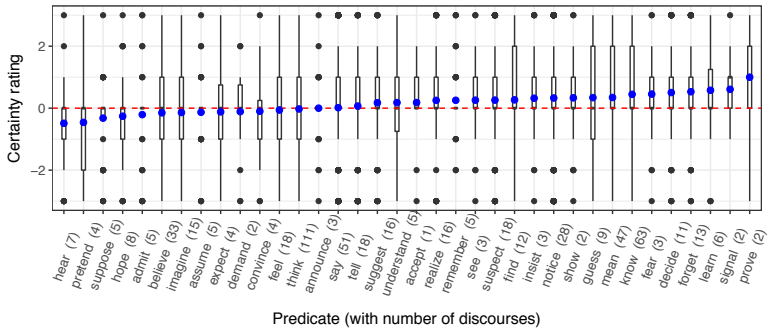

Figure 3: Boxplot of certainty ratings for the first 2 sentences of each discourse by predicate (mean ratings in blue).