Projection variation: Is the family of sentences really a family?[1]

Lisa HOFMANN — *University of Stuttgart*
Marie-Catherine DE MARNEFFE — *FNRS, UCLouvain*
Judith TONHAUSER — *University of Stuttgart*

Abstract. Under the 'family of sentences' diagnostic for projection, the projection of content is investigated by embedding the expression that contributes the content in the scope of negation, polar questions, epistemic possibility modals, and conditional antecedents. This paper reports on the results of a set of experiments designed to investigate whether there is variation in the projection of content from under these four types of entailment-canceling operators. The contents investigated are the contents of the complements of 20 English clause-embedding predicates. The results of the experiments suggest (i) that the by-operator variation is small when aggregating over the 20 contents, but (ii) that the effect of operator differs between the clause-embedding predicates. The results of these experiments also extend a result of Degen and Tonhauser 2022, that projection ratings in polar questions do not categorically distinguish factive from non-factive predicates, to cases with negation, the epistemic possibility modal *perhaps*, and conditional antecedents. The observed by-predicate and by-operator variation is not captured by existing theoretical accounts of projection (e.g., Heim 1983, van der Sandt 1992, Abrusán 2011, Schlenker 2021). Our results suggest that an empirically adequate projection analysis must consider interactions between predicates and operators.

Keywords: Projection variation, entailment-canceling operators, (non)factive predicates.

1. Introduction

The 'family of sentences' diagnostic is the standard way of diagnosing whether a content is projective (e.g., Chierchia and McConnell-Ginet 1990). For instance, in (1), the content of the clausal complement of *discover* (that Julian dances salsa) is diagnosed as projective content, if it is typically implied not just by an utterance of (1), but also by utterances of the variants in (2), where (1) is embedded under an entailment-canceling operator, such as negation (2a), a polar question (2b), an epistemic possibility modal (2c), or in a conditional antecedent (2d).

(1) Cole discovered that Julian dances salsa.

(2) a. **Negation:** Cole didn't discover that Julian dances salsa.
b. **Polar Question:** Did Cole discover that Julian dances salsa?
c. **Modal:** Perhaps Cole discovered that Julian dances salsa.
d. **Conditional:** If Cole discovered that Julian dances salsa, Logan will be joyful.

Some research, however, suggests that entailment-canceling operators may affect projection differentially. For instance, Karttunen (1971) proposed distinguishing English factive predicates (e.g., *regret*) from semi-factives (e.g., *discover*). Based on (3), he argued that the content

[1]We thank Taylor Mahler for assistance in collecting the data presented here as well as valuable comments. We gratefully acknowledge the National Science Foundation grant BCS #1452674 (to Marie-Catherine de Marneffe, Craige Roberts, and Judith Tonhauser), which provided financial support for this project. Marie-Catherine de Marneffe is a Research Associate of the Fonds de la Recherche Scientifique – FNRS.

 In: Baumann, Geraldine, Daniel Gutzmann, Jonas Koopman, Kristina Liefke, Agata Renans, and Tatjana Scheffler (eds.) 2024. Proceedings of Sinn und Bedeutung 28. Bochum: Ruhr-University Bochum, 422-440.

of the complement (CC) of true factives consistently projects across the operators in (2), but that of semi-factives does not always project from under polar questions, modals, or conditionals.

(3) Karttunen 1971: (22, 24–26)
   a. John didn't {regret/discover} that he had not told the truth.
   b. Did you {regret/discover} that you had not told the truth?
   c. If I {regret/discover} later that I have not told the truth, I will confess it to everyone.
   d. It is possible that I {regret/discover} later that I have not told the truth.

There have been two experimental investigations of by-operator projection variation. First, Smith and Hall (2014) investigated projection from under negation and conditional antecedents for various types of English projective contents. They found that the expressive content of epithets (e.g., *idiot*) and the CC of *know* was more projective under negation than conditionals. In contrast, the content of appositive relative clauses and the preparatory content of *win* showed the opposite pattern, and the existential presupposition of clefts showed no difference. Second, Sieker and Solstad (2022) compared the projection of the CCs of German factives (*wissen* 'know', *bereuen* 'regret', *enthüllen* 'reveal') and semi-factives (*bemerken* 'notice', *entdecken* 'discover', *herausfinden* 'find out') from under the four operators in (2). Their results suggest that the CCs project more from under negation than from under the other three operators. Contrary to what Karttunen (1971) suggested, a comparison of the factive and semi-factive predicates did not reveal that the CCs of factive predicates project more from under polar questions, modals, or antecedents of conditionals than the CCs of semi-factive predicates.[2]

This paper reports on the results of a set of experiments that were designed to compare projection from under the four entailment-canceling operators in (2) in English. Our experiments extend the empirical scope of prior research on by-operator projection variation by investigating projection for a larger set of contents, namely the contents of the complements of the 20 English clause-embedding predicates in (4), from Degen and Tonhauser (2022).

(4) a. (Semi-)factive predicates: *be annoyed, know, reveal, discover, see*
   b. Non-factive predicates: *acknowledge, admit, announce, confess, confirm, establish, hear, inform, prove, be right, demonstrate, pretend, say, suggest, think*

The five predicates in (4a) have been characterized as factive or semi-factive. Our set of predicates also includes the 15 non-factives in (4b). Including non-factive predicates in investigations of projection is motivated by the empirical investigations in de Marneffe et al. 2019 and Degen and Tonhauser 2022, which suggest that the CCs of non-factives may also project and that projection ratings do not categorically distinguish factive and non-factive predicates.

The results of our investigation suggest that the projection of the CCs of these 20 predicates is affected differently by the four entailment-canceling operators in (2), but not in a way that is consistent with Karttunen's 1971 factive/semi-factive distinction. The results also replicate a result of Degen and Tonhauser 2022, namely that projection from under polar questions does not categorically distinguish factive and non-factive predicates. We also extend this result to projection from under the other three entailment-canceling operators, thereby solidifying their

[2]The factive/semi-factive distinction is also called into question by naturally occurring examples where the CCs of factive predicates do not project from under the four operators (see Beaver 2010, de Marneffe, Manning, and Potts 2012, de Marneffe, Simons, and Tonhauser 2019). For experimental research on the distinction, see Djärv, Zehr, and Schwarz 2018.

claim that "research on projective content has a much broader empirical scope than previously assumed" (p.585), as this scope includes the CCs of both factive and non-factive predicates.

2. Experiments

To assess the effect of entailment-canceling operator and clause-embedding predicate on projection, we collected projection judgments for the CCs of the 20 clause-embedding predicates in four sets of experiments. The predicates were embedded under negation in Exps. 1, under polar questions in Exps. 2, under the epistemic possibility modal *perhaps* in Exps. 3, and in conditional antecedents in Exps. 4. Each set of experiments contained three experiments that differed in the at-issueness measure that was used in a separate block. In this paper, we limit our attention to the projection ratings collected in these twelve experiments.[3]

In all twelve experiments, projection was measured with the 'certain that' diagnostic, which has been used to measure projection with both polar interrogative and declarative sentences (see, e.g., Tonhauser 2016, Stevens, de Marneffe, Speer, and Tonhauser 2017, Tonhauser, Beaver, and Degen 2018, Mahler 2019, Djärv and Bacovcin 2020, Mahler 2020, de Marneffe et al. 2019, Sieker and Solstad 2022).[4] Under this diagnostic, participants are presented with utterances like those in (5), and asked to rate whether the (named) speaker is certain of the CC.

(5) a. Christopher: "Cole didn't discover that Julian dances salsa."
b. Christopher: "Did Cole discover that Julian dances salsa?"
c. Christopher: "Perhaps Cole discovered that Julian dances salsa."
d. Christopher: "If Cole discovered that Julian dances salsa, Logan will be joyful."
Projection question: Is Christopher certain that Julian dances salsa?

We assume, following Tonhauser et al. 2018 and Degen and Tonhauser 2022, that judgments of speaker certainty about the embedded content reflect speaker commitment to that content, that is, projection. If a participant interprets utterances like (5a–d) in a way that the speaker (here, Christopher) is certain of the CC, the CC is assumed to project. If a participant does not take the speaker to be certain of the CC, the CC is taken to not project.

Participants. We recruited 250-300 participants for each of the 12 experiments. Participants for one experiment were recruited on Amazon's Mechanical Turk platform. These participants were required to have U.S. IP addresses and at least 99% of previously approved HITs. Participants for the remaining experiments were recruited on Prolific. These participants were required to reside in the US, to be born in the US, to have English as their first language, and to have an approval rating of at least 99%. See Supplement D (in the repository linked to in footnote 3) for further information about the participants.

Materials. The target sentences consisted of the 400 combinations of the 20 clause-embedding predicates in (4) with 20 embedded clauses (provided in Supplement A). As mentioned

[3]The experiments, data and analysis scripts, as well as the supplements referred to in this paper can be found in the following GitHub repository: `https://github.com/judith-tonhauser/CommitmentBankPlus`.
[4]For other diagnostics of projection see, e.g., Smith and Hall 2011, Xue and Onea 2011, and Tonhauser, Beaver, Roberts, and Simons 2013, and discussion in Tonhauser et al. 2018.

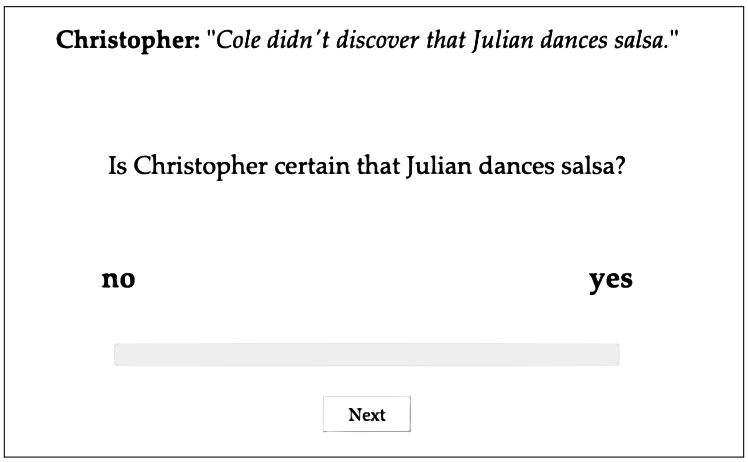

Figure 1: A sample trial from Exps. 1 (negation). In the other experiments, participants were presented with utterances with a different entailment-canceling operator.

above, the predicates were embedded under negation in Exps. 1, under polar questions in Exps. 2, under the epistemic possibility modal *perhaps* in Exps. 3, and in conditional antecedents in Exps. 4, for a total of 400 target stimuli in each of the four sets of experiments. To assess whether participants were attending to the task, each experiment included six control stimuli. For details on the six control stimuli, see Supplement C.

Each participant saw a random set of 26 stimuli: Each set contained one target stimulus for each of the 20 clause-embedding predicates (each with a unique complement clause) and the same six control stimuli.[5] Trial order was randomized.

Procedure. Participants were asked to imagine that they are at a party and that, when walking into the kitchen, they overhear somebody say something to somebody else. On each trial, they read an utterance and gave a response to the 'certain that' question on a slider marked 'no' (coded as 0) at one end and 'yes' (coded as 1) at the other. A sample trial is shown in Figure 1. Following Tonhauser et al. 2018, higher ratings of speaker certainty could reflect one of two things. First, higher certainty ratings could reflect greater speaker commitment towards the CC, and therefore greater projection. This assumes that speaker commitment is interpreted in a gradient way. Second, higher certainty ratings could reflect a higher probability that an interpreter takes the speaker to be committed to the CC. On this interpretation, speaker commitment may be a binary, categorical property and projection variation is a result of uncertainty about speaker commitment. In this paper, we remain agnostic about the underlying interpretation of projection as a gradient property (for discussion, see Grove and White 2023).

At the end of the experiment, participants filled out a short optional demographic survey. To encourage truthful responses, they were told that they would be paid no matter what answers they gave in the survey.

Data exclusion. Data were excluded based on self-declared non-native speaker status and other criteria given in Supplement D. The data from 2,682 participants entered into the analysis.

[5]Each participant saw their set of 26 stimuli twice, once in the projection block and once in the at-issueness block. Block order was randomized. As mentioned above, we focus here on the projection ratings.

3. Results and discussion

We first address by-operator variation (Section 3.1) and then the question of whether there is by-predicate variation in the observed by-operator variation (Section 3.2). Finally, in Section 3.3, we relate our results to those of Degen and Tonhauser 2022, that projection from under polar questions does not categorically distinguish factive and non-factive predicates.

3.1. By-operator variation

Figure 2 shows the mean certainty ratings by entailment-canceling operator, aggregating over the clause-embedding predicates. As shown, there is projection variation by operator: The CCs of the clause-embedding predicates were relatively more projective when embedded in the antecedent of a conditional than in a polar interrogative, where they were relatively more projective than when they were embedded under negation or the epistemic modal *perhaps*.

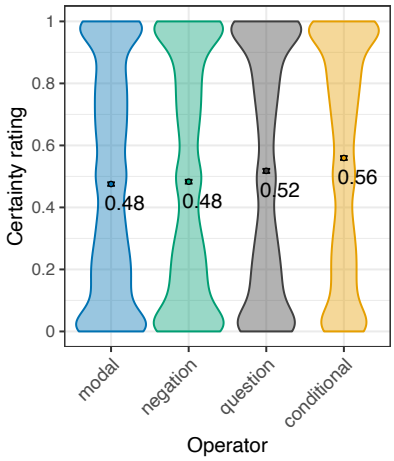

Figure 2: Mean certainty ratings by operator. Error bars indicate 95% bootstrapped confidence intervals. Violin plots indicate the kernel probability density of participants' individual ratings.

These observations are supported by a post-hoc pairwise comparison of the estimated means for each entailment-canceling operator using the `emmeans` package (Lenth 2023) in R (R Core Team 2016). The input to the pairwise comparison was a Bayesian mixed-effects beta regression model that was fit using the `brms` package (Bürkner 2017) with weakly informative priors. The model predicted certainty ratings[6] from a fixed effect of entailment-canceling operator (with treatment coding and 'modal' as the reference level) and included a random by-predicate intercept.[7] The output of the pairwise comparison were 95% highest density intervals (HDIs) of estimated marginal mean differences between each of the operators. We assume that two operators differ in certainty ratings if the HDI of their pairwise comparison does not include 0.

Table 1 provides the output of the pairwise comparison on a logit scale. As shown, the analysis suggests differences between each pair of operators. That is, certainty ratings are higher for

[6]To model the certainty ratings using a beta regression, the ratings were first transformed from the interval [0,1] to the interval (0,1) using the method proposed in Smithson and Verkuilen 2006.
[7]See Supplement E for details on the model.

CCs embedded in conditional antecedents than those embedded under polar questions, certainty ratings for CCs embedded under polar questions are higher than those for CCs embedded under negation, and finally certainty ratings for CCs embedded under negation are higher than those for CCs embedded under the epistemic modal *perhaps*.

contrast	estimate	lower 95% CrI	upper 95% CrI
conditional - negation	0.21	0.18	0.24
conditional - question	0.13	0.10	0.16
modal - conditional	-0.26	-0.29	-0.23
modal - negation	-0.05	-0.08	-0.02
modal - question	-0.14	-0.17	-0.11
negation - question	-0.09	-0.12	-0.06

Table 1: Output of the pairwise comparison of entailment-canceling operators. The 'contrast' column identifies entailment-canceling operators pairs, 'estimate' the estimated marginal mean difference, and 'lower/upper 95% CrI' provide the lower/upper bounds of the HDIs.

These results suggest that certainty ratings for the CCs of the English clause-embedding predicates we investigated vary by entailment-canceling operator. In contrast to Sieker and Solstad 2022 for German, the results of our experiments do not suggest that projection is strongest from under negation. Recall, however, that they only investigated projection of the CCs of (semi-) factive predicates. Since there is by-predicate variation in the effect of entailment-canceling operator on projection (as we show in the next section), this difference between the results of their experiment and ours might be due to the types of predicates investigated. Finally, the differences in mean certainty ratings between the four entailment-canceling operators are very small. This suggests that, when abstracting away from individual predicates and contents, projection from under the four entailment-canceling operators is very similar. In other words, when abstracting away from individual contents, the family of sentences really are a family.

3.2. By-predicate variation in the effect of entailment-canceling operator

Figure 3 shows mean certainty ratings by entailment-canceling operator for the 20 predicates, with predicates ordered by their overall mean certainty rating. As shown, there is by-operator projection variation for each predicate. Further, the effect of operator differs between predicates. For instance, the five (semi-)factive predicates (highlighted in pink) exhibit four different patterns. First, the CC of *be annoyed* projects most from under questions, less from under negation, followed by conditionals, and least from under the modal *perhaps* (Q > N > C > M).[8] Second, the CC of *know* projects most from under questions, less from conditionals and negation, and least from under *perhaps* (Q > {N, C} > M). The CCs of *discover* and *see* exhibit a third pattern: They project most from under questions and conditionals, less from under negation, and least from under *perhaps* ({Q, C} > N > M). Finally, the CC of *reveal* projects most from conditionals, less from questions, and least from under negation and *perhaps* (C >

[8]Under this notation, > indicates a non-zero difference between two adjacent levels of operator, when operators are ordered by mean projection ratings.

Q > {M, N}). Thus, the CCs of none of our (semi-)factive predicts project uniformly from under all four entailment-canceling operators (contrary to what Karttunen 1971 suggested for factive predicates) and the purported semi-factive predicates *discover* and *reveal* do not project more from under negation than the other three entailment-canceling operators.

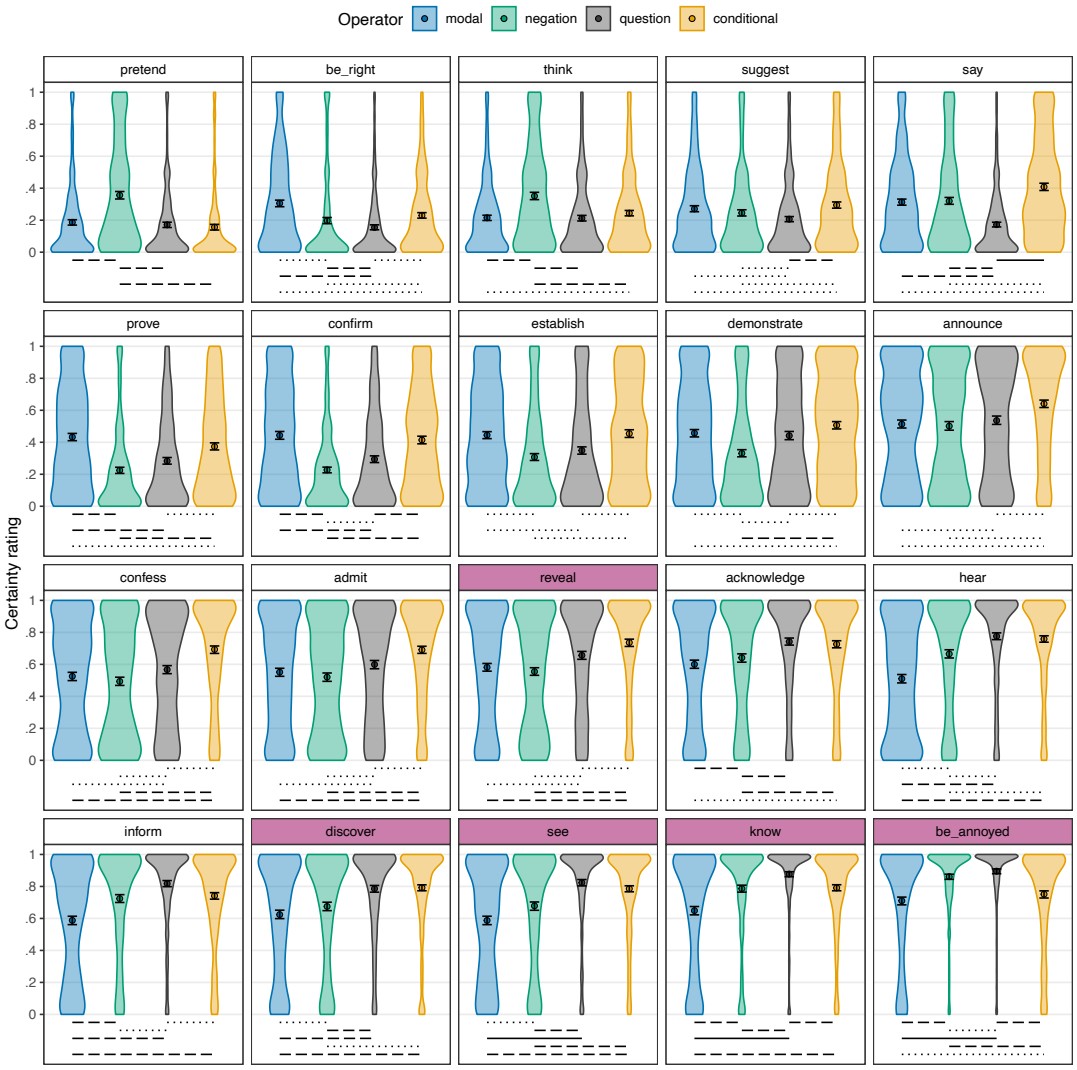

Figure 3: Mean certainty ratings by predicate ((semi-)factive, non-factive) and operator (modal, negation, polar questions, conditional antecedents) with 95% bootstrapped confidence intervals. Violin plots indicate kernel probability density of individual participants' ratings. Predicate facets are ordered by the predicate mean certainty rating (aggregating across operators). Below each facet, a line spanning two operators indicates a non-zero difference according to the pairwise comparison of operators. The line type indicates whether the difference d is ≥ 1 (solid line: —), $0.5 \leq d < 1$ (dashed line: – –), $0 \leq d < 0.5$ (dotted line: . . .).

We also observe by-operator projection variation for non-factive predicates. Some of this variation aligns with that observed for factive predicates: For instance, the CC of *inform* exhibits the same pattern as the CC of *know*, and the CC of *hear* the same pattern as those of *discover* and *see*. Other non-factive predicates exhibit other patterns: The CCs of *admit*, *confess* and *announce* project most from the conditional antecedents than the other three operators.

These observations are supported by post-hoc pairwise comparisons of the estimated means for each entailment-canceling operator within each predicate using the `emmeans` package (Lenth 2023). The input to the pairwise comparisons for each predicate were 20 Bayesian mixed-effects beta regression models that were fit using the `brms` package (Bürkner 2017) with weakly informative priors. The models for each predicate predicted certainty ratings[9] from a fixed effect of entailment-canceling operator (with treatment coding and 'modal' as the reference level) and included a random by-item intercept and a random slope for operator by item.[10] The output of the pairwise comparisons were 95% highest density intervals (HDIs) of estimated marginal mean differences between each of the operators for each predicate. We assume that two operators differ in certainty ratings for a given predicate if the HDI of their pairwise comparison does not include 0. These non-zero differences between two operators are indicated by lines spanning the two operators below each predicate facet in Figure 3.

Our findings align with those of Smith and Hall 2014, who also observed by-expression variation in the effect of operator. However, while they found that the CC of *know* projects more from under negation than the antecedent of a conditional, we did not find a difference here. We hypothesize that this difference is due to the difference in projection diagnostic used. Our results differ, however, from those of Sieker and Solstad 2022. While their work did not find differences in by-operator projection variation between factive and semi-factive predicates, our results suggest four different patterns of by-operator variation for the five (semi-)factive predicates. As Sieker and Solstad 2022 also used the 'certain that' diagnostic for projection, this difference in results is not likely due to the projection diagnostic. Other factors that varied between our experiments are the language under investigation (German vs. English), the clause-embedding predicates investigated, and the CCs that the predicates were paired with. Future research will need to establish which of these factors are implicated in the observed differences.

3.3. Factive vs. non-factive predicates

Lexical approaches to projection assume that factive predicates are ones that presuppose the CC, while the CC of non-factive predicates is not presupposed (e.g., Kiparsky and Kiparsky 1970, Karttunen 1971, Schlenker 2009, Abrusán 2011).[11] Because presuppositions are assumed to typically project from under entailment-canceling operators, this definition predicts that factive predicates are distinguished from non-factives by the projection of their CCs: The CCs of factive predicates are expected to be categorically more projective than those of non-factives. This expectation was investigated in Degen and Tonhauser 2022 based on the 20 clause-embedding predicates in (4) embedded in polar questions. Contrary to expectation, Degen and Tonhauser's 2022 Exps. 1 found that the CCs of the five (semi-)factive predicates varied in projection, that the CCs of the 15 non-factive predicates were projective compared to the non-projective main clause contents, and that the CCs of some non-factives were as projective, or even more projective, than those of some factive predicates. In short, projection of

[9] To model the certainty ratings using a beta regression, the ratings were first transformed from the interval [0,1] to the interval (0,1) using the method proposed in Smithson and Verkuilen 2006.

[10] See Supplement F for further details on the models.

[11] Some of these works additionally assume that the CC of factive predicates is entailed. For details see Degen and Tonhauser 2022.

the CC from under polar questions did not categorically distinguish factive from non-factive predicates. Further support for this result came from analyses of projection ratings in three additional datasets, namely the CommitmentBank (de Marneffe et al. 2019), the VerbVeridicality dataset (Ross and Pavlick 2019), and the MegaVeridicality dataset (White and Rawlins 2018).

The results of the experiments reported on in this paper replicate Degen and Tonhauser's 2022 result. As shown in Figure 4a, there is variation between the five (semi-)factive predicates in the polar question condition, and projection from under polar questions does not categorically distinguish factive from non-factive predicates. Furthermore, our results suggest that this result can be extended to the three other entailment-canceling operators. As shown in Figures 4b-4d, there is variation in the projection of the CCs of the five (semi-)factive predicates from under *perhaps*, negation, and conditional antecedents, and projection ratings in these conditions do not show a categorical difference between factive and non-factive predicates either. These results lend further support to the conclusion of Degen and Tonhauser 2022 that there is, to date, no empirical evidence for a coherent class of factive predicates.

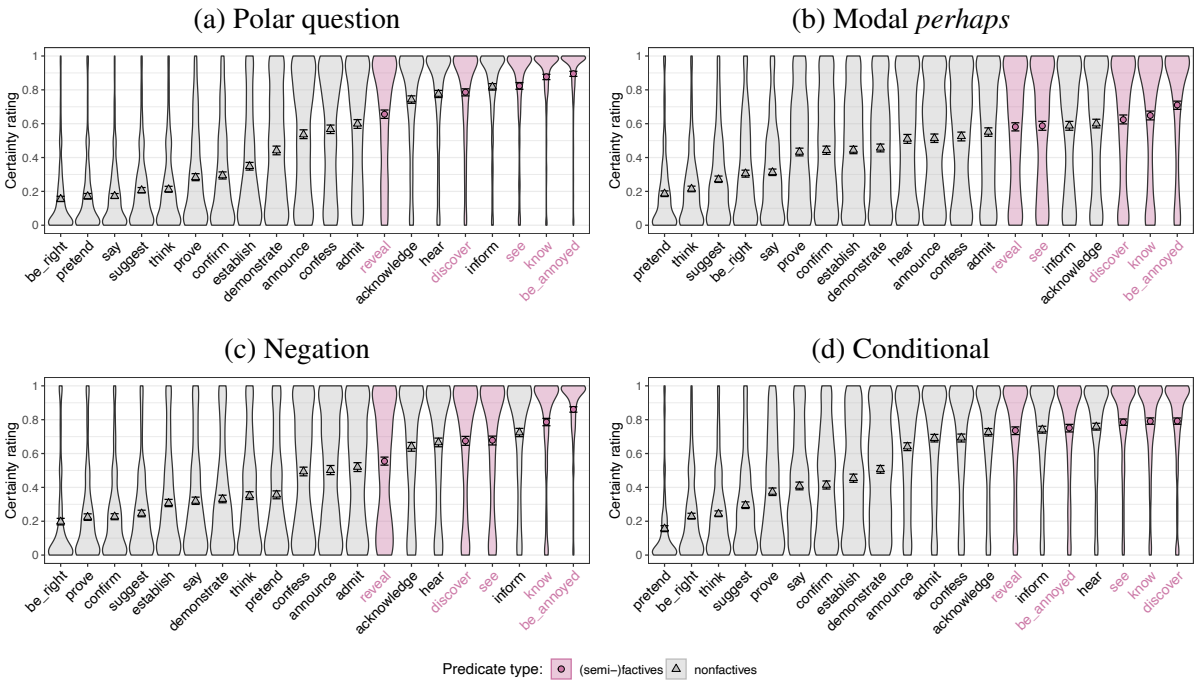

Figure 4: Mean certainty ratings by predicate, with (semi-)factive predicates in pink, for (a) polar questions, (b) the modal *perhaps*, (c) negation, (d) conditional antecedents. Error bars indicate 95% bootstrapped confidence intervals. Violin plots indicate kernel probability densities of the individual participants' ratings.

3.4. Summary

The results of our experiments suggest that there is little by-operator variation when aggregating over clause-embedding predicates, but that the CCs of different clause embedding predicates exhibit by-operator projection variation. Crucially, the effect of operator on projection differs by predicate, but not in ways that align with prior claims about differences between factive and semi-factive predicates (e.g., Karttunen 1971). Finally, the results of our experiments provide

further support for the results of Degen and Tonhauser 2022, who did not find empirical support for a class of factive predicates based on the projection of the CC from under polar questions. Our experiments suggest that projection of the CC from under negation, the antecedent of conditionals and epistemic modals also do not provide empirical support for a natural class of factive predicates. Before discussing the methodological and theoretical implications of these results in Section 5, we provide converging evidence from a different dataset in Section 4.

4. Converging evidence for the by-predicate variation in the effect of operator

We provide converging evidence for the by-predicate variation in the effect of operator based on the MegaVeridicality dataset (White and Rawlins 2018). This dataset contains projection ratings for the CCs of 517 English clause-embedding predicates. As shown in (6) for *know*, the predicates were combined with (what the authors refer to as) "low lexical content". The predicates were embedded under negation as in (6a), in the antecedent of a conditional as in (6b), and under negation in the antecedent of a conditional as in (6c). To assess projection, participants were asked to respond to the question *did that thing happen?*. The response options were 'yes' (indicating projection), 'maybe or maybe not', and 'no' (no projection).

(6) a. Somebody didn't know that a particular thing happened. Did that thing happen?
b. If somebody knows that a particular thing happened, did that thing happen?
c. If somebody didn't know that a particular thing happened, did that thing happen?

To investigate by-predicate projection variation, we recoded the responses as 1 ('yes'), -1 (no), and 0 'maybe or maybe not'). We calculated the mean projection ratings for 25 predicates under the three types of operator combinations shown in (6). We used the 14 non-factive and 5 factive predicates from our experiments that are in MegaVeridicality (*be right* is not included). As Djärv et al. 2018 suggested that the factive/semi-factive distinction can be understood as a difference between cognitive and emotive predicates, we included the emotive (*be disappointed*, *be surprised*) and cognitive (*realize*, *find out*) predicates from their experiments that are in MegaVeridicality. We also added two other predicates suggested by Karttunen as factive (*regret*) and semi-factive (*notice*).

Figure 5 shows mean projection ratings by embedding operator(s) and predicate. As shown, the effect of operator varies by predicate: For many, though not all, predicates, ratings are (at least numerically) lower when embedded under negation, or under negation and in a conditional antecedent, than when embedded in a conditional antecedent. In addition, there is by-predicate variation in the effect of entailment-canceling operator. For instance, the CCs of *acknowledge* and *hear* are more projective (at least numerically) when embedded under negation than when embedded under negation and in a conditional antecedent. The CCs of *demonstrate* and *prove*, on the other hand, exhibit the opposite pattern.[12] These observations suggest that there is by-predicate variation in the effect of entailment-canceling operator even when projection ratings are collected with a different measure and different materials than in our experiments.

[12] As discussed in White and Rawlins 2018: §3.2, the question in (6a) measures projection or, rather, global accommodation, given that the relevant content is not entailed by the common ground. However, since the predicates are embedded in conditional antecedents in stimuli like (6b) and (6c), and the questions that participants respond to in the conditional consequents, (6b) and (6c) might measure local accommodation of the content in the scope of the conditional antecedent. This does not change the point we make in this section, namely that the MegaVeridicality data suggest that there is by-predicate variation in the effect of entailment-canceling operator.

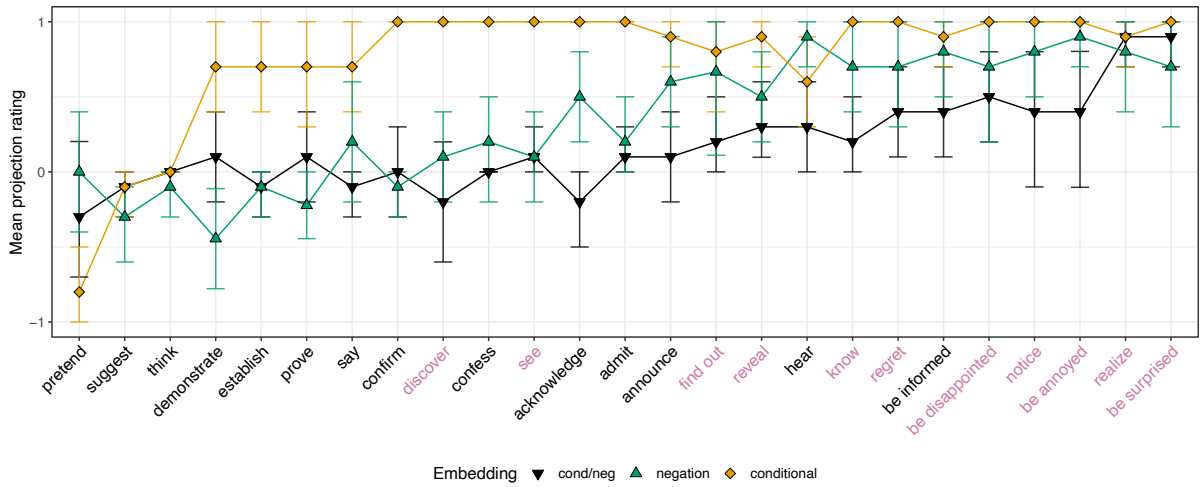

Figure 5: Mean projection ratings by entailment-canceling operator(s) and predicate in the MegaVeridicality dataset. Error bars indicate 95% bootstrapped confidence intervals.

5. General discussion

We now point out methodological implications of our results (Section 5.1), discuss whether contemporary projection analyses can capture the observed variation (Section 5.2), and speculate about lexical differences between the clause-embedding predicates that might predict the by-operator variation observed (Section 5.3).

5.1. Methodological implications

The results of the experiments reported on in Section 2 suggest that there is little by-operator projection variation when aggregating observations for the CCs of the clause-embedding predicates, but that there is by-operator variation that cannot be neglected when we do not aggregate. These results have two methodological implications. First, when initially investigating the projection of a content (or teaching projection to students), the family-of-sentences can indeed be treated as a family, as standard textbooks assume. However, it is advisable to apply the diagnostic with all four operators and attend to possible by-operator projection variation.

A second methodological implication of the results of our experiments is that research on projection should take into account possible by-operator variation when choosing which entailment-canceling operator to use. Further, results about projection variation should be relativized to the entailment-canceling operator under which the expressions and contents were embedded, as they may very well depend on this choice. For instance, two contents that exhibit projection variation when embedded under polar questions might not exhibit such differences with a different entailment-canceling operator, and two contents that do not exhibit projection variation from under negation might exhibit variation when embedded under some other operator.

5.2. Theoretical implications

This section discusses whether contemporary analyses can capture the observed by-predicate projection variation and the interactions between predicate and entailment-canceling operator.

5.2.1. Heim 1983 and van der Sandt 1992

On the analyses in Heim 1983 and van der Sandt 1992, the CCs of factive predicates are lexically specified as presupposed, in contrast to the CCs of non-factive predicates. Presupposed CCs project to the global context, except when that would produce an inconsistency, in which case they are accommodated to the local context of the operator.

These analyses do not predict the observed by-predicate variation, for several reasons. One reason is variation among the factive predicates: For instance, when embedded under polar questions, negation, or *perhaps,* the CC of *be annoyed* is more projective than that of *discover*, which in turn is more projective than that of *reveal*. (In conditional antecedents, the order of *be annoyed* and *discover* is reversed, see Figure 4.) This variation is not predicted because the CCs of factive predicates are invariably specified as presupposed, with no mechanism to predict projection variation between factive predicates. (Recourse to local accommodation is not possible because the stimuli in our experiments were presented in minimal contexts that were not inconsistent with the CCs.) A second reason is that the CCs of some non-factive predicates (e.g., *inform, acknowledge*) are just as projective as those of some factive predicates, or even more projective. Analyses like those in Heim 1983 and van der Sandt 1992 fail to make predictions about the projection of the CCs of non-factive predicates. Thus, as discussed in detail in Degen and Tonhauser 2022, the factive/non-factive distinction is not sufficiently fine-grained to adequately predict the observed by-predicate projection variation.

The analyses in Heim 1983 and van der Sandt 1992 also do not predict differential effects of entailment-canceling operators on projection. This is because negation and conditional antecedents are given a semantics that derives their behavior as presuppositional holes, in the sense of Karttunen 1971. For instance in Heim 1983, presupposition triggers under negation or in conditional antecedents are evaluated relative to a local context that is equal to the global context. Therefore, presuppositions are not expected to project differently from under negation or conditional antecedents. While the analysis does not explicitly address epistemic possibility modals or polar questions, we might expect that they would be treated as presuppositional holes along the same lines, therefore also not predicting any by-operator projection variation, or by-predicate variation in the effect of entailment-canceling operators.

5.2.2. Abrusán 2011 and Simons, Beaver, Roberts, and Tonhauser 2017

Abrusán (2011) and Simons et al. (2017) do not assume that presuppositions are lexically specified. In Abrusán 2011, a lexical entailment of a (possibly embedded) sentence is a presupposition, if it is about a time that is not the event time of the matrix predication and it is not at-issue with respect to the Question Under Discussion (QUD, see Roberts 2012) addressed by

the utterance. For instance, the CC of B's utterance in (7), that Phil's ballet class is canceled, is predicted to be a presupposition (and therefore to project) because it is a lexical entailment of the modal prejacent and not at-issue with respect to A's interrogative utterance.

(7) Adapted from Simons et al. 2017: 188
Context: It's early on Saturday morning. A and B are talking about their son.
A: Why is Phil up already?
B: Perhaps he forgot that his ballet class is canceled today.

In Simons et al. 2017, the CC of a clause-embedding predicate projects if it is entailed by the Current Question of the utterance (where, roughly, the Current Question is the question that is congruent with the utterance).[13] In (7), the Current Question of B's utterance might be the set of propositions {Phil forgot that his ballet class is canceled today, Phil is aware that his ballet class is canceled today}. If so, the Current Question entails that Phil's ballet class is canceled today, and the CC may therefore project under Simons et al.'s 2017 analysis.

Under the analyses in Abrusán 2011 and Simons et al. 2017, CCs that are entailed based on the literal content of an utterance may or may not project, depending on whether they contribute to the main point of the utterance. This opens up the possibility for projection variation among veridical predicates (including those often characterized as factive), representing an improvement compared to the lexical specification of factivity assumed in Heim 1983 and van der Sandt 1992. However, these accounts still do not predict the observed by-predicate variation. First, to make explicit predictions about the observed by-predicate variation, these types of analyses would need to make explicit assumptions about how different predicates contribute to how the question under discussion is chosen. Second, the analyses do not make systematic predictions for non-entailed CCs, and hence they do not predict that the CCs of some non-factive predicates project at least as much or even more than those of some factive predicates.

Both of these analyses set aside the effect of entailment-canceling operators in determining the main point of an utterance. Therefore they do not make predictions about by-operator differences or by-predicate variation in the effect of entailment-canceling operators.

5.2.3. Schlenker 2021

Under the view put forward in Schlenker 2021, the CC of a sentence S like (8a), is presupposed in a context c, if the CC is presupposed by the sentence S' under the entailment-canceling operator, that is (8b), in its local context c'. For the CC to be presupposed in (8b), two conditions must be met: (i) S' contextually entails the CC relative to c'; and (ii) If we consider "a generic agent" who believes the propositions in c' and who has now learned about the truth of S$'$, then the probability that this generic agent already believed the CC is above a contextual threshold α; more colloquially, condition (ii) requires that the generic agent "typically antecedently believes" the CC (p.6) upon interpreting S$'$ in c'. Based on Heim 1983 and Schlenker 2009, the the local context under negation c' is assumed to be identical to the global context c. There-

[13] The Current Question is defined in Simons et al. 2017: 194 as follows: "The Current Question for an utterance is a privileged subset of the focal alternative set of the uttered sentence (given a structural analysis of that sentence, including focus marking)" which meets the conditions that "(i) the proposition expressed is a member of the Current Question and (ii) the Current Question has at least one additional member."

fore, epistemic preconditions relative to c' will also apply to the global context. Applied to (8), Schlenker's 2021 analysis predicts that the CC is presupposed if (i) (8b) contextually entails that Julian dances salsa, and (ii) a generic agent "typically antecedently believes" that Julian dances salsa upon interpreting (8b) in the contexts provided to our participants.

(8) a. Cole doesn't know that Julian dances salsa.
b. Cole knows that Julian dances salsa.

Condition (i) is met under the assumption that the CC of (8b) is an entailment. Schlenker 2021 also assumes the condition (ii) is met: "in many cases, one's knowledge of facts will precede one's knowledge of [Cole's] beliefs about them . . . believing that [Julian dances salsa] is often an epistemic precondition for believing that" Cole knows that Julian dances salsa (p.6). One might, however, challenge this assumption on the basis of the corpus investigation presented in Spenader 2002, which showed that the CCs of the majority of the utterances of sentences with factive verbs (namely 81 out of 109) had to be accommodated (i.e., were not contextually entailed). In other words, utterances of sentences with the factive predicates investigated by Spenader (2002), which included *know*, were "generally used to communicate information the speaker thought was hearer-new" (p.99). This result might therefore suggest that one cannot assume that a generic agent typically antecedently believes the CC of *know*.

Schlenker's 2021 analysis does not incorporate differential effects of entailment-canceling operators on projection. As discussed for Heim 1983 above, the local contexts under negation and in conditional antecedents are both assumed to be the global context (and similarly in Schlenker 2009). Therefore, the analysis does not make predictions about by-operator projection variation or about by-predicate variation in the effect of entailment-canceling operator.

Does the account capture the observed by-predicate variation? The account takes the subjective conditional probabilities associated with expressions and the contextual threshold α as given, and therefore does not offer explicit or systematic predictions about projection ratings for different expressions. We, therefore, assume that it is an open, empirical question which predicates are such that the probability of a generic agent antecedently believing the CC is above the contextual threshold α in the minimal contexts we provided our participants (and, of course, what that threshold might be).

Throughout the discussion, however, Schlenker 2021 appears to assume that there are two classes of predicates: those where the probability is usually above the threshold (including *know, inform*, and *announce*), and those where it is not (including *demonstrate* and *establish*; see p.12 and appendix I). As this division does not fall along the lines of traditionally assumed classes of (non-)factives, an advantage of Schlenker's 2021 analysis over those reviewed in Sections 5.2.1 and 5.2.2 is that it predicts the projection of the CCs of (certain) non-factive predicates (modulo the open questions about condition ii). The analysis correctly predicts that the CCs of *know, inform*, and *announce* are more projective than the CCs of *demonstrate* and *establish*, by virtue of the CCs of the former being usually presupposed, in contrast to the CCs of the latter. It is not clear, however, that the analysis is able to predict the observed variation, because the analysis – even though it does not divide predicates into factive and non-factive ones – nevertheless imposes a binary, categorical distinction between predicates. Crucially, the by-predicate variation reported in Sections 3.2 and 3.3 is not captured by an analysis that does not make more fine-grained distinctions between the meanings of clause-embedding predicates,

as discussed extensively in Degen and Tonhauser 2022. First, there is projection variation between the supposedly presupposed CCs (under all four operators, the CC of *know* is more projective than that of *inform*, which is more projective than that of *announce*). Furthermore, as discussed in Degen and Tonhauser 2022, the CCs of the supposedly nonpresupposed CCs (of *establish* and *demonstrate*) are projective when compared to nonprojective main clause content. In the experiments reported on in Section 2, the mean certainty rating of the CC of *announce*, which is assumed to usually be presupposed, is .55, and that of *demonstrate*, which is not assumed to be presupposed, is .43 (both means are aggregated across entailment-canceling operators). It is not clear that this particular difference in mean certainty rating motivates analyzing the CC of *announce* as presupposed in contrast to that of *demonstrate*.

5.2.4. Interim summary

There is currently no projection analysis on the market that predicts the by-predicate and by-operator variation we observed in the experiments from Section 2. Degen and Tonhauser 2022 suggested that an empirically adequate analysis to capture the observed by-predicate variation requires consideration of "more fine-grained distinctions [between clause-embedding predicates] that are based on the lexical meaning and discourse use of clause-embedding predicates" (p. 585). Further, an account of the observed by-predicate variation in the effects of entailment-canceling operators will need to consider how semantic and pragmatic properties of predicate meanings interact with entailment-canceling operators. We offer some suggestions based on the data from our experiments in the next section.

5.3. Lexical patterns

The results of our experiments suggest that there is by-predicate variation in the effect of entailment-canceling operator. In this section, we suggest that the observed variation may be predictable, at least in part, from lexical semantic and pragmatic properties shared by subsets of the 20 clause-embedding predicates featured in our experiments. Specifically, we observe that the projection of the CCs of the 20 predicates we investigated exhibit (what we call) 'projection patterns', such that the CCs of several predicates exhibit the same or at least a similar by-operator projection variation. The six patterns we identified based on our results are shown in Figure 6, which shows mean certainty ratings for the four operators by projection pattern.

We suggest that these patterns are not accidental, but that the predicates that share a pattern also share lexical semantic and pragmatic meanings. For instance, the non-veridical predicates *pretend* and *think* exhibit the 'Negation high' pattern, shown in panel (a) of Figure 6. These are the only predicates that are most projective under negation compared to all other operators. This generalization might be derivable from the observation that there is an anti-veridical inference for the CC of both *think* and *pretend*. For *think*, this anti-veridical inference can arise from the fact that *think* has a veridical alternative in *know* (e.g., Heim 1991, Chemla 2008). We tentatively hypothesize that this can lead to an inference that the CC is false in many environments, but not under negation. For *pretend*, one might either assume a similar alternative or investigate whether this predicate entails that the speaker assumes that the CC is false.

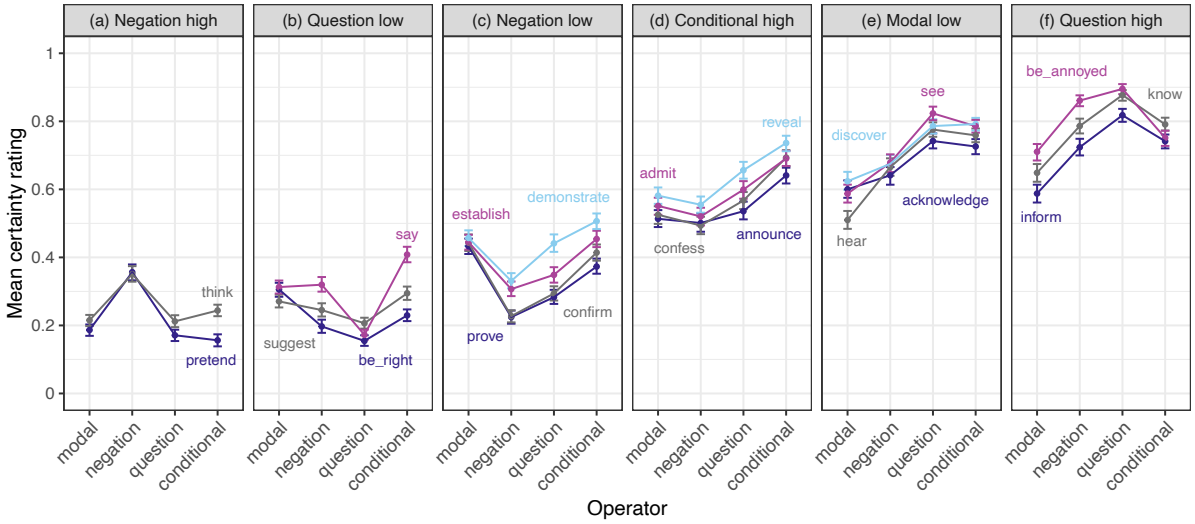

Figure 6: Mean certainty ratings by operator and predicate, by projection pattern. Error bars indicate 95% bootstrapped confidence intervals.

The predicates *suggest, be right*, and *say* ('Question low' pattern, panel (b)) are the only predicates which are least projective in polar questions. We tentatively suggest that these predicates can interact with the pragmatics of polar questions in a way that can lead to an inference of incredulity towards the CC. However, there are also some differences between these predicates. The veridical *be right* is most projective under modals (M > C > N > Q), whereas *say* and *suggest* are most projective in conditional antecedents (C > {M, N} > Q).

The two ∪-shaped patterns in panels (c) and (d) are characterized by relatively low projection ratings with negation and high ratings with conditionals. While there is fine-grained variation with regard to which by-operator differences were statistically supported, (c) and (d) can be distinguished by their ratings for the modal. The the CCs of *admit*, *announce*, *confess* and *reveal* exhibit relatively low projection ratings when embedded under *perhaps* (C > Q > {M, N}). In contrast, the modal ratings for the CCs of *confirm, demonstrate, establish*, and *prove* are relatively high (*confirm*, *prove*: M > C > Q > N; *demonstrate:* C > {M, Q} > N; *establish:* {C, M} > {Q, N}). For an explanation of the high conditional ratings, one might examine how the discourse effect of a conditional interacts with a change-of-state meaning component of these inferential and communicative predicates. A possible explanation for the relatively low negation ratings is that these predicates can be interpreted relative to contextual assumptions that lead to a neg-raising type inference more readily than others (so that, for instance, not announcing *p* amounts to communicating not *p*, or not proving *p* amounts to inferring not *p*).

Finally, the two ∩-shaped patterns in (e) and (f) are characterized by relatively high projection ratings with questions, and low ratings with *perhaps*. The two patterns differ in their ratings for conditionals. The predicates *acknowledge, discover, hear,* and *see*, which are associated with a change of some informational state, show relatively high ratings with conditionals ({Q, C} > N > M). The ratings for conditionals are relatively lower for the predicates *inform*, *know* (Q > {N, C} > M), and *be annoyed* (Q > N > C > M), whose CCs are among the most projective.

Although the patterns we have identified here are tentative and only based on few predicates, we believe that future investigations into shared projection patterns across a wider range of predicates is a fruitful enterprise for future investigations of projection inferences.

6. Conclusion

This paper investigated variation in projection from under the four entailment-canceling operators that have traditionally been used in the family-of-sentences diagnostic for projection, namely negation, polar questions, epistemic modals, and conditional antecedents. The results of our experiments suggest that the projection of the contents of the clausal complements of clause-embedding predicates varies across these operators. As discussed, there is currently no projection analysis on the market that is able to predict the observed by-predicate variation or the by-operator variation. The results of our experiments also extend a result of Degen and Tonhauser 2022, that projection ratings in polar questions do not categorically distinguish factive from non-factive predicates, to cases with negation, the epistemic possibility modal *perhaps*, and conditional antecedents. This results strengthens the conclusion of Degen and Tonhauser 2022 that there is, to date, no empirical evidence for a coherent class of factive predicates.

References

Abrusán, M. (2011). Predicting the presuppositions of soft triggers. *Linguistics and Philosophy 34*, 491–535.

Beaver, D. (2010). Have you noticed that your belly button lint colour is related to the colour of your clothing? In R. Bäuerle, U. Reyle, and T. Zimmermann (Eds.), *Presuppositions and Discourse: Essays offered to Hans Kamp*, pp. 65–99. Oxford: Elsevier.

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software 80*(1), 1–28.

Chemla, E. (2008). An epistemic step for anti-presuppositions. *Journal of Semantics 25*(2), 141–173.

Chierchia, G. and S. McConnell-Ginet (1990). *Meaning and Grammar: An Introduction to Ssemantics*. Cambridge, MA: MIT Press.

de Marneffe, M.-C., C. D. Manning, and C. Potts (2012). Did it happen? The pragmatic complexity of veridicality assessment. *Computational Linguistics 38*(2), 301–333.

de Marneffe, M.-C., M. Simons, and J. Tonhauser (2019). The CommitmentBank: Investigating projection in naturally occurring discourse. *Proceedings of Sinn und Bedeutung 23*(2), 107–124.

Degen, J. and J. Tonhauser (2022). Are there factive predicates? An empirical investigation. *Language 98*(3), 552–591.

Djärv, K. and H. A. Bacovcin (2020). Prosodic effects on factive presupposition projection. *Journal of Pragmatics 169*, 61–85.

Djärv, K., J. Zehr, and F. Schwarz (2018). Cognitive vs. emotive factives: An experimental differentiation. *Proceedings of Sinn und Bedeutung 21*(1), 367–385.

Grove, J. and A. S. White (2023). Factivity, presupposition projection, and the role of discrete

knowledge in gradient inference judgments. Unpublished ms., University of Rochester, retrieved on Feb 10, 2024 from `https://lingbuzz.net/lingbuzz/007450`.

Heim, I. (1983). On the projection problem for presuppositions. In P. H. Portner and B. H. Partee (Eds.), *Formal Semantics: The Essential Readings*, pp. 249–260. New York: Wiley.

Heim, I. (1991). Artikel und Definitheit. In A. von Stechow and D. Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, pp. 487–535. de Druyter: Berlin.

Karttunen, L. (1971). Some observations on factivity. *Research on Language & Social Interaction 4*(1), 55–69.

Kiparsky, P. and C. Kiparsky (1970). Fact. In M. Bierwisch and K. E. Heidolph (Eds.), *Progress in Linguistics*, pp. 143–173. The Hague: Mouton.

Lenth, R. V. (2023). *emmeans: Estimated marginal means, aka least-squares means*. R package version 1.8.8.

Mahler, T. (2019). Does at-issueness predict projection? It's complicated! *Proceedings of NELS 49*(1), 245–255.

Mahler, T. (2020). The social component of projection behavior of clausal complements. *Linguistic Society of America 5*, 777–791.

R Core Team (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics 5*, 1–69. Reprint of 1998 revision of 1996 OSU Working Papers publication.

Ross, A. and E. Pavlick (2019). How well do NLI models capture verb veridicality? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 2230–2240.

Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics 2*(3), 1–78.

Schlenker, P. (2021). Triggering presuppositions. *Glossa: A Journal of General Linguistics 6*(1), 35.

Sieker, J. and T. Solstad (2022). Projective variability of (semi)factive verbs in family of sentence contexts: A rating study. In *Proceedings of the 23rd Amsterdam Colloquium*, pp. 282–289.

Simons, M., D. Beaver, C. Roberts, and J. Tonhauser (2017). The best question: Explaining the projection behavior of factives. *Discourse Processes 54*(3), 187–206.

Smith, E. A. and K. C. Hall (2011). Projection diversity: Experimental evidence. In *Proceedings of the ESSLLI 2011 Workshop on Projective Content*, pp. 156–170.

Smith, E. A. and K. C. Hall (2014). The relationship between projection and embedding environment. *Proceedings of CLS 48*, 563–578.

Smithson, M. and J. Verkuilen (2006). A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods 11*(1), 54.

Spenader, J. (2002). *Presuppositions in Spoken Discourse*. PhD Thesis, Stockholm University.

Stevens, J., M.-C. de Marneffe, S. Speer, and J. Tonhauser (2017). Rational use of prosody predicts projectivity in manner adverb utterances. In *39th Annual Meeting of the Cognitive Science Society (CogSci 2017)*, pp. 1144–1149.

Tonhauser, J. (2016). Prosodic cues to presupposition projection. *Proceedings of SALT 26*,

934–960.

Tonhauser, J., D. Beaver, C. Roberts, and M. Simons (2013). Toward a taxonomy of projective content. *Language 89*(1), 66–109.

Tonhauser, J., D. I. Beaver, and J. Degen (2018). How projective is projective content? Gradience in projectivity and at-issueness. *Journal of Semantics 35*(3), 495–542.

van der Sandt, R. A. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics 9*(4), 333–377.

White, A. S. and K. Rawlins (2018). The role of veridicality and factivity in clause selection. *Proceedings of NELS 3*, 221–234.

Xue, J. and E. Onea (2011). Correlation between presupposition projection and at-issueness: An empirical study. In *Proceedings of the ESSLLI 2011 Workshop on Projective Meaning*, pp. 171–184.