

Prior beliefs modulate projection

Judith Degen^{*} and Judith Tonhauser[°]

^{*}Stanford University

[°]The Ohio State University / University of Stuttgart

February 20, 2021

Abstract

Beliefs about the world affect language processing and interpretation in several empirical domains. In two experiments, we tested whether subjective prior beliefs about the probability of utterance content modulate *projection*, that is, listeners' inferences about speaker commitment to that content. We find that prior beliefs predict projection at both the group and the by-participant level: the higher the prior belief in a content, the more speakers are taken to be committed to it. This result motivates the integration of formal analyses of projection with cognitive theories of language understanding.

Keywords: experimental semantics; experimental pragmatics; projection

Word count: 3923 words without main body reference citations

1 Introduction

Psycholinguistic work has documented several ways in which probabilistic beliefs about the world, often termed *world knowledge*, affect language processing (e.g., Chambers, Tanenhaus, Eberhard, Filip, & Carlson, 2002; Hagoort, Hald, Mastiaansen, & Petersson, 2004; Hald, Steenbeek-Planting, & Hagoort, 2007; Warren & McConnell, 2007), including syntactic ambiguity resolution (e.g., Chambers, Tanenhaus, & Magnuson, 2004; Bicknell & Rohde, 2014), reference resolution (e.g., Winograd, 1972; Hanna & Tanenhaus, 2004), genericity (e.g., Tessler & Goodman, 2019), scalar implicature (e.g., Degen, Tessler, & Goodman, 2015), underinformativity implicatures (Kravtchenko & Demberg, 2015), and the production of redundant referring expressions (Mitchell, Reiter, & Van Deemter, 2013; Westerbeek, Koolen, & Maes, 2015; Rubio-Fernández, 2016; Degen, Hawkins, Graf, Kreiss, & Goodman, 2020). In contrast, formal linguistic research on meaning in the tradition of Montague (1973), which is devoted to specifying how meanings of expressions are computed from the meanings of the parts of the expressions, the way the parts are combined, and the contexts in which the expressions are used, has often sidelined world knowledge, as non-linguistic, encyclopedic knowledge that must enter into the meaning computation, but whose effect has eluded systematic investigation and formalization (for relevant discussion see, e.g., Dowty, 1986; Peeters, 2000; Beaver, 2001; Hobbs, 2019).¹ In this paper, we provide empirical evidence from English that *projection*, a key topic in linguistic research on meaning, is systematically modulated by listeners'² subjective beliefs about the world. This provides further impetus for accounts of meaning computation to include a mechanism for integrating subjective prior beliefs. We provide a sketch of such an account at the end of this paper.

¹Because *knowledge* implies justified true belief but subjective beliefs need not be accurate to affect language processing in systematic ways, we henceforth avoid the term *world knowledge* and instead refer to (*subjective prior*) *beliefs about the world*.

²We include readers, writers, and signers in the terms *listener* and *speaker*.

To introduce projection, consider first that speakers can present themselves, through their utterances, as believing that particular content is true, that is, as committed to that content. Listeners, in turn, regularly draw inferences about which content speakers present themselves as committed to. For instance, if a speaker utters *Sam knows that it's raining*, listeners typically infer that the speaker is committed to the following two contents: (i) the content of the complement (CC) of *know*, that it's raining; and (ii) the content of the matrix clause, that Sam knows (i). In formal research on meaning, the inference to (i) is attributed to the speaker having uttered the sentence, and the inference to (ii) is attributed to the lexical meaning of *know*, specifically, that only true content can be known (e.g., Chierchia & McConnell-Ginet, 1990). The puzzle is that the inference to (i) persists even when the speaker inquires about what Sam knows, as in *Does Sam know that it's raining?*, or when the speaker denies Sam's knowledge, as in *Sam doesn't know that it's raining*. Because Sam's knowledge is questioned or even denied in these variants, that is, the inference to (ii) does not persist, these inferences to (i) cannot be attributed to the aforementioned lexical meaning of *know*. This phenomenon of speaker commitment to utterance content that occurs in negated sentences or questions is termed *projection*. Decades of research in formal semantics have aimed to explain why content projects (e.g., Langendoen & Savin, 1971; Beaver & Geurts, 2014).

While content is standardly taken to either project or not (Beaver & Geurts, 2014), recent experimental work suggests that projection is gradient: listeners' inferences about speaker commitment to utterance content vary in strength and are modulated by the expression (e.g., *know* vs. *discover* vs. *announce*) and by contextual factors including the discourse status of the content and the prosody of the utterance (for an overview see Tonhauser, Beaver, & Degen, 2018). The hypothesis that listeners' prior beliefs modulate projection was initially put forth by Stevens, de Marneffe, Speer, and Tonhauser (2017) and Tonhauser et al. (2018), who observed by-item projection variability for different CCs of clause-embedding predicates like *know* and *discover*. They argued that one source of the observed variability may be that more a priori likely content (*Kim flew to New York*) projects more strongly than less a priori likely content (*Kim flew to the moon*) when realized as the CC of a clause-embedding predicate (as in *Did John discover that Kim flew to New York/the moon?*). This idea can straightforwardly be made sense of under recent Bayesian accounts that treat pragmatic utterance interpretation as a matter of combining uncertain prior beliefs about the world with uncertain beliefs about likely speaker production choices via Bayes' rule (Goodman & Frank, 2016; Degen et al., 2015): a CC that is more likely a priori (*before* observing an utterance) is also more likely a posteriori (*after* observing an utterance).

There is conflicting evidence for the hypothesis that prior beliefs modulate projection. Support for the hypothesis comes from Mahler (2020), who investigated the projection of politically charged CCs of English clause-embedding predicates. For example, the politically charged content in (1) is that Obama improved/damaged the American economy. The prior probability of the content was manipulated by the speaker (Cindy in (1)) speaking at the club meeting of either the College Republicans or Democrats.

- (1) Cindy, at the College Republicans/Democrats club meeting:
 Ben doesn't know that . . .
 a. . . . Obama improved the American economy.
 b. . . . Obama damaged the American economy. (Mahler, 2020, 784f.)

Higher prior probability content (e.g., a liberal content like (1a) uttered by a Democrat) was more projective than lower prior probability content (e.g., a liberal content uttered by a Republican).

In contrast, Lorson (2018) did not find empirical support for the hypothesis that listeners' prior beliefs modulate projection in a study of the projection of the pre-state content of the English change of state verb *stop*. Prior probability was manipulated through gender stereotypes reported in Boyce, von der Malsburg, Poppels, and Levy (2018). For instance, because men are more likely than women to be plumbers, the pre-state content of (2a), that James has worked as a plumber, was hypothesized to be more projective than the

pre-state content of (2b), that Linda has worked as a plumber.

- (2) a. Did James stop working as a plumber?
- b. Did Linda stop working as a plumber? (Lorson, 2018, 38)

Several differences between Mahler’s and Lorson’s investigations could be implicated in the differential support for the hypothesis: a) the projective content investigated (CCs vs. pre-state content of *stop*); b) stimulus type (negated sentences vs. questions); c) the manipulation of prior beliefs (political party affiliation vs. gender stereotypes); and d) how explicitly the prior-manipulating information was provided to participants (statement of political party affiliation vs. use of a male or female name to indicate gender). The two experiments reported on in this paper provide additional support for the hypothesis that prior beliefs modulate projection. The experiments included 20 clause-embedding predicates (rather than just 7, as in Mahler, 2020) and the prior belief manipulation involved 20 properties of individuals (rather than just political party affiliation, as in Mahler, 2020, or gender, as in Lorson, 2018). Furthermore, we tested the hypothesis both at the level of the individual and of the group: Exp. 1 investigated the effect of prior beliefs on projection by measuring prior probability and projection in a within-participant design. In Exps. 2a and 2b, prior probability and projection were measured in separate groups, as in Mahler, 2020 and Lorson, 2018.

2 Experiment 1

This experiment tested whether higher prior probability content is more likely to project. Prior probability and projection ratings were collected for the contents of 20 clauses that realized the complements of 20 clause-embedding predicates.³

2.1 Methods

Participants 300 participants with U.S. IP addresses and at least 99% of previous HITs approved were recruited on Amazon’s Mechanical Turk platform (ages: 18-82, median: 35.5; 119 female, 179 male, 1 other, 1 undeclared). They were paid \$1.80.

Materials and procedure The prior probability and projection of the contents of 20 clauses were measured in separate blocks. Each clause (e.g., *Julian dances salsa*) was paired with two facts between participants: The content of the clause was expected to have a higher prior probability in the presence of one fact (e.g., *Julian is Cuban*) than of the other (e.g., *Julian is German*). See Supplement A for the full set of clauses and facts.

In the prior block, the 20 clauses were realized as the complements of *How likely is it that...?* questions. As shown in Fig. 1a, each target stimulus consisted of one of the two facts for that clause and the *How likely is it that...?* question. Participants read the fact and assessed the likelihood of the content, given the fact. They gave their responses on a slider marked ‘impossible’ at one end (coded as 0) and ‘definitely’ at the other (coded as 1).

In the projection block, target stimuli consisted of a fact and a polar question that was uttered by a named speaker, as shown in Fig. 1b. The polar questions were formed by realizing the 20 clauses as the complements of the 20 clause-embedding predicates in Fig. 1c. Participants were told to imagine that they are at a party and that, on walking into the kitchen, they overhear somebody ask somebody else a question.

³The experiments, data and R code for generating the figures and analyses of the experiments reported on in this paper are available at <https://github.com/judith-tonhauser/projective-probability>. Exp. 1 was pre-registered: <https://osf.io/vd9ru/>. All experiments were conducted with approval from the IRB of The Ohio State University and informed consent was obtained.

Projection was measured using the ‘certain that’ diagnostic (Djäv & Bacovcin, 2017; Tonhauser et al., 2018; Lorson, 2018; Mahler, 2020): participants were asked to rate whether the speaker was certain of the CC, taking into consideration the fact. They gave their responses on a slider marked ‘no’ at one end (coded as 0) and ‘yes’ at the other (coded as 1). Greater speaker commitment to the CC should result in higher slider ratings.

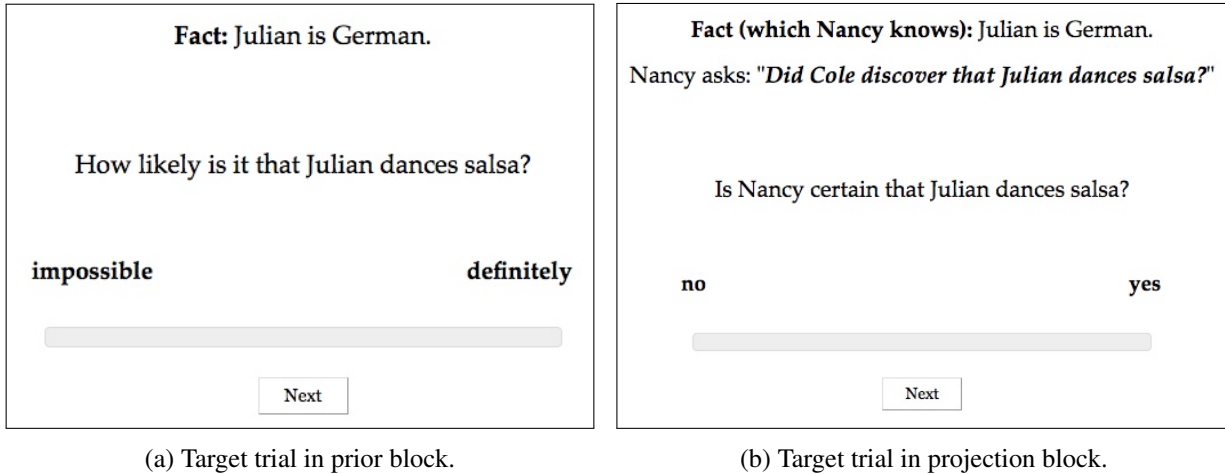


Figure 1: Sample trials and 20 clause-embedding predicates in Exp. 1.

The projection block also included 6 control trials, which functioned as attention checks. The content of these stimuli was expected not to project: For example, in Fig. 1e, the speaker is not committed to the main clause content, that Samantha has a new hat. The same 6 main clauses were also used to form 6 filler trials in the prior block; a sample stimulus is given in Fig. 1d. These filler stimuli were not used to assess participants’ attention. For the full set of stimuli see Supplement A.

Each participant’s stimulus set was semi-randomly generated by first randomly pairing up the 20 predicates and clauses. Half of the stimuli were then randomly assigned the respective clause’s higher-probability fact, and half its lower-probability fact. Participants completed a total of 52 trials: 20 target trials in each

block, 6 control trials in the projection block, and 6 filler trials in the prior block. Each participant completed the same 6 filler and control trials. Block order and within-block trial order were randomized.

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

Data exclusion Data was excluded based on self-declared non-native speaker status and other criteria given in Supplement B, leaving 5,720 data points from 286 participants to be analyzed (ages 18-82; median: 35.5; 116 female, 186 male, 1 other, 1 undeclared).

2.2 Results and discussion

Prior beliefs. Fig. 2 shows the mean prior probabilities of the 20 contents by fact. We conducted a mixed-effects linear regression predicting slider rating from dummy-coded fact type (reference level: ‘lower probability’) and random by-item and by-participant intercepts and slopes for fact type.⁴ Each content’s mean prior probability was rated as higher when it was presented with its higher probability fact than when it was presented with its lower probability fact ($\beta = 0.45$, $SE = 0.01$, $t = 31.12$, $p < .0001$). This suggests that the manipulation of the prior probability of the 20 contents was successful.




Figure 2: Mean prior probability by content and fact in Exp. 1. Error bars indicate 95% bootstrapped confidence intervals. Transparent dots indicate individual participant ratings.

Do prior beliefs modulate projection? Fig. 3 shows the mean certainty ratings for the CCs by predicate and by fact, as well as the mean certainty rating for the main clause controls (abbreviated ‘MC’). We conducted a mixed effects linear regression predicting certainty ratings from dummy-coded fact type (reference

⁴All analyses were conducted in R (R Core Team, 2016) using the lme4 package (Bates, Mächler, Bolker, & Walker, 2015).

level: ‘lower probability’) and random by-item and by-participant intercepts and slopes for fact type. The mean certainty ratings were higher for contents presented with higher probability facts than for contents presented with lower probability facts ($\beta = 0.14$, $SE = 0.01$, $t = 12.24$, $p < .0001$). The same was true when using the group-level by-item mean prior belief as a predictor ($\beta = 0.31$, $SE = 0.02$, $t = 12.58$, $p < .0001$). This suggests that participants’ prior beliefs about content probability systematically modulated the extent to which they take the speaker to be committed to that content.

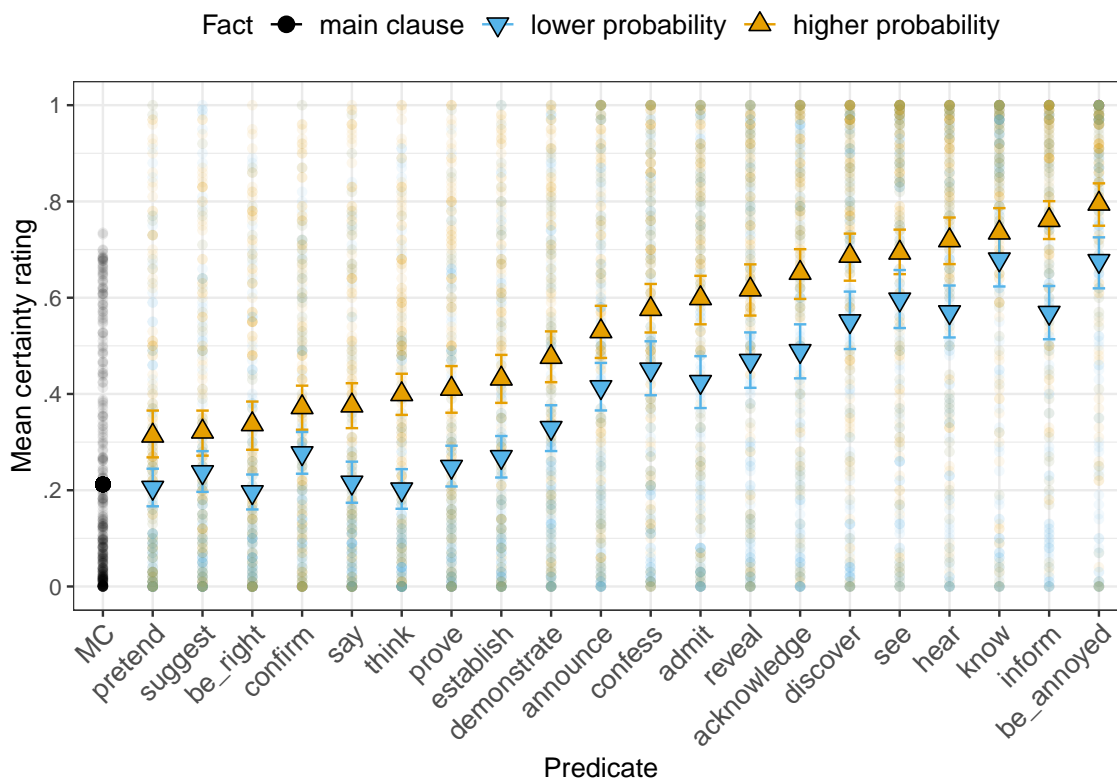


Figure 3: Mean certainty ratings by predicate and prior probability of the content of the complement in Exp. 1. Error bars indicate 95% bootstrapped confidence intervals. Light dots indicate participants’ ratings.

We also replicated Tonhauser and Degen’s (under review) result of by-predicate variability in the projection of the CC: for instance, the CC of *be annoyed* was more projective than that of *discover*, which in turn was more projective than that of *announce*. The Spearman rank correlation between the mean certainty ratings in Exp. 1 (collapsing over facts) and Tonhauser and Degen’s (under review) Exp. 1a is .991; see Supplement C for a visualization. Exp. 1 thereby also provides further evidence for the systematic influence of the predicate on projection. Crucially, the effect of the prior was observable independently of predicate.

Closer inspection of Fig. 2 reveals by-participant variability in prior probability ratings, suggesting that individual participants’ prior beliefs may not align with the prior probability classification assumed in Fig. 3. For example, given a particular content (*Julian dances salsa*), it is possible that one participant’s prior probability rating was lower than that of another participant, even though the first participant was presented with the higher probability fact (*Julian is Cuban*) and the second one with the lower probability fact (*Julian is German*). Fig. 4 shows participants’ certainty ratings by their individual prior probability ratings. To investigate whether prior beliefs modulate projection at the by-participant level, we conducted the same mixed-effects analysis reported above, but used participants’ individual, continuous prior probability ratings as the fixed effect prior predictor. Again, higher-prior-probability CCs were more likely to project ($\beta =$

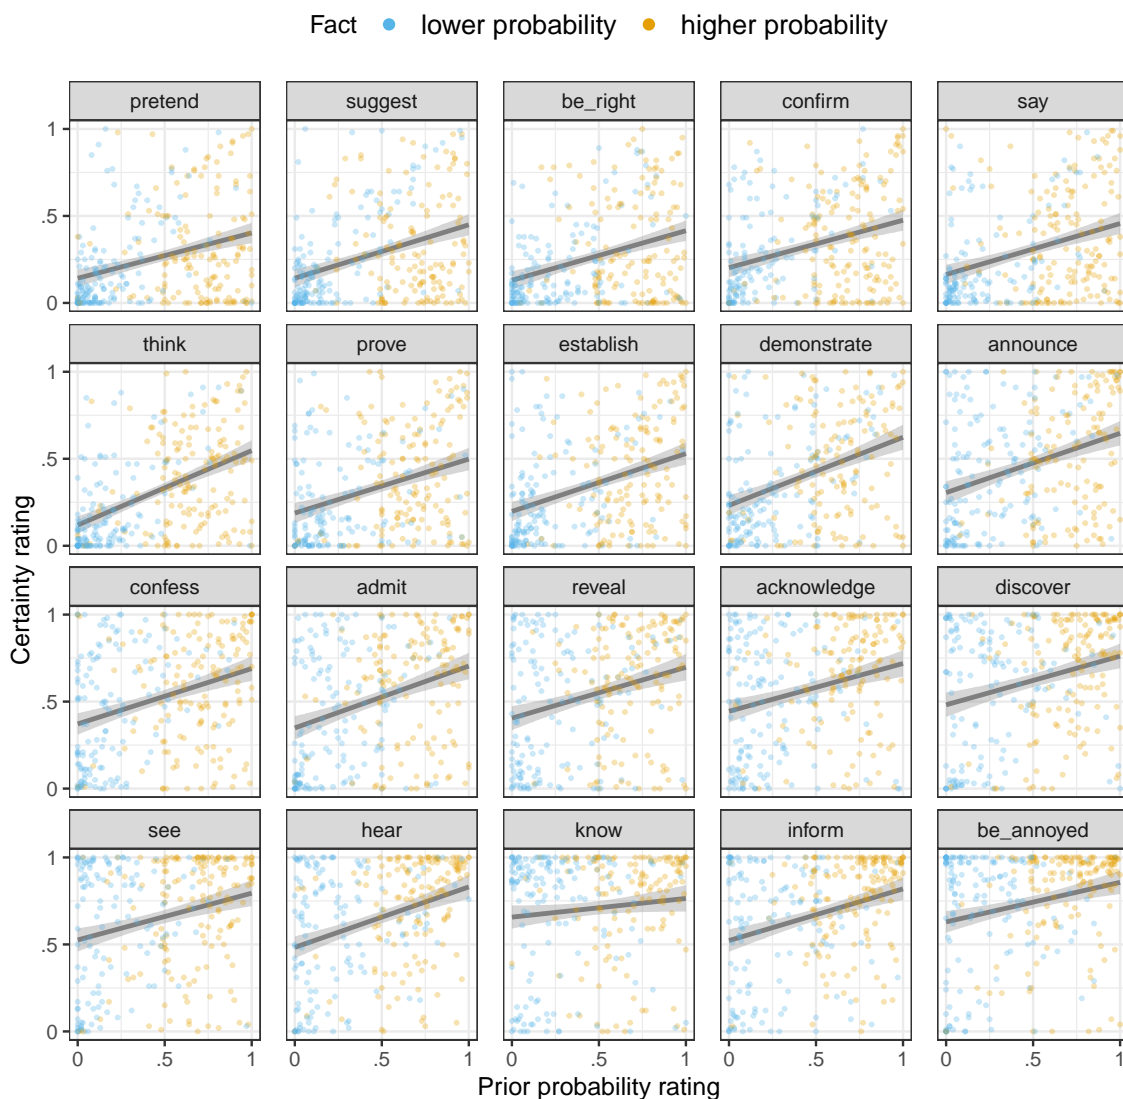


Figure 4: Certainty ratings against individual prior probability ratings for each predicate in Exp. 1. Linear smoothers with 95% confidence intervals are overlaid.

0.28, $SE = 0.02$, $t = 13.85$, $p < .0001$). This suggests that prior beliefs modulate projection even at the by-participant level. In fact, a Bayesian Information Criterion (BIC) model comparison revealed that the individual-level model better captured the variance in the data (categorical model BIC: 2654; group-level model BIC: 2586; individual-level model BIC: 2291), suggesting that individual listeners' prior beliefs systematically modulate the extent to which they take the speaker to be committed to a content: the more they believe it, the more they take the speaker to believe it.

The results of Exp. 1 provide empirical support for the hypothesis that higher prior probability content is more likely to project. It is possible, however, that the within-participant design resulted in participants' responses on either block influencing their responses on the other block. To guard against this possibility, we replicated Exp. 1 by collecting prior probability and projection ratings from different groups.

3 Experiment 2

Exps. 2a and 2b measured the prior probability and the projection of the 20 contents of Exp. 1, respectively.

3.1 Methods

Participants Participants with U.S. IP addresses and at least 99% of previous HITs approved were recruited on Amazon’s Mechanical Turk platform. The 95 participants in Exp. 2a (ages: 21-75, median: 33; 45 female, 50 male) were paid 55 cents. The 300 participants in Exp. 2b (ages: 21-72, median: 36; 145 female, 154 male, 1 undeclared) were paid 85 cents.

Materials and procedures The target stimuli of Exp. 2a were identical to those of the prior block of Exp. 1. Each participant saw two control stimuli as attention checks (see Supplement D). The materials of Exp. 2b were identical to those of the projection block of Exp. 1. Trial order in both experiments was random. The procedures of Exps. 2a and 2b were identical to those of the prior and projection blocks of Exp. 1, respectively.

Data exclusion We excluded data based on the criteria given in Supplement B, leaving data from 75 participants to be analyzed in Exp. 2a (1,500 data points; ages 21-75; median: 35; 34 female, 41 male) and from 266 participants in Exp. 2b (5,320 data points; ages 21-72; median: 36; 129 female, 136 male, 1 undeclared).

3.2 Results and discussion

Prior beliefs. Exp. 2a successfully replicated the prior probability manipulation of Exp. 1: contents were rated as more likely when presented with a higher probability fact ($\beta = 0.54$, $SE = 0.04$, $t = 15.07$, $p < .0001$). Fig. 5 shows contents’ mean prior probability ratings in Exp. 2a against those of Exp. 1. The Spearman rank correlation was very high, at $r = .977$. For a visualization of the by-content prior ratings see Supplement D.

Do prior beliefs modulate projection? Mean certainty ratings were higher for contents presented with higher prior probability facts than for contents presented with lower prior probability facts (see Fig. 6). This was true when the prior predictor was entered as a categorical predictor (reference level: ‘lower probability’; $\beta = 0.18$, $SE = 0.01$, $t = 12.81$, $p < .0001$) and when it was entered as a continuous predictor representing group-level prior means ($\beta = 0.34$, $SE = 0.03$, $t = 13.27$, $p < .0001$). Thus, Exp. 2b replicates the critical result of Exp. 1 that prior content probability modulates its projection.⁵ The replication suggests that the result of Exp. 1 is not an artifact of the within-participant design of Exp. 1.

4 General discussion and concluding remarks

We tested whether listeners’ prior beliefs modulate projection. While previous research on this question has yielded conflicting results (Mahler, 2020; Lorson, 2018), we showed in two experiments that content is more likely to project the more a priori likely it is, thus confirming Mahler (2020)’s results and expanding on them in several ways. First, while Mahler (2020) manipulated only the political party affiliation of the speaker,

⁵28 participants took Exp. 2b after taking Exp. 2a two weeks before. Analyses that excluded these participants’ data did not change the results. Exp. 2b also replicated Tonhauser and Degen’s (under review) result that there is by-predicate variability in the projection of the CC; see Supplement C.

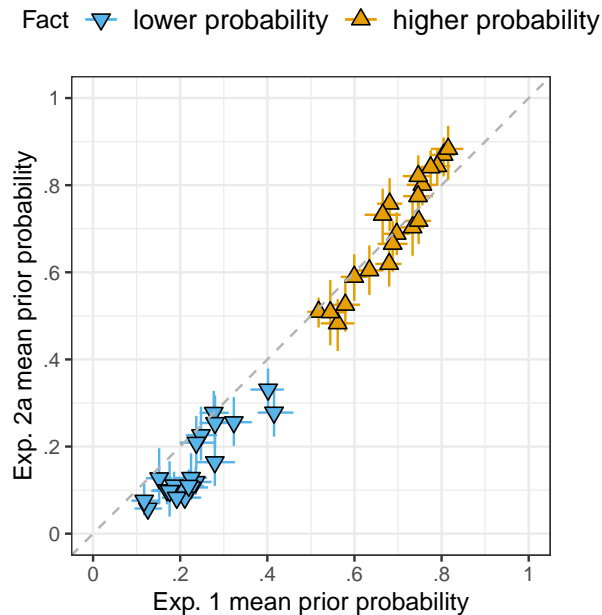


Figure 5: Mean prior probability ratings in Exp. 2a against those of Exp. 1. Error bars indicate 95% bootstrapped confidence intervals.

the manipulation in Exps. 1 and 2 relied on 20 distinct properties of individuals (e.g., whether Julian is more likely to dance salsa if he is German or Cuban). Thus, the results of Exps. 1 and 2 suggest a general effect of prior beliefs on projection. Second, our experiments show that prior beliefs modulate projection for a wider cross-section of clause-embedding predicates, including cognitive (e.g., *know*), emotive (e.g., *be annoyed*), communication (e.g., *announce*), and inferential (e.g., *prove*) predicates. Finally, the within-participant design of Exp. 1 shows that individuals’ prior beliefs better predict projection than group-level beliefs, which in turn better predict projection than the binary categorical beliefs that Mahler (2020) investigated. This suggests that at least some by-participant variability observed in previous projection experiments (see, e.g., Tonhauser et al., 2018; Tonhauser & Degen, under review) may be due to participants assigning different prior probabilities to investigated content.

Our results have two broader implications. First, they suggest that the purview of projection analyses is wider than assumed by current analyses, which typically limit their attention to a narrow subset of clause-embedding predicates, like factive ones (e.g., Heim, 1983; van der Sandt, 1992; Abrusán, 2011, 2016; Romoli, 2015; Simons, Beaver, Roberts, & Tonhauser, 2017). Second, they motivate the development of projection analyses that consider listeners’ variable subjective beliefs about the world. Given the gradient nature of the measured (prior and posterior) beliefs and the uncertainty inherent in the different factors that have been shown to modulate projection (e.g., at-issueness, prosody), probability theory suggests itself as a representational framework within which to model projection. To date, only few probabilistic models of projection have been developed (Qing, Goodman, & Lassiter, 2016; Stevens et al., 2017). In these models, projection is the result of listeners’ reasoning about the common ground that the speaker is assuming and the likely question that was being addressed, respectively. While neither investigated the effect of prior beliefs explicitly, both models are couched within the Rational Speech Act (RSA) framework (Franke & Jäger, 2016; Goodman & Frank, 2016), which standardly assumes that utterance interpretation is modulated by listeners’ prior beliefs. The RSA framework is thus equipped to capture the effects reported here. We see the implementation of projection analyses within RSA as a promising avenue for formalizing the intricate inter-




Figure 6: Mean certainty ratings by predicate and prior probability of the content of the complement in Exp. 2b. Error bars indicate 95% bootstrapped confidence intervals. Light dots indicate participants' ratings.

play of semantic and pragmatic factors in the projection of contents of complements of clause-embedding predicates, including the conventional contribution of predicates, content at-issueness, and subjective prior beliefs about content.

References

- Abrusán, M. (2011). Predicting the presuppositions of soft triggers. *Linguistics & Philosophy*, 34, 491–535. doi: 10.1007/s10988-012-9108-y
- Abrusán, M. (2016). Presupposition cancellation: Explaining the ‘soft-hard’ trigger distinction. *Natural Language Semantics*, 24, 165-202. doi: 10.1007/s11050-016-9122-7
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01
- Beaver, D. (2001). *Presupposition and assertion in dynamic semantics*. Stanford, CA: CSLI Publications.
- Beaver, D., & Geurts, B. (2014). Presupposition. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Stanford University.
- Bicknell, K., & Rohde, H. (2014). Dynamic integration of pragmatic expectations and real-world event knowledge in syntactic ambiguity resolution. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (p. 1216-1221). Austin, TX: Cognitive Science Society.
- Boyce, V., von der Malsburg, T., Poppels, T., & Levy, R. (2018). Implicit gender in the production and comprehension of pronominal references. In *Proceedings of the 31th Annual CUNY Conference on*

Human Sentence Processing.

- Chambers, C., Tanenhaus, M., Eberhard, K., Filip, H., & Carlson, G. (2002). Circumscribing referential domains during real-time language comprehension. *Journal of Memory and Language*, 47, 30-49.
- Chambers, C., Tanenhaus, M., & Magnuson, J. (2004). Actions and affordances in syntactic ambiguity resolution. *Journal of Experimental Psychology*, 30, 687-696.
- Chierchia, G., & McConnell-Ginet, S. (1990). *Meaning and grammar*. Cambridge, MA: MIT Press.
- Degen, J., Hawkins, R. X. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A Bayesian approach to 'overinformative' referring expressions. *Psychological Review*, 127, 591621. doi: 10.1037/rev0000186
- Degen, J., Tessler, M., & Goodman, N. (2015). Wonky worlds: Listeners revise world knowledge when utterances are odd. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Djävrv, K., & Bacovcin, H. (2017). Prosodic effects on factive presupposition projection. *Semantics and Linguistic Theory*, 27, 116-133. doi: 10.3765/salt.v27i0.4134
- Dowty, D. (1986). The effects of aspectual class on the temporal structure of discourse: Semantics or pragmatics? *Linguistics & Philosophy*, 9, 37-61.
- Franke, M., & Jäger, G. (2016). Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft*, 35(1), 3-44.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818-829. doi: 10.1016/j.tics.2016.08.005
- Hagoort, P., Hald, L., Mastiaansen, M., & Petersson, K. (2004). Integration of word meaning and world knowledge in language comprehension. *Science*, 304, 438-441.
- Hald, L. A., Steenbeek-Planting, E., & Hagoort, P. (2007). The interaction of discourse context and world knowledge in online sentence comprehension. Evidence from the N400. *Brain research*, 1146, 210-218.
- Hanna, J., & Tanenhaus, M. (2004). Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive Science*, 28, 105-115.
- Heim, I. (1983). On the projection problem for presuppositions. *West Coast Conference on Formal Linguistics*, 2, 114-125.
- Hobbs, J. (2019). Word meaning and world knowledge. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics* (p. 154-180). Berlin: de Gruyter.
- Kravtchenko, E., & Demberg, V. (2015). Semantically underinformative utterances trigger pragmatic inferences. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Langendoen, T., & Savin, H. (1971). The projection problem for presuppositions. In C. Fillmore & T. Langendoen (Eds.), *Studies in Linguistic Semantics* (p. 54-60). New York: Holt, Rinehart and Winston.
- Lorson, A. (2018). *The influence of world knowledge on projectivity*. (M.A. thesis, University of Potsdam)
- Mahler, T. (2020). The social component of projection behavior of clausal complements. *Linguistic Society of America*, 5, 777-791. doi: 10.3765/plsa.v5i1.4703
- Mitchell, M., Reiter, E., & Van Deemter, K. (2013). Typicality and object reference. *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*, 3062-3067.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, & P. Suppes (Eds.), *Approaches to natural language: Proceedings of the 1970 Stanford workshop on grammar and semantics*. Dordrecht: Reidel.
- Peeters, B. (2000). Setting the scene: Some recent milestones in the lexicon-encyclopedia debate. In B. Peeters (Ed.), *The lexicon-encyclopedia interface* (p. 1-52). New York: Elsevier.

- Qing, C., Goodman, N. D., & Lassiter, D. (2016). A rational speech-act model of projective content. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- R Core Team. (2016). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Romoli, J. (2015). The presuppositions of soft triggers are obligatory scalar implicatures. *Journal of Semantics*, *32*, 173-291. doi: 10.1093/jos/fft017
- Rubio-Fernández, P. (2016, February). How Redundant Are Redundant Color Adjectives? An Efficiency-Based Analysis of Color Overspecification. *Frontiers in Psychology*, *7*. doi: 10.3389/fpsyg.2016.00153
- Simons, M., Beaver, D., Roberts, C., & Tonhauser, J. (2017). The Best Question: Explaining the projection behavior of factive verbs. *Discourse Processes*, *3*, 187-206. doi: 10.1080/0163853X.2016.1150660
- Stevens, J., de Marneffe, M.-C., Speer, S., & Tonhauser, J. (2017). Rational use of prosody predicts projectivity in manner adverb utterances. *Annual Meeting of the Cognitive Science Society*, *39*, 1144-1149.
- Tessler, M. H., & Goodman, N. (2019). The language of generalization. *Psychological Review*, *126*, 395-436.
- Tonhauser, J., Beaver, D., & Degen, J. (2018). How projective is projective content? Gradiance in projectivity and at-issueness. *Journal of Semantics*, *35*, 495-542. doi: 10.1093/jos/ffy007
- Tonhauser, J., & Degen, J. (under review). *Are there factive predicates? An empirical investigation*. (Manuscript under review, Stuttgart University and Stanford University, retrieved from <https://ling.auf.net/lingbuzz/005360>)
- van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, *9*, 333-377. doi: 10.1093/jos/9.4.333
- Warren, T., & McConnell, K. (2007). Investigating effects of selectional restriction violations and plausibility violation severity on eye-movements in reading. *Psychonomic Bulletin & Review*, *14*(4), 770-775.
- Westerbeek, H., Koolen, R., & Maes, A. (2015, July). Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology*, *6*. doi: 10.3389/fpsyg.2015.00935
- Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, *3*(1), 1-191.

Supplemental material for *Prior beliefs modulate projection*

A Experiment 1: Target and control stimuli

This list shows the 20 clauses of the target stimuli alongside their lower and higher probability facts, respectively:

1. Mary is pregnant. Facts: Mary is a middle school student / Mary is taking a prenatal yoga class
2. Josie went on vacation to France. Facts: Josie doesn't have a passport / Josie loves France
3. Emma studied on Saturday morning. Facts: Emma is in first grade / Emma is in law school
4. Olivia sleeps until noon. Facts: Olivia has two small children / Olivia works the third shift
5. Sophia got a tattoo. Facts: Sophia is a high end fashion model / Sophia is a hipster
6. Mia drank 2 cocktails last night. Facts: Mia is a nun / Mia is a college student
7. Isabella ate a steak on Sunday. Facts: Isabella is a vegetarian / Isabella is from Argentina
8. Emily bought a car yesterday. Facts: Emily never has any money / Emily has been saving for a year
9. Grace visited her sister. Facts: Grace hates her sister / Grace loves her sister
10. Zoe calculated the tip. Facts: Zoe is 5 years old / Zoe is a math major
11. Danny ate the last cupcake. Facts: Danny is a diabetic / Danny loves cake
12. Frank got a cat. Facts: Frank is allergic to cats / Frank has always wanted a pet
13. Jackson ran 10 miles. Facts: Jackson is obese / Jackson is training for a marathon
14. Jayden rented a car. Facts: Jayden doesn't have a driver's license / Jayden's car is in the shop
15. Tony had a drink last night. Facts: Tony has been sober for 20 years / Tony really likes to party with his friends
16. Josh learned to ride a bike yesterday. Facts: Josh is a 75-year old man / Josh is a 5-year old boy
17. Owen shoveled snow last winter. Facts: Owen lives in New Orleans / Owen lives in Chicago
18. Julian dances salsa. Facts: Julian is German / Julian is Cuban
19. Jon walks to work. Facts: Jon lives 10 miles away from work / Jon lives 2 blocks away from work
20. Charley speaks Spanish. Facts: Charley lives in Korea / Charley lives in Mexico

In the target stimuli of the projection block of Exp. 1, eventive predicates, like *discover* and *hear*, were realized in the past tense and stative predicates, like *know* and *be annoyed*, were realized in the present tense. The direct object of *inform* was realized by the proper name *Sam*. Each clause-embedding predicate was paired with a unique subject proper name. The speaker of the target stimuli was realized by a randomly sampled unique proper name.

The following list shows the six clauses that were used in the control and filler stimuli of Exp. 1, with their facts. In the prior block, these six clauses were embedded under *How likely is it that...?*. The projection block featured polar questions variants of the clauses.

1. Zack is coming to the meeting tomorrow. Fact: Zack is a member of the golf club.
2. Mary's aunt is sick. Fact: Mary visited her aunt on Sunday.
3. Todd played football in high school. Fact: Todd goes to the gym 3 times a week.
4. Vanessa is good at math. Fact: Vanessa won a prize at school.
5. Madison had a baby. Fact: Trish sent Madison a card.
6. Hendrick's car was expensive. Fact: Hendrick just bought a car.

B Data exclusion

Table A1 presents how many participants' data were excluded from the analyses based on the exclusion criteria. The first column records the experiment, the second ('recruited') how many participants were recruited, and the final column ('remaining') how many participants' data entered the analysis. The 'Exclusion criteria' columns show how many participants' data were excluded based on the two exclusion criteria:

- 'language': Participants' data were excluded if they did not self-identify as native speakers of American English.
- 'controls': In Exps. 1 and 2b, participants' data were excluded if their response mean on the 6 control items was more than 2 sd above the group mean. In Exp. 2a, participants' data were excluded if their response to (3a) was more than 2 sd below the group mean or if their response to (3b) was more than 2 sd above the group mean.

	recruited	Exclusion criteria		remaining
		language	controls	
Exp. 1	300	3	11	286
Exp. 2a	95	8	12	75
Exp. 2b	300	23	11	266

Table A1: Data exclusion in Exps. 1 and 2

C Projection comparisons

Fig. A1 compares the mean certainty ratings of the predicates and main clause controls in Exp. 1, Exp. 2b, and Tonhauser & Degen, under review Exp. 1a (abbreviated 'Exp. 1a TD'). The Spearman rank correlations were .986 (Exp. 2b vs. Exp. 1a TD), .988 (Exp. 1 vs. Exp. 2b) and .991 (Exp. 1a TD vs. Exp. 1).

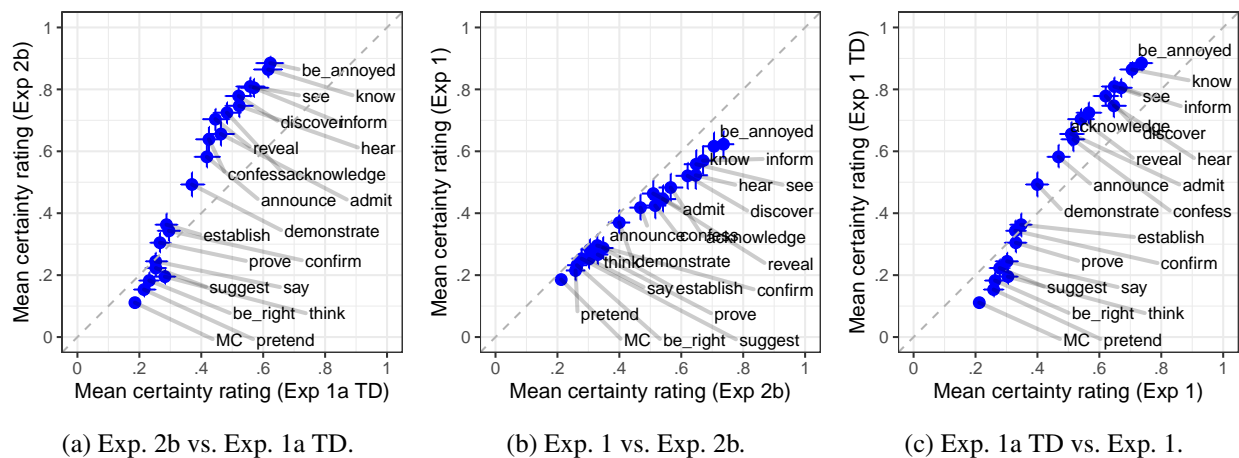


Figure A1: Comparisons of mean by-predicate certainty ratings from Exp. 1, Exp. 2b, and Tonhauser & Degen's Exp. 1a (abbreviated 'Exp. 1a TD'). Error bars indicate 95% bootstrapped confidence intervals. Range of means from current experiments appears compressed compared to that of Tonhauser & Degen because means are collapsed across fact type.

D Experiment 2 supplements

The control items in Exp. 2a are given in (3).

- (3) a. **Fact:** Barry lives in Germany.
How likely is it that Barry lives in Europe?
- b. **Fact:** Tammy is a rabbit.
How likely is it that Tammy speaks Italian and Greek?

Fig. A2 plots the mean prior probabilities of the 20 contents by fact. Participants' ratings are given as light dots. The mean prior probability rating for each content was higher when the content was presented with the higher probability fact than when it was presented with the lower probability fact.

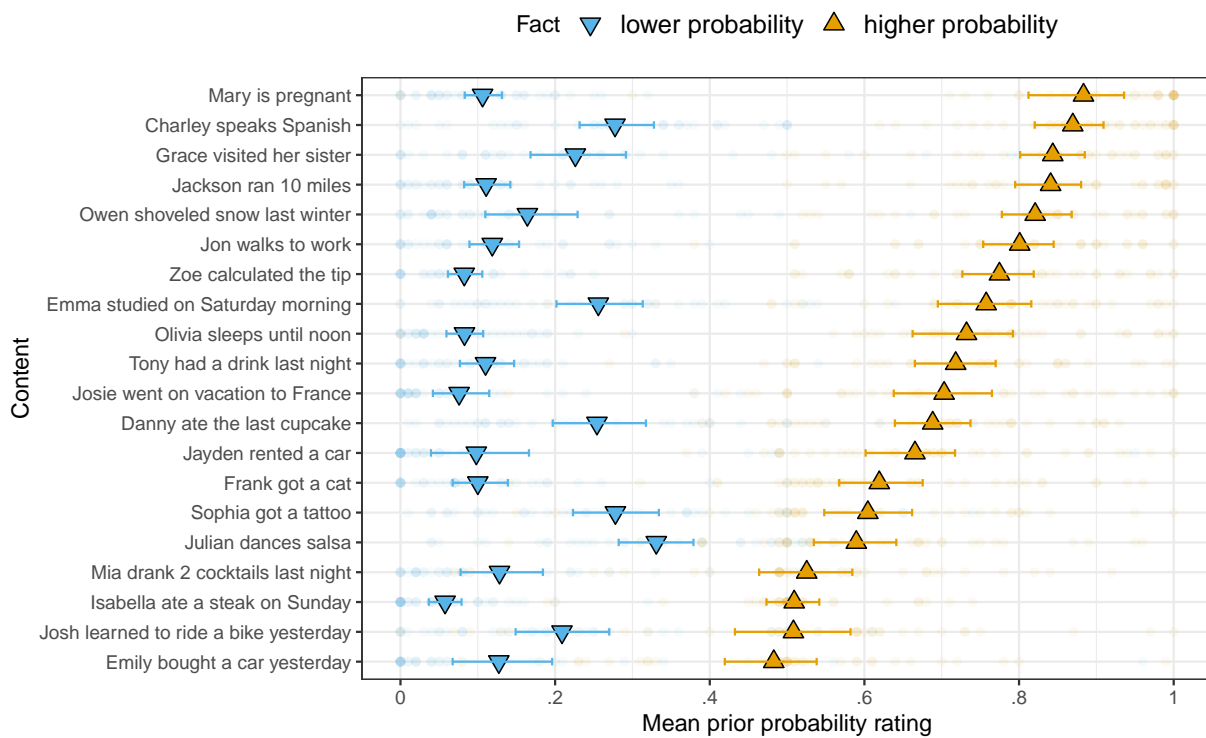


Figure A2: Mean prior probability by content and fact in Exp. 2a. Error bars indicate 95% bootstrapped confidence intervals. Light dots indicate participants' ratings.