Glossa a journal of general linguistics

Tonhauser, Judith, et al. 2020. Evaluative adjective sentences: A question-based analysis of projection. *Glossa: a journal of general linguistics* 5(1): 87. 1–33. DOI: https://doi.org/10.5334/gjgl.701

RESEARCH

Evaluative adjective sentences: A question-based analysis of projection

Judith Tonhauser[1], Marie-Catherine de Marneffe[2] and Judith Degen[3]

[1] Department of English Linguistics, University of Stuttgart, Stuttgart, DE

[2] Department of Linguistics, The Ohio State University, Ohio, US

[3] Department of Linguistics, Stanford University, Stanford, CA, US

Corresponding author: Judith Tonhauser (judith.tonhauser@ling.uni-stuttgart.de)

Two contents of evaluative adjective sentences, like *Kim was smart to watch the movie*, are the prejacent (that Kim watched the movie) and the generalization (that the degree to which Kim watching the movie is smart was higher than the contextual standard of *smart*). The prejacent is standardly analyzed as a presupposition (e.g., Norrick 1978; Barker 2002; Oshima 2009; Kertz 2010). This paper argues against such analyses of the prejacent because, among other things, they do not capture an interaction between the prejacent and the generalization that has not yet been observed for projective content: when the prejacent projects, the generalization does not, and when the prejacent does not project, the generalization does. We develop an analysis according to which the prejacent is not a lexically specified presupposition but is projective to the extent that it is not at-issue with respect to the question addressed by the utterance of the evaluative adjective sentence. In addition to capturing the interaction between the prejacent and the generalization, our question-based projection analysis extends previous such analyses (e.g., Beaver & Clark 2008; Beaver et al. 2017; Simons et al. 2017) by incorporating a novel constraint on the question addressed by an utterance: the more the interpreter takes the truth of content *c* to follow from the common ground a priori, the less likely the question is about *c*. We provide experimental evidence for the analysis and argue that it improves on that of Karttunen et al. (2014), according to which evaluative adjectives are ambiguous.

Keywords: presupposition projection; evaluative adjectives; questions under discussion; experiments

1 Introduction

In an evaluative adjective sentence (EAS) like (1), an evaluative adjective[1] (*stupid*) subcategorizes for a non-pleonastic subject noun phrase (*Feynman*) and a *to*-infinitive (*to dance on the table*); the subject of the predicate (*be stupid*) is the understood subject of the *to*-infinitive (e.g., Wilkinson 1970; Norrick 1978; Barker 2002; Kertz 2010). One of the contents standardly discussed in the literature is (what we refer to as) the prejacent: in (1), the prejacent is that Feynman danced on the table.

(1) Feynman was stupid to dance on the table. (Barker 2002: 18)

The prejacent has traditionally been analyzed as a presupposition (e.g., Norrick 1978; Barker 2002; Oshima 2009; Kertz 2010). Thus, under formal analyses in this tradition,

[1] Evaluative adjectives are adjectives like *stupid, rude* or *fortunate* that, in sentences like *Sam was stupid/rude/fortunate*, convey the speaker's positive or negative evaluation of or attitude towards the denotation of the subject noun phrase.

evaluative adjectives like *stupid* lexically specify that the prejacent must be entailed by or satisfied in the common ground of the interlocutors in order for an EAS to be interpretable (e.g., Heim 1983; van der Sandt 1992). Such analyses are motivated by and straightforwardly account for the variants of (1) in (2) in which the prejacent may project over entailment-canceling operators, such as negation in (2a), a polar question in (2b), the possibility adverb *perhaps* in (2c) and the antecedent of the conditional in (2d). That is, speakers who utter (2a)–(2d) may be committed to the truth of the prejacent, that Feynman danced on the table, even though the evaluative adjective is embedded under an entailment-canceling operator.

(2) Barker (2002: 18f.)
  a. Feynman wasn't stupid to dance on the table.
  b. Was Feynman stupid to dance on the table?
  c. Perhaps Feynman was stupid to dance on the table.
  d. If Feynman was stupid to dance on the table, then tell him.

Recently, Karttunen et al. (2014) provided naturally occurring examples that show that utterances of sentences in which the evaluative adjective is embedded under negation can receive an interpretation according to which the prejacent does not project, i.e., can be interpreted in the scope of negation. For instance, the speaker of (3a) is not committed to living close to their parents and instead communicates that they do not live close to their parents. Likewise, the speaker of (3b) communicates that they did not go stumbling through the junkyard and get hurt.[2]

(3) Karttunen et al. (2014: 235)
  a. I wasn't fortunate to live extremely close to my Mom and Dad for most of my adult life. The closest was when I was in Denver and they were in Garden City, KS.
  b. Now I knew someone was in the junkyard and the cold wind was carrying the cries. I wasn't stupid to go stumbling through the junkyard in the dark and get hurt.

Before discussing why such examples are problematic for analyses of the prejacent as a lexically specified presupposition, we would first like to acknowledge that there is variation in the population of native speakers of American English: whereas speakers are generally able to interpret EASs in which the prejacent does not project, i.e., are able to retrieve the intended interpretations of examples like (3), a sizeable portion strongly prefers to realize such interpretations with variants that include *enough*, as in *I wasn't stupid enough to go stumbling through the junkyard in the dark* for (3b); see also Karttunen (2013) and Karttunen et al. (2014) for this observation. Nevertheless, EASs in which the prejacent

[2] Non-projection of the prejacent is also observed with other entailment-canceling operators, such as the antecedent of a conditional in B's utterance in (i) or the possibility modal *perhaps* in (ii). While this paper limits its attention to EASs embedded under negation, we expect the analysis to generalize to other entailment-canceling operators.

(i) A: I wonder how many people would be banned for using SAM on Killing Floor.
B: If they were stupid to leave it running then maybe a few.
steamcommunity.com/app/730/discussions/0/540744934462316309/

(ii) I am seaching [sic] for a remote cabine in Finland that is available for renting. [...] Region wise I would prefer Lapland or Lakeside. Perhaps we are lucky to see northern lights.
www.tripadvisor.com/ShowTopic-g189896-i442-k11448709-Remoterentalcottage-Finland.html

does not project are part of American English, as evidenced by the existence of naturally occurring examples like (3). Additional evidence comes from the fact that there are native speakers who judge such EASs to be perfectly acceptable[3] and who produce examples like (3). We therefore assume that native speakers of American English are generally able to retrieve both interpretations of EASs even if they might not produce EASs in which the prejacent does not project. Our goal in this paper is to analyze the interpretation of EASs; we briefly return to the observed production variation after developing our analysis.

To account for examples like (3), in which the prejacent does not project, presupposition analyses appeal to local accommodation, a process whereby a presupposition is added to a local context, such as that created by negation: presuppositions can be locally accommodated if the default global accommodation – adding the presupposition to the common ground of the interlocutors – would result in a contradiction, uninformativity or problems with binding (Heim 1982; van der Sandt 1992). In example (3a), for instance, the prejacent, that the speaker lived close to their parents, is locally accommodated under negation because global accommodation would result in a contradiction: according to the context, the closest that the speaker lived to their parents was when the speaker lived in Denver, Colorado, i.e., about 300 miles (480 km) away from their parents in Garden City, Kansas. The prejacent is correctly predicted to be locally accommodated in (3a) and thereby is not a commitment of the speaker.

Negated evaluative adjective sentences (NEASs) like (3) point to two problems for analyses according to which the prejacent is a presupposition. The first problem concerns the prejacent. In example (3b), globally accommodating the prejacent does not result in a contradiction, uninformativity or problems with binding: if the prejacent, that the speaker went stumbling in the junkyard, was added to the common ground, (3b) would mean that the speaker knew that someone was in the junkyard, that they stumbled through the junkyard and got hurt, and that the speaker does not consider these actions stupid (perhaps because these actions led to the person in the junkyard receiving help). Because such an interpretation does not result in a contradiction, uninformativity or problems with binding, analyses of the prejacent as a presupposition incorrectly predict that the prejacent of (3b) is globally accommodated, i.e., is a commitment of the speaker.[4]

The second problem concerns content other than the prejacent: here, the two prior analyses that explicitly consider content other than the prejacent, Oshima (2009) and Barker (2002), do not make sufficiently strong predictions. Consider first Oshima (2009). On this analysis, the prejacent is a presupposition and what is asserted is the following complex content: from the prejacent it can be inferred that the denotation of the subject is in the extension of the evaluative adjective (p.371). That is, (2a) is predicted to presuppose that Feynman danced on the table and to assert that it cannot be inferred from Feynman dancing on the table that Feynman is stupid. Similarly then, (3b) is predicted to convey that the speaker did not stumble through the junkyard (if the prejacent is locally accommodated) and to assert that it cannot be inferred from the speaker stumbling through

[3] The audiences to which this research was presented over the years always included native speakers that judged examples like (3) to be acceptable. More systematic evidence for the existence of such speakers comes from the acceptability rating study presented in Supplementary file 1: of the 94 self-reported native speakers of American English that participated in the study, about 20–30% judged negated evaluative adjective sentences without *enough* to be acceptable under a non-projecting interpretation of the prejacent.

[4] It may be possible to account for the non-projection of the prejacent on the basis of plausibility considerations: von Fintel (2008: 162), for instance, proposed that "what gets accommodated depends on the best guess of the hearers about what the speaker might have intended as the adjustment to the common ground that would admit the asserted sentence". The cues to the projection of the prejacent identified in this paper may help flesh out the reasoning process that interpreters/readers undergo in interpreting EASs. Plausibility considerations do not, however, address the second problem that presuppositional analyses face.

the junkyard that the speaker is stupid. This assertion is too weak: (3b) means that the speaker stumbling through the junkyard would be stupid.

Next consider Barker (2002). This analysis considers the prejacent, which is taken to be presupposed, and a content that we refer to as the generalization: the generalization of (2a) is that the degree to which Feynman dancing on the table is stupid was higher than the contextual standard of *stupid*; that of (3b) is that the degree to which the speaker stumbling through the junkyard in the dark is stupid was higher than the contextual standard of *stupid*. On Barker's (2002) dynamic semantic analysis, the update effect of a NEAS like (2a) is to check that, for each world, the presupposed prejacent is true and to filter out those worlds in which the generalization is false; thus, only those worlds remain in which Feynman dancing on the table was not considered stupid. Similarly then, the update effect of (3b), in which the prejacent is locally accommodated, is to filter out those worlds in which the contextual standard of *stupid* is too high for the speaker's participation in the event of stumbling through the junkyard to have counted as stupid. The worlds that remain are ones in which the speaker did not stumble through the junkyard and in which the speaker stumbling through the junkyard was not stupid. This, again, is not what (3b) is understood to mean.

The examples in (2) and (3) reveal a remarkable interaction that has not yet been observed for projective content and that was not noted in Karttunen et al. (2014): when the prejacent projects, the generalization does not, as in (2b), and when the prejacent does not project, the generalization does, as in (3b). To illustrate that this behavior is strikingly different from the interpretation of utterances of other sentences that give rise to projective content consider (4). When the content of the complement of *know*, that the meeting was canceled, projects, what is denied is Sam's knowledge of this content. The critical difference between EASs and examples like (4) comes out when the content of the complement does not project, as is brought out, for instance, by continuing (4) with …*he, like all of us, is in the dark about whether the meeting will take place*. In this case, Sam's knowledge of the content that the meeting was canceled is still denied and the speaker is not committed to either its truth or its falsity. Thus, the two contents of (4) do not exhibit the interaction observed between the projection of the prejacent and the generalization.

(4) Sam doesn't know that the meeting was canceled.

The observed interaction has strong implications for analyses of EASs according to which the prejacent is a lexically specified presupposition. First, when the prejacent is locally accommodated under negation, the speaker is committed to the falsity of the prejacent; this is in contrast to a locally accommodated factive presupposition, for which the speaker is not committed to its truth or falsity. Second, when the prejacent is locally accommodated under some operator, the generalization is not interpreted under that operator, in contrast to the attitude ascription with *know*, which is always interpreted in the scope of the operator. These observations are problematic for advocates of analyses according to which the prejacent is a lexically specified presupposition and projection is assumed to be governed by the standard mechanisms of presupposition projection (e.g., Heim 1983; van der Sandt 1992).

Karttunen et al.'s (2014) analysis of EASs made progress over previous analyses by assuming that evaluative adjectives are ambiguous between the two schematic lexical entries in (5): the prejacent is specified as presupposed in the lexical entry adj_1 in (5a) but not in that of adj_2 in (5b).

(5) Karttunen et al. (2014: 249): Presuppositions and assertions of EASs of the form 'NP was Adj to VP'

a. adj_1
Presupposed content: NP VPed
Asserted content: For NP to VP would be Adj

b. adj_2
Presupposed content: For NP to VP would be Adj & for NP not to VP would not be Adj
Asserted content: What NP did about VPing was Adj

These two lexical entries correctly predict an interpretation of the NEAS in (2a) in which the prejacent projects and an interpretation of the NEAS in (3b) according to which the prejacent does not project. Furthermore, the lexical entry in (5b) predicts that the NEAS in (3b) has an interpretation according to which it is presupposed that for the speaker to stumble through the junkyard would be stupid and for the speaker to not stumble through the junkyard would not be stupid; thus, in contrast to Oshima's (2009) and Barker's (2002) analyses, this analysis correctly predicts that (3b) conveys that the speaker stumbling through the junkyard would be stupid. Finally, Karttunen et al.'s analysis captures the interaction between the projection of the prejacent and the generalization: when the prejacent projects, as in the lexical entry in (5a), the generalization does not, by virtue of being coded in the asserted content; when the prejacent does not project, as in the lexical entry in (5b), the generalization does, by virtue of being part of the presupposed content.

There are, however, concerns with Karttunen et al.'s analysis. First, because it is not formalized, it is not clear what it predicts. Take (2a), which conveys that Feynman dancing on the table was not stupid. This meaning is assumed in Karttunen et al. (2014: 248) to follow from the presupposed prejacent (Feynman danced on the table) and the asserted content (it is not the case that for Feynman to dance on the table would be stupid). But whether this meaning follows depends on how the asserted meaning is formalized. A second concern is that Karttunen et al.'s analysis comes at a high cost: to derive the two interpretations, evaluative adjectives are systematically ambiguous. Finally, in addition to hardwiring the projection of the prejacent into one lexical entry but not the other, the analysis also hardwires the interaction between the projection of the prejacent and the generalization. Preferably, this interaction would fall out of the analysis.

In this paper, we develop an analysis that derives the two interpretations of EASs and the interaction between the projection of the prejacent and the generalization from a single lexical entry for evaluative adjectives. This is achieved, in a nutshell, by making the projection of the prejacent and the generalization sensitive to their discourse status rather than lexically specifying either as presupposed. In doing so, our analysis builds on previous analyses according to which the projection of content is derived from its discourse status, such as being backgrounded or not-at-issue; such analysis have been developed for the prejacent of manner adverbs (e.g., that Sam ran in *Sam ran quickly*), the pre-state content of *stop* (e.g., that Sam whistled in *Sam stopped whistling*) and the content of the complement of factive predicates like *know* in (4) (e.g., Abrusán 2011; 2013; 2016; Simons et al. 2010; 2017; Beaver et al. 2017; Stevens et al. 2017; Tonhauser et al. 2018; 2019).

Empirical motivation to develop analyses that derive the projection of content from its discourse status comes from observing that the content is less projective than one would expect if projection was derived from its conventional specification as a presupposition (e.g., Heim 1983; van der Sandt 1992) or as a conventional implicature (e.g., Potts 2005). The idea that some content is less projective than other content has been in the

literature for a long time: Karttunen (1971), for instance, suggested that the content of the complement of *regret* in (6a) is more projective than the content of the complement of *discover* in (6b); similarly, Schlenker (2010: 139) referred to *announce* as a "part-time trigger" because the content of its complement may, but often does not, project. For further remarks on projection variability see, e.g., Kadmon (2001), Simons (2001), Abusch (2002; 2010), Potts (2005), Beaver (2010), Abrusán (2011; 2016) and Tonhauser et al. (2013).

(6) Karttunen (1971: 63)
  a. John didn't regret that he had not told the truth.
  b. John didn't discover that he had not told the truth.

More recently, experimental research has provided empirical evidence for projection variability. Xue & Onea (2011), for instance, observed that the content of the complement of German *wissen* 'know' is less projective than the content of the complement of *erfahren* 'find out', both of which are less projective than the relevant projective contents of sentences with *auch* 'too' (that a parallel event is contextually salient) and *wieder* 'again' (that the relevant event has happened before). The most comprehensive investigation of projection variability to date is Tonhauser et al. (2018), which found extensive projection variability among 19 contents associated with American English expressions. They found, for instance, that the content of the complement of the factive predicate *be annoyed* is more projective than that of the factive predicate *discover*, which in turn is more projective than that of the factive predicate *reveal*. For further experimental evidence for projection variability see Smith & Hall (2011), de Marneffe et al. (2019) and Tonhauser et al. (2019) on American English and Tonhauser (in press) on Paraguayan Guaraní (Tupí-Guaraní).

The empirical observation that there is variability among projective contents in how projective they are gives rise to the question of what it means for one content to be more projective than another. According to Tonhauser et al. (2018: 498f.), there are at least two interpretations of gradience in projectivity:

> "On a first interpretation, a listener's (or reader's) judgment that a content is projective to a certain extent means that the listener takes the speaker (or writer) to be committed to the content to that extent. On this interpretation, projectivity being a gradient property is a consequence of speaker commitment being a gradient property. On a second interpretation, a listener's judgment that a content is projective to a certain extent reflects the probability with which they believe the speaker to be committed to the content. On this interpretation, speaker commitment may be a binary, categorical property and projection variability arises from the listener's uncertainty about … whether the speaker is committed."

Like Tonhauser et al. (2018), we remain agnostic about the underlying interpretation of projectivity as a gradient property, though our discussion will be in line with the first interpretation. To illustrate, consider the content of the complement of the examples in (6), that John had not told the truth. The observation that the content of the complement of *regret* in (6a) is more projective than that of *discover* in (6b) will be taken to mean that a speaker is more committed to John not having told the truth if the speaker utters (6a) than if the speaker utters (6b).

For the prejacent of EASs, there is empirical evidence that it is not highly projective. This empirical evidence provides motivation for developing an analysis according to which its projectivity is not derived from its status as a lexically specified presupposition but from its discourse status. First, Tonhauser et al.'s (2018) experimental investigation

found that the projectivity of the prejacent of EASs with *stupid* was significantly lower than that of non-restrictive relative clauses and appositive content, as well as of the content of the complement of the factive predicates *be annoyed* and *know*. A second piece of evidence that the prejacent is not highly projective comes from naturally occurring data. In a corpus-based web study, we collected 'certainty' ratings on a 7-point Likert scale from 226 native speakers of American English for the prejacents of 59 naturally occurring NEASs; the higher the certainty rating, the more projective the prejacent. Figure 1, which plots the mean certainty ratings of the 59 NEASs by evaluative adjective, shows that the prejacent exhibits projection variability: for some NEASs, the prejacent is highly projective, for some it is barely projective, and for others it is in-between. Across the 59 NEASs, the prejacent is not highly projective: the mean certainty rating was only 3.2.[5] Details on this corpus-based web study are given in Supplementary file 2.

In the next section, we build on Simons et al.'s (2017) question-based analysis of projective content to develop an analysis that derives the projection of the prejacent and the generalization from their status as not at-issue with respect to the question addressed by the EAS. Section 3 provides experimental evidence for two predictions of the analysis. After briefly considering the aforementioned interspeaker variation in Section 4, the paper concludes in Section 5.

2 A question-based analysis of the projective content of EASs

In this section, we develop a question-based projection analysis for EASs. We start in Section 2.1 by introducing the question-based analysis developed in Simons et al. (2017) for the content of the complement of factive predicates. In Section 2.2, we extend this analysis to EASs and show how the interaction between the projection of the prejacent and the generalization falls out of the analysis. In Section 2.3, we extend the analysis by incorporating a constraint on the questions addressed by utterances of EASs.

2.1 Utterance content projects if it is not at-issue with respect to the Discourse Question

The content of the complement of factive predicates like *know* or *discover* has long been observed to be able to project (e.g., Kiparsky & Kiparsky 1970; Karttunen 1971). Thus, speakers and writers can be committed to the content of the complement of *discover*, as

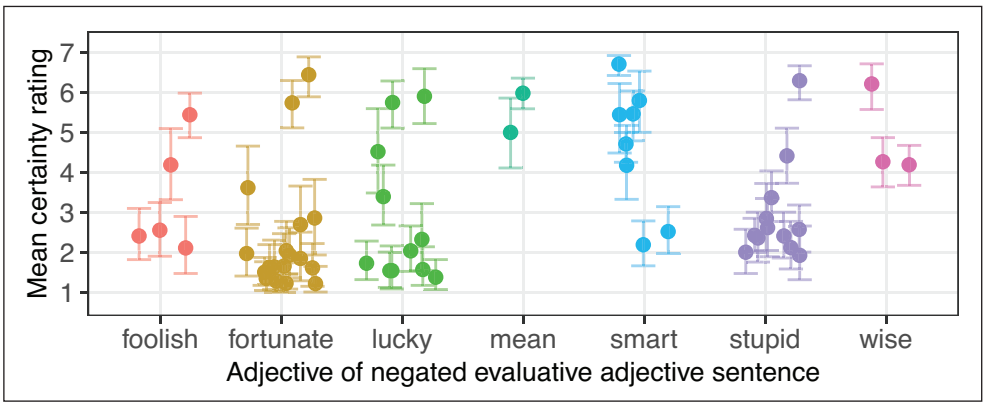

Figure 1: Mean certainty ratings of 59 naturally occurring NEASs by evaluative adjective. Error bars indicate bootstrapped 95% confidence intervals.

[5] In contrast to the prejacent of EASs like *Kim was smart to watch the movie*, the prejacent of sentences like *It was smart of Kim to watch the movie* appears to be more highly projective according to native speaker intuitions and, therefore, may be best analyzed as a lexically specified presupposition, as in Oshima (2009).

in (7a), but they don't have to be, as in (7b). See also the discussion around example (4) with *know* in Section 1.

(7) a. Caroline stepped forward after a moment. "There is, of course, also the issue of Alex and the other Enhanced to tackle. If anyone **discovers** that Alex is here, it will be a disaster difficult to avert." Alex frowned. L.C. Mawson, *Pandora: Freya Snow*, #13

b. [Mattress springs] also work well to deter rabbits & foxes from digging into the chook-pen (Hen-run). Dig a shallow trench the width of a single mattress, then place the springs flat into the trench. Drive your fence posts in the mid-line, so half the spring is outside & half inside the pen. I haven't tried this with wombats, though & if anyone **discovers** that the method is also wombat-proof, I'd really like to know. (Beaver 2010: 79)

Traditionally, projection of the content of the complement of factive predicates has been derived by specifying the content as presupposed in the lexical entry of the factive predicate (e.g., Heim 1983; van der Sandt 1992); as discussed in Section 1, non-projection is attributed to local accommodation under such analyses. In view of the observation that the projectivity of the content of the complement is weaker than that of other projective content, some analyses have abandoned the assumption that the content projects because it is lexically specified as presupposed: on some analyses, projection is derived from a lexically specified set of alternatives to the factive predicate in combination with pragmatic principles (e.g., Abusch 2002; 2010; Romoli 2015) and on others it is derived from the discourse status of the content of the complement (e.g., Abrusán 2011; 2016; Simons et al. 2010; 2017). Analyses that derive the projection of the content of the complement from its discourse status are also motivated by the observation that projection is sensitive to information structure. To illustrate, consider the examples in (8): Beaver (2010) hypothesized that the content of the complement of *discover* is less projective in (8a), which is realized with focus on *plagiarized*, than in (8b), which is realized with focus on *discover*. For experimental evidence consistent with this hypothesis see Cummins & Rohde (2015), Tonhauser (2016), Djärv & Bacovcin (2017) and Mahler (2019).

(8) Examples adapted from Beaver (2010: 93)

a. If the T.A. discovers that your work is $[\text{plagiarized}]_F$, I will be forced to notify the Dean.

b. If the T.A. $[\text{discovers}]_F$ that your work is plagiarized, I will be forced to notify the Dean.

To capture the information-structure sensitivity of the projection of the content of the complement, Simons et al. (2010, 2017) assumed that the content of the complement is a lexical entailment of sentences with factive predicates that projects if it is not at-issue with respect to the Discourse Question addressed by the utterance of the sentence.[6] The Discourse Question "provides the topic of a segment of discourse and imposes relevance

[6] Following Abrusán (2011), lexical entailments are entailments of sentences, not predicates. They differ from logical entailments, which can be derived independently from the meaning of the sentence. For instance, the disjunction of the content of the complement and some random proposition, like the proposition that it is raining, is a logical entailment of a sentence with a factive predicate, but not a lexical one. For the assumption that at least some projective contents, including presuppositions, are entailments that may project see, e.g., Gazdar (1979), Barker (2002), Schlenker (2010), Abrusán (2011; 2016) and Anand & Hacquard (2014).

constraints on conversational contributions" (Simons et al. 2017: 192). Utterance content is at-issue with respect to the Discourse Question of the utterance if the content addresses the Discourse Question, i.e., entails at least a partial answer to the question (Roberts 2012/1996). To illustrate, consider the content of the complement of *discover* in B's utterances in (9a) and (9b), that Harriet was at Princeton for a job interview. In these examples, the Discourse Questions that B's utterances address are made explicit by A's interrogative utterances. In (9a), the content of the complement of *discover* is at-issue because it addresses the Discourse Question: that Harriet was at Princeton for a job interview is an answer to A's interrogative utterance of where Harriet was yesterday. In (9b), on the other hand, the content of the complement is not at-issue because it is not an answer to A's interrogative utterance; rather, here the Discourse Question is addressed by the main clause content of B's utterance.

(9) Examples adapted from Simons (2007: 1035)

a. A: Where was Harriet yesterday?
B: Henry discovered that she was at Princeton for a job interview.

b. A: Why is Henry in such a bad mood?
B: He discovered that Harriet was at Princeton for a job interview.

Now consider the example in (10), in which *discover* is embedded in a polar question: consequently, the speaker is not committed to the content of the main clause; rather, the speaker is asking whether Henry discovered the content of the complement.

(10) Did Henry discover that Harriet was at Princeton for a job interview?

In contrast to the main clause content, the content of the complement of *discover* may project out of the polar question, as noted above. According to Simons et al. (2017), one of the conditions under which the content of the complement projects is when it is not at-issue with respect to the Discourse Question addressed by the utterance. Consider the examples in (11), where A's interrogative utterances again make explicit the Discourse Questions addressed by B's utterances. In (11a), the content of the complement of B's utterance does not address the Discourse Question, i.e., is not at-issue. Here, B is committed to the content of the complement: to make sense of B's utterance as a relevant response to A's question, there must be a connection between the possibility of Henry discovering something about Harriet and Henry's mood; an obvious connection is that B is committed to the content of the complement, that Harriet was at Princeton for a job interview. Thus, the content of the complement projects. In (11b), by contrast, the content of the complement addresses the Discourse Question, i.e., is at-issue. Here, B need not be committed to the content of the complement to make sense of how B's utterance addresses A's question and the content of the complement does not project.

(11) a. Context: Henry and Harriet are an academic couple that lives on the West Coast.
A: Why is Henry in such a bad mood?
B: Did he discover that Harriet was at Princeton for a job interview?

b. Context: Henry is a nosy colleague of Harriet and well-informed about her whereabouts.
A: Where was Harriet yesterday?
B: Did Henry discover that she was at Princeton for a job interview?

In sum, according to Simons et al. (2017), utterance content projects if it is not at-issue with respect to the Discourse Question addressed by the utterance. The question of why not-at-issue content projects has received several answers. According to Potts (2005), it projects because it is contributed to a separate dimension of meaning and, according to Simons et al. (2010), because it is not targeted by operators like negation and thereby projects over such operators. In this paper, we follow Abrusán (2011; 2016) in assuming that not-at-issue content is backgrounded and projects as a result of its discourse status.[7]

In naturally occurring discourse, the Discourse Question addressed by an utterance is more often implicit than explicit. When the Discourse Question is implicit, the utterance itself and the discourse context in which it is made provide cues to the Discourse Question, and thereby to the at-issueness and projectivity of utterance content. For instance, for utterances of sentences with factive predicates, the information structure of the utterance, in particular prosodically marked focus, has been shown to constrain the question addressed by the utterance and, hence, the at-issueness and projection of the content of the complement (e.g., Beaver 2001; Cummins & Rohde 2015; Tonhauser 2016; Simons et al. 2017; Djärv & Bacovcin 2017; Mahler 2019). As noted in Tonhauser et al. (2018: 501), when the Discourse Question is implicit, an interpreter's determination of the Discourse Question of an utterance "requires the integration of prosodic, lexical and structural cues from the uttered sentence as well as cues from high-level properties of discourse" (for discussion see, e.g., Tonhauser 2016; Beaver et al. 2017; Simons et al. 2017). We assume, with Tonhauser et al. (2018), that variable at-issueness may reflect uncertainty on part of the interpreter about the Discourse Question that the speaker intended to address with their utterance. In other words, the extent to which any given utterance content is at-issue reflects the extent to which the various sentence-level and high-level discourse cues jointly point to the speaker's utterance addressing a Discourse Question relative to which the content is at-issue. Given that the projectivity of utterance content is also variable, Tonhauser et al. (2018) modified Simons et al.'s (2017) proposal, according to which content projects if it is not at-issue, by arguing instead that content projects to the extent that it is not at-issue:

(12) **Gradient Projection Principle:** If content *C* is expressed by a constituent embedded under an entailment-canceling operator, then *C* projects to the extent that it is not at-issue. (Tonhauser et al. 2018: 499)

Tonhauser et al. (2018) provided experimental evidence for the Gradient Projection Principle on the basis of an experimental investigation of 19 projective contents associated with American English expressions. Specifically, Tonhauser et al. (2018) found a positive correlation between not-at-issueness and projection such that content that is more not-at-issue is also more projective. For instance, as shown in Figure 2 for 9 of these 19 contents

[7] We do not adopt Abrusán's analysis otherwise because it does not appear to make correct predictions for EASs. Abrusán (2011) proposed that lexical entailments that are about the running time of the main event are the default main point, i.e., what we have referred to as the at-issue content. This analysis, however, does not make correct predictions for EASs: given Abrusán's notion of aboutness, both the prejacent and the generalization of EASs are about the running time of the main event. For instance, Abrusán (2011: 508) took the entailment of (i), that John solved the exercise, to be "non-accidentally (i.e., necessarily) about the matrix event time". By the same argument, the prejacent of (ii) would be about the matrix event time, which means that neither the prejacent nor the generalization are predicted by Abrusán (2011) to be the default at-issue content of EASs.

(i) John managed (at time t1) to solve the exercise (at t1). (Abrusán 2011: 508)

(ii) John was smart (at time t1) to solve the exercise (at t1).

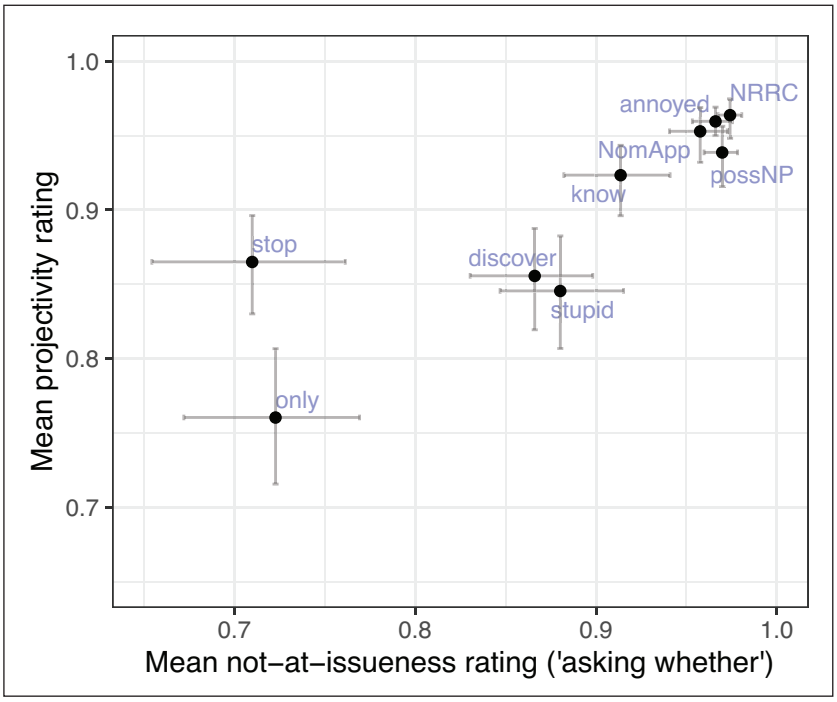

Figure 2: Tonhauser et al. (2018: 509): mean projectivity against mean not-at-issueness by target expression/projective content. Error bars indicate bootstrapped 95% confidence intervals.

(r = .85), the prejacent of EASs with *stupid* is not only less projective than the content of the complement of *be annoyed*, but also less not at-issue.

In the next section, we build on the question-based projection analysis developed in Simons et al. (2017) and Tonhauser et al. (2018) in developing a question-based projection analysis of EASs.

2.2 Lexical entailments of evaluative adjective sentences

We assume, with Barker (2002), that the prejacent and the generalization are lexical entailments of unembedded EASs (but not that the prejacent is lexically specified as presupposed).[8,9] We also assume that the generalization is not lexically specified as presupposed. Translations of the prejacent and the generalization of an unembedded EAS of the form 'NP be.TENSE Adj to VP' are given in (13). In the translation of the prejacent in (13a), the translation of the NP is the constant *np* (of type *e*, for entities) and of the VP is *VP′* (of type ⟨*e*, ⟨*ev*, *t*⟩⟩, where *ev* is the type of eventualities and *t* is the type of truth values). The run time of the event *e* of *NP VP*ing, given as $\tau(e)$ (type *i*), is temporally located at the reference time *rt*, whose temporal location is constrained by the tense of the EAS. In the translation of the generalization in (13b), the constant *deg′* combines with the transla-

[8] As lexical entailments, the prejacent and the generalization are entailments of evaluative adjective sentences, not of the evaluative adjectives. This is not to say that the evaluative adjectives do not give rise to entailments: as discussed in detail in Barker (2002: §4.2), they give rise to subject-related entailments, namely that the denotation of the subject noun phrase is capable of volition and, regarding the situation described by the *to*-infinitive, has the power to bring it about and intends for it to come about. We assume, with Barker (2002), that these subject-related entailments are presuppositions that are triggered by the evaluative adjectives.

[9] Some EASs with *will* may not entail the prejacent: for example, some speakers do not judge (i) to entail that Johnson will take what he can get. We thank David Beaver (p.c.) for this point, which we sidestep here.

(i) With more teams denying interest in Johnson, he will be smart to take what he can get. (*http://www.sportsworldreport.com/articles/28999/20140407/chris-johnson-rumors-ny-jets-release-mike-goodson-free-agent-signs-dallas-cowboys-demarco-murray-rams-falcons.htm*)

tion of the adjective (of type $\langle i, \langle\langle ev, t\rangle, d\rangle\rangle$) and maps it to the contextual standard of the adjective (a degree, type d).

(13) Unembedded EAS of the form 'NP be.TENSE Adj to VP'

a. Prejacent: $\exists e(VP'(np)(e) \wedge at(rt, \tau(e)))$
"There is an event of *NP VP*ing and the run time of that event is located at the reference time."

b. Generalization: $adj'(rt, VP'(np), deg'(adj'))$
"At the reference time, the degree to which *NP VP*ing is *Adj* is higher than the contextual standard for *Adj*."

By (13), the past tense EAS in (1) *Feynman was stupid to dance on the table* entails that Feynman danced on the table (the prejacent) and that, at the past reference time, the degree to which Feynman dancing on the table is stupid was higher than the contextual standard for *stupid* (the generalization). The events involved in the generalization need not be actual: for instance, it does not follow from the generalization of (1) that an event of Feynman dancing on the table took place; this follows from the prejacent.

The proposal that the prejacent and the generalization are lexical entailments predicts that unembedded EASs are judged to be unacceptable if either of them is false. The examples in (14) show that that prediction is borne out. Consider (14a), whose prejacent is false and whose generalization is true (under the assumption that the degree to which anybody, including Kim, being born into poverty is unfortunate is higher than the contextual standard of *unfortunate*): this example is correctly predicted to be unacceptable because the prejacent is false. In (14b), on the other hand, the prejacent is true and the generalization is false (under the assumption that the degree to which anybody, including Sandy, being born into poverty is lucky is lower than the contextual standard of *lucky*). This sentence is correctly predicted to be unacceptable because the generalization is false.

(14) a. What is true: Kim was born to rich parents
#Kim was unfortunate to be born into poverty.

b. What is true: Sandy was born into poverty
#Sandy was lucky to be born into poverty.

Having motivated that both the prejacent and the generalization are lexical entailments of EASs, we now turn to the discourse status of these contents. We argue that the discourse status of the prejacent and the generalization is not conventionally specified by the EAS: either can be entailed by the common ground at the time at which the EAS is uttered or new information at that time; in other words, neither content is associated with Tonhauser et al.'s (2013) Strong Contextual Felicity constraint. To illustrate, consider the naturally occurring EASs in (15). The EAS in (15a) is acceptable in a context in which the prejacent follows from the common ground, i.e., the interlocutors already know that Trump gave Wolff unlimited access, and in which the writer conveys the generalization as new information, i.e., that they took the degree to which Trump giving Wolff unlimited access is stupid to be higher than the contextual standard of *stupid*. The EAS in (15b) illustrates the reverse: (15b) is acceptable in a context in which the generalization follows from the common ground, whereas the prejacent, that MacLean grabbed Kouassi in his private parts is new information. And, finally, there are also EASs in which both the prejacent and the generalization may be new information: (15c) is acceptable in a context in which neither the generalization nor the prejacent follow from the common ground.

When a speaker utters this EAS, they are thereby committing to the truth of the prejacent, that they bought an Xbox, as well as to the truth of the generalization, that the degree to which them buying an Xbox 360 elite is stupid was higher than the contextual standard of *stupid*.

(15) a. Trump was stupid to give Wolff unlimited access.[10]
b. Steven MacLean was stupid to grab Eboue Kouassi in his private parts.[11]
c. I was stupid to buy the Xbox 360 elite.[12]

Not only can both the prejacent and the generalization of an EAS be new information, they can also both be at-issue with respect to the Discourse Question addressed by an utterance of the EAS. Of course, given the characterization of at-issue content as addressing the Discourse Question, only one of them is at-issue in any given utterance of an EAS (see also, e.g., Abbott 2000). We further assume that the prejacent and the generalization are the only contenders for at-issue content of an EAS, which means that exactly one of them is at-issue in any given utterance of an EAS; the other one is not-at-issue.[13] To illustrate, consider the examples in (16), in which utterances of the EAS *Sam/she was smart to buy a ticket/one the day they went on sale* address four distinct Discourse Questions. The Discourse Questions in (16a) and (16b) are about the prejacent: both the question of whether Sam got a ticket in (16a) and who got a ticket in (16b) are answered by the prejacent of B's utterances, that Sam bought a ticket.[14] Thus, in (16a) and (16b), the prejacent is at-issue and the generalization, that the degree to which Sam buying a ticket the day the tickets went on sale is smart was higher than the contextual standard of *smart* is not at-issue. The Discourse Questions in (16c) and (16d), on the other hand, are not about the prejacent: neither the question about B's assessment of Sam buying a ticket when they went on sale (16c) nor the question of whether Sam was smart to buy a ticket the day tickets went on sale in (16d) are answered by the prejacent of B's utterances; rather they are answered by the generalization. Thus, in (16c) and (16d), the generalization is at-issue and the prejacent is not at-issue.

(16) a. A: The show was sold out. Did Sam get a ticket?
B: She was smart to buy one the day they went on sale.
b. A: There were so few tickets for the show! Who got a ticket?
B: Sam was smart to buy one the day they went on sale.
c. A: How do you assess Sam's buying of a ticket the day tickets went on sale?
B: Sam was smart to buy one the day they went on sale.

[10] https://www.eastbaytimes.com/2018/01/07/letter-trump-was-stupid-to-give-wolff-unlimited-access/.
[11] https://www.heraldscotland.com/sport/17189197.brendan-rodgers-steven-maclean-was-stupid-to-grab-eboue-kouassi-in-his-private-parts/.
[12] www.gamespot.com/forums/xbox-association-1000003/i-was-stupid-to-buy-the-xbox-360-elite-26937557.
[13] The subject-related entailments of EASs mentioned in footnote 8 are, by virtue of being lexically specified presuppositions, conventionally specified as not at-issue. We leave for future research the question of whether there are discourse questions that make the conjunction of the prejacent and the generalization the at-issue content of the utterance.
[14] Some native speakers of American English prefer to produce B's utterances in (16a) and (16b) with *enough*: *She was smart enough to buy one the day tickets went on sale*. We hypothesize that such speakers disprefer producing EASs in which the prejacent is at-issue. Given our hypothesis that the prejacent projects to the extent that it is not at-issue, we would expect such speakers to also disprefer producing NEASs like (3) in which the prejacent does not project. Crucially, as discussed in Section 1, even speakers who prefer to produce (16a) and (16b) with *enough* can retrieve the intended interpretations of the variants without *enough*.

d. A: Was Sam smart to buy a ticket for the show the day tickets went on sale? The price would have gone down after a few days!
B: Sam was smart to buy one the day they went on sale. The show sold out the day tickets went on sale.

We are now ready to return to our question-based projection analysis of EASs, according to which the more the prejacent and the generalization are not at-issue with respect to the Discourse Question, the more projective they are. To illustrate this analysis, consider first the minimal pair in (17a) and (17b), where B utters the NEAS *Sam/she wasn't smart to buy a ticket the day tickets went on sale*. The Discourse Questions relative to which B's utterances are interpreted are given by A's interrogative utterances; crucially, these Discourse Questions differ with respect to whether the prejacent or the generalization are at-issue. In (17a), A's interrogative utterance is about the prejacent, which means that the prejacent is at-issue and the generalization is not. Thus, our analysis predicts that the prejacent of B's utterance does not project, but that the generalization does. Accordingly, B's utterance is correctly predicted to mean that Sam didn't buy a ticket the day tickets went on sale (the prejacent does not project) and that the degree to which Sam buying a ticket the day tickets went on sale is smart was higher than the contextual standard of *smart* (the generalization projects). What follows is that it would have been smart for Sam to buy a ticket on the day tickets went on sale. In (17b), on the other hand, A's interrogative utterance is about the generalization, which means the generalization is at-issue and the prejacent is not. Thus, our analysis predicts that the prejacent projects and that the generalization does not. Accordingly, B's utterance is interpreted to mean that Sam bought a ticket the day tickets went on sale (the prejacent projects) and that the degree to which Sam buying a ticket the day tickets went on sale is smart was not higher than the contextual standard of *smart* (the generalization does not project). What follows is that Sam having bought a ticket on the day tickets went on sale wasn't smart.

(17) a. A: The show was sold out. Did Sam get a ticket?
B: She wasn't smart to buy a ticket the day they went on sale. (So, she didn't go to the show.)

b. A: You keep criticizing Sam for doing some not-so-smart things. I'm not sure I agree. Can you give me an example (of something not-so-smart that Sam did)?
B: Sure! Remember the rock show we talked about? Sam wasn't smart to buy a ticket the day they went on sale. The price went down a couple of days later.

In the examples we have entertained so far, the Discourse Question was realized by an interrogative utterance, which made explicit whether the prejacent or the generalization is at-issue. However, as noted above, the Discourse Question is often implicit in naturally occurring discourse and, when it is, interpreters may have uncertainty about the Discourse Question that the speaker intended to address with their utterance. In such cases, the prejacent and the generalization are at-issue to the extent to which the various sentence-level and high-level discourse cues jointly point to the speaker's utterance addressing a Discourse Question relative to which either content is at-issue. Our proposal predicts that both the prejacent and the generalization may project and that the more they are not at-issue, the more projective they are. Given that either the prejacent or the generalization must be at-issue in any given utterance of an EAS, the interaction in their projection falls

out of the analysis: the more the generalization is at-issue, the more projective the prejacent is, and the more the prejacent is at-issue, the more projective the generalization is.

Previous work has identified focus marking as a constraint on the question that an utterance can be taken to address (see references above); in this paper, we do not entertain focus marking because the prosody of EASs hasn't been investigated yet. Instead, the next section identifies a novel constraint on the Discourse Question of utterances, including utterances of EASs.

2.3 At-issue content is non-redundant

We propose that one of the constraints on the Discourse Question comes from a felicity requirement that is found in different guises in the literature: an utterance of an indicative sentence is felicitous only if the sentence is informative in the context in which it is uttered. Stalnaker (1999: 88), for instance, maintained that "[a] proposition asserted is always true in some but not all of the possible worlds in the context set". And Groenendijk (1999: 144) formulated this as the requirement that indicative sentences be non-redundant. We assume here that a speaker's utterance of a sentence is felicitous only if the at-issue content of the uttered sentence does not follow from the common ground of the interlocutors prior to utterance of the sentence.[15] Since content is at-issue with respect to the Discourse Question, this means that the Discourse Question cannot be about content that is entailed by the common ground. What follows is that the more an utterance content is taken to follow from the common ground, the less likely it is that the Discourse Question of the utterance is about that content:

(18) **Non-redundancy principle for at-issue content**
The more the interpreter takes the truth of content *c* to follow from the common ground before observing an utterance of an expression that conveys *c*, the less likely it is that the interpreter takes the speaker to have intended the Discourse Question of the utterance to be about *c*, i.e., the less *c* is at-issue.[16]

The principle in (18) has the following consequences for the interpretation of EASs. The more the interpreter takes the truth of the prejacent to follow from the common ground, the more she will assume that the Discourse Question does not concern the prejacent, i.e., the more not at-issue the prejacent is and the more projective it is. And likewise for the generalization. To illustrate, consider the naturally occurring examples in (19) and (3b) with the evaluative adjective *stupid*: the prejacent of (19) is highly projective according to the ratings obtained in the corpus study, whereas that of (3b) is not, as discussed in Section 1.

(19) God offers Hope to Hispanics! In His pages are solutions to every immigration problem. God loves citizens and immigrants equally. His solutions are for all of us. They are practical. They work. He is not stupid to think so.

[15] Whereas at-issue content may not be entailed by the common ground, not-at-issue content varies in its discourse status: some not-at-issue content, like factive presuppositions, may be new information or already entailed by the common ground, anaphoric presuppositions must be entailed by the common ground and, finally, conventional implicatures have been argued in Potts (2005) to be required to be new. We follow works like Potts (2005) and Murray (2014) in assuming that conventional implicatures are always not at-issue. Thus, while both the at-issue content and conventional implicatures may be new information, they differ in at-issueness.

[16] This principle is meant as a conditional, not a bi-conditional: it is not the case that the less likely it is that the Discourse Question is about a particular content, the more the truth of that content is taken to follow from the common ground.

(3b) Now I knew someone was in the junkyard and the cold wind was carrying the cries. I wasn't stupid to go stumbling through the junkyard in the dark and get hurt.

The prejacent of the NEAS in (19) is that God thinks that his solutions are practical and work. The strength of the inference from the common ground to the truth of this prejacent is presumably high. Thus, the principle in (18) predicts that the Discourse Question of (19) is not likely to be about the prejacent, which is therefore highly not-at-issue and highly projective, as observed. Now consider the NEAS in (3b), whose generalization is that the degree to which the speaker stumbling through the junkyard in the dark and getting hurt is stupid was higher than the contextual standard of *stupid*. The strength of the inference from the common ground to the truth of the generalization is quite high. Thus, the principle in (18) predicts that the Discourse Question of (3b) is not likely to be about the generalization; rather, it is expected to be about the prejacent, which, as predicted, is observed to not project.

2.4 Interim summary and predictions

This section developed a question-based projection analysis for EASs that builds on previous analyses in assuming that utterance content projects to the extent that it is not at-issue (Beaver et al. 2017; Simons et al. 2017; Tonhauser et al. 2018). To apply a question-based projection analysis to EASs, this section showed that the prejacent and the generalization are lexical entailments of EASs and that both can but need not be at-issue with respect to the Discourse Question. In contrast to analyses according to which the prejacent is lexically specified as a presupposition (e.g., Barker 2002; Oshima 2009), neither the prejacent nor the generalization are lexically specified as such on our analysis. Instead, the projectivity of the prejacent and the generalization are derived from their discourse status: the more the prejacent and the generalization are not at-issue, the more projective they are. Because only one utterance content is at-issue, the analysis also predicts the remarkable interaction between the prejacent and the generalization that has not been observed before in the literature on projective content: when the prejacent projects, the generalization does not, and when the prejacent does not project, the generalization does. The analysis also incorporates a novel constraint on Discourse Questions, the 'Non-redundancy principle for at-issue content'.

As discussed in Section 1, the analysis developed in Karttunen et al. (2014) captures the projection of the prejacent and the generalization, as well as the interaction in their projection, by assuming two lexical entries for each evaluative adjective: one according to which the prejacent projects and the generalization does not, and one according to which the generalization projects and the prejacent does not. There are three reasons why we think our analysis compares favorably to that presented in Karttunen et al. (2014). First, our analysis is more parsimonious: whereas both analyses build on prior analyses of projective content (i.e., do not introduce additional machinery), evaluative adjectives are required to be ambiguous under Karttunen et al.'s (2014) analysis but not ours. Second, the interaction between the prejacent and the generalization is stipulated in Karttunen et al. (2014), but falls out of the independently-motivated assumption of our analysis that only one utterance content is at-issue. Finally, as shown by the discussion in Section 2.1, the fact that projective content can but need not project is generally not attributed to lexical ambiguity, neither under analyses that take projective content to be conventionally specified nor on analyses that derive projection differently. Thus, when evaluated against the broader landscape of projection analyses, Karttunen et al.'s (2014) analysis is more idiosyncratic than ours.

The next section provides experimental evidence in support of our analysis. Specifically, the two experiments manipulated the extent to which the generalization was at-issue and investigated the following two predictions of the analysis:

(20) The more the generalization is at-issue,
  a. … the more projective the prejacent is. (Experiment 1)
  b. … the less at-issue the prejacent is. (Experiment 2)

To investigate the predictions in (20), we measured the projectivity of the prejacent (Exp. 1) and the at-issueness of the prejacent (Exp. 2).

3 Empirical evidence for the question-based projection analysis

According to the question-based projection analysis developed in Section 2, the prejacent and the generalization of EASs are predicted to project to the extent that they are not at-issue. Preliminary evidence for this prediction comes from Tonhauser et al.'s (2018) finding that at-issueness predicts projection for 19 projective contents. One of the contents they investigated was the prejacent of EASs with *stupid* embedded under the polar question operator. The four EASs they investigated are given in (21):

(21) a. Was Raul stupid to cheat on his wife?
  b. Were John's kids stupid to be in the garage?
  c. Is Mary's daughter stupid to be biting her nails?
  d. Is Richie stupid to be a stuntman? (Tonhauser et al. 2018, Appendix A)

Tonhauser et al. (2018) collected projection and at-issueness ratings on a scale: Figure 3 shows participants' projectivity ratings for the prejacents of the four items in (21) against their not-at-issueness ratings.[17] There is a clear relationship between at-issueness and projectivity: the more a participant rated a prejacent as not-at-issue, the more projective they rated it ($r = .91$; when not collapsing over the four questions $r = .57$).

Tonhauser et al.'s (2018) finding provides preliminary evidence that the projectivity of the prejacent of *stupid* is sensitive to its at-issueness. However, in Tonhauser et al.'s (2018) experiment, at-issueness was only measured, not manipulated, and only the prejacents of EASs with *stupid* were investigated. In the experiments reported on in this section, we

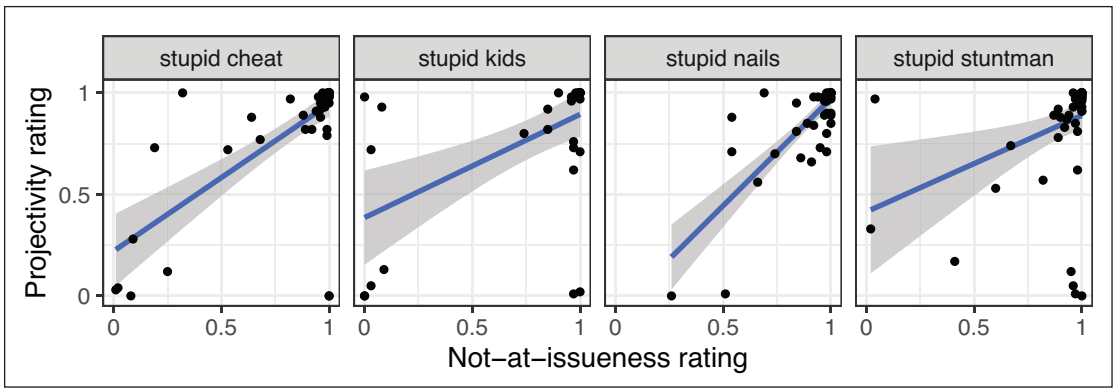

Figure 3: Projectivity ratings against not-at-issueness ratings for the prejacents of four EASs in Tonhauser et al. (2018). Each dot represents one participant's ratings. Linear smoothers with 95% confidence intervals overlaid.

[17] We re-plotted their data, obtained at https://github.com/judith-tonhauser/how-projective.

provide more direct evidence for the question-based projection analysis by manipulating at-issueness and including a wider variety of evaluative adjectives and items.

Further preliminary evidence for our analysis comes from an experiment reported on in Karttunen et al. (2014), which investigated the projection of the prejacents of NEASs with 19 evaluative adjectives (*arrogant, brave, careless, cruel, evil, foolish, fortunate, heroic, humble, lucky, mean, nice, polite, rude, sensible, smart, stupid, sweet, wise*). The materials included one triple for each evaluative adjective, like the triple for *smart* in (22). The experiment manipulated whether there was a predisposition to a content that is related to our generalization. Specifically, each triple included a NEAS referred to by Karttunen et al. (2014) as 'consonant', which means that "there is a predisposition to assume or grant that for NP to VP would be Adj" (p.237). For instance, the NEAS in (22a) is consonant because for Paul to take the best piece would be smart; this is comparable to what we have characterized as the generalization following from the common ground. Each triple also included a NEAS that Karttunen et al. (2014) referred to as 'dissonant', which means that "there is a predisposition to assume or grant that for NP to VP would not be Adj" (*ibid*). The NEAS in (22c) is dissonant because for the man to take the worst piece would not be smart; this is comparable to the falsity of the generalization following from the common ground. Finally, the third NEAS in each triple was considered 'neutral', i.e., neither consonant or dissonant, like (22b).

(22) Sample stimuli from Karttunen et al. (2014: 241)
  a. Paul wasn't smart to take the best piece. [consonant]
  b. Sally wasn't smart to take the middle piece. [neutral]
  c. The man wasn't smart to take the worst piece. [dissonant]

Karttunen and his colleagues found that the prejacent of dissonant NEASs was more projective than the prejacent of neutral ones, and that the prejacent of neutral ones was more projective than that of consonant ones. These findings support the prediction of our analysis, that the discourse status of the generalization influences the projectivity of the prejacent. There are, however, some concerns with their experiment. First, the experiment included only one triple for each evaluative adjective, and so the findings potentially have limited generalizability. Second, the triples were not normed to establish that native speakers of American English share Karttunen and his colleagues' assumptions about consonance and dissonance. Relatedly, the stimuli were presented to the participants without a context even though context can influence whether the generalization follows from the common ground. For instance, if Sally is on a diet and *middle piece* in (22b) is understood as *middle piece of the cake*, then (22b) is not neutral, but dissonant: for Sally to take the middle piece (or any piece) is not smart given that she is on a diet. To address these concerns, our Experiment 1 included a greater variety of items, which were normed with native speakers of American English, and presented to the participants with a context.[18]

3.1 Experiment 1: Projectivity of the prejacent

The experiment we report on in this section investigated the prediction in (20a), that the more at-issue the generalization is, the more projective the prejacent is. The at-issueness of the generalization was manipulated by having the truth of the generalization be more

[18] The experiments reported on in this paper and the Supplementary files were conducted with approval from the Human Research Protection Program at The Ohio State University. The data and the R code for generating the figures and analyses of the experiments reported on in this paper and the Supplementary files are available at https://github.com/judith-tonhauser/evaluative-adjectives.

likely to follow from the common ground in one condition than the other. We collected gradient projection ratings for prejacents in the two conditions.

3.1.1 Methods

3.1.1.1 Participants

152 participants with US IP addresses and at least 97% of HITs approved were recruited on Amazon's Mechanical Turk platform (ages: 18–69, mean age: 33). They were paid 45 cents.

3.1.1.2 Materials

Stimuli consisted of two-sentence discourses. In the target stimuli, the first sentence was a context sentence and the second sentence was a NEAS with one of the following 10 evaluative adjectives: *stupid, smart, wise, fortunate, lucky, brave, polite, mean, foolish* and *rude*. For each adjective, there were 6 pairs of target stimuli, for a total of 60 pairs of target stimuli. We used a 2 x 2 within-participant design: one factor was whether the truth of the generalization was taken to follow from the common ground (levels: less likely vs. more likely); the other factor was whether the content of the NEAS or of the context sentence determined whether it was likely that the truth of the generalization follows from the common ground (levels: content vs. context). The sample stimuli in (23) and (24) illustrate these two factors. For a pair of stimuli in the Content condition, shown in (23), the context sentences were identical and the two NEASs differed in how likely the truth of the generalization is taken to follow from the common ground. In (23a), the generalization is that the degree to which Sally losing her wallet is fortunate was higher than the contextual standard of *fortunate* (to be clear: the generalization is established on the EAS, not the NEAS). Here, it is less likely that the truth of the generalization is taken to follow from the common ground. In (23b), the generalization is that the degree to which Sue speaking French is fortunate was higher than the contextual standard of *fortunate*. Here, it is more likely that the truth of the generalization is taken to follow from the common ground. For pairs of stimuli in the Context condition, shown in (24), the NEASs were identical and the context sentences differed in how likely the truth of the generalization is taken to follow from the common ground. In (24a), the generalization is that the degree to which Jane prank-calling the police is smart was higher than the contextual standard of *smart*; here, it is less likely that the truth of the generalization is taken to follow from the common ground. In (24b), the generalization is that the degree to which Jane calling the police is smart was higher than the contextual standard of *smart*; here, it is more likely that the truth of the generalization is taken to follow from the common ground. In both the Content and Context conditions, we expected the projectivity of the prejacent to be higher when the truth of the generalization was less likely to follow from the common ground, i.e., when the generalization is more at-issue.

(23) Sample stimuli in the Content condition

a. Sue was traveling in France. She wasn't fortunate to lose her wallet. [generalization less likely]

b. Sue was traveling in France. She wasn't fortunate to speak some French. [generalization more likely]

(24) Sample stimuli in the Context condition

a. Jane was prank-calling people. She wasn't smart to call the police. [generalization less likely]

b. Jane saw a man with a gun. She wasn't smart to call the police. [generalization more likely]

To assess the projectivity of the prejacent, participants were asked a polar question version of the prejacent on each trial, e.g., 'Did Sue lose her wallet?' in (23a), 'Did Sue speak some French?' in (23b), and 'Did Jane call the police?' in (24a) and (24b).[19]

Of the 6 pairs of target stimuli per evaluative adjective, 3 pairs were in the Content condition and 3 in the Context condition. We selected these 60 pairs of target stimuli from a total of 120 pairs of potential target stimuli for which we collected ratings in a norming study from native speakers of American English about whether the truth of the generalization was taken to follow from the common ground. For each evaluative adjective, we chose the 3 pairs of target stimuli in the Content and Context conditions such that the generalization was most likely to be taken to follow from the common ground in one member of the pair and least likely to follow in the other member of the pair. The norming study is described in detail in Supplementary file 3; the full set of stimuli is provided in Supplementary file 4.

The 120 target stimuli were distributed across 12 lists of 10 target stimuli so that each evaluative adjective occurred once per list. Each list included 5 target stimuli in which the generalization is more likely to be taken to follow from the common ground and 5 in which the generalization is less likely to be taken to follow from the common ground. Each list included 5 stimuli from the Content condition and 5 from the Context condition. To assess whether participants were attending to the task, the same 6 control stimuli were added to each list, for a total of 16 stimuli per list (see Supplementary file 4 for the control stimuli).

3.1.1.3 Procedure

Participants were told that they would read short descriptions of scenarios and were asked a question about each scenario. They were randomly assigned to a list and presented with the 16 stimuli, one after the other, in random order. As shown in Figure 4, they gave their response to the polar question on a 7-point Likert scale labeled at four points: No (coded as 1), Possibly no, Possibly yes, Yes (coded as 7). We assume that the higher a participant's response to a polar question, the more projective the prejacent is, i.e., the more the speaker/writer is committed to the truth of the prejacent.

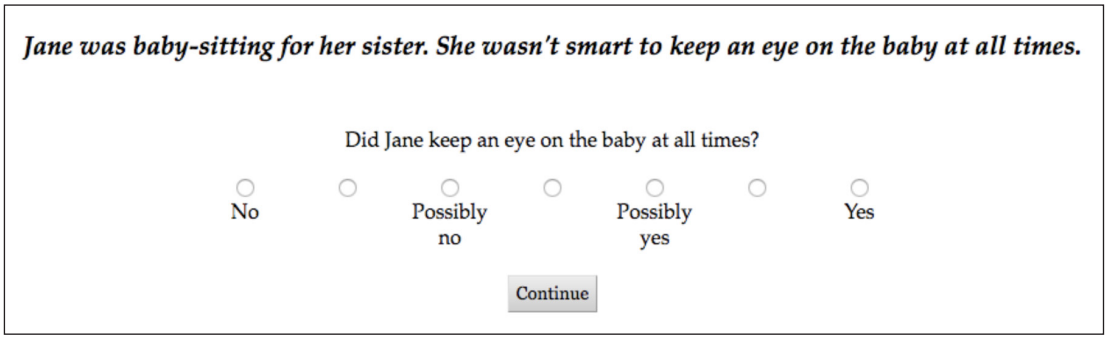

Figure 4: A sample trial in Experiment 1.

[19] This diagnostic for projection differs from the 'certain that' diagnostic used in the web-based corpus study or in Tonhauser et al. (2018). Experiment 1 relied on a different diagnostic because it was run before the 'certain that' diagnostic was developed. Under the assumption that participants only take the prejacent to be true if it follows from the two-sentence discourse, i.e., if the author of the two-sentence discourse is committed to the prejacent, we take the two diagnostics to be comparable in diagnosing projection. See Tonhauser et al. (2018) for a discussion of other diagnostics for projection.

After rating the 16 stimuli, participants completed a brief questionnaire about their age, their native language(s) and, if English is a native language, whether it is American English, as opposed to, e.g., Indian or Australian English. Participants were told that they would be paid no matter how they responded to these questions, in order to encourage them to answer truthfully.

3.1.1.4 Data exclusion

The data from 5 participants who did not self-identify as native speakers of American English were excluded. 13 participants were excluded based on their responses to the control stimuli (see Supplementary file 4 for details), leaving data from 134 participants (ages: 18–69; mean age: 33).

3.1.2 Results

Each of the 120 target stimuli received between 9 and 14 ratings (mean: 11.2). Figure 5 shows the mean projectivity ratings for prejacents of NEASs in the Content condition (left panel) and the Context condition (right panel) by how likely the truth of the generalization was taken to follow from the common ground. As expected, prejacents were more projective when the truth of the generalization was less likely to follow from the common ground than when it was more likely to follow from the common ground: in the Content condition, the mean projectivity ratings were 5.2 and 2.6, respectively; in the Context condition, the mean projectivity ratings were 4.7 and 2.9, respectively. As shown by the overlaid adjective means in Figure 5, the effect of the discourse status of the generalization was observed for all of the adjectives, albeit to varying degrees. These

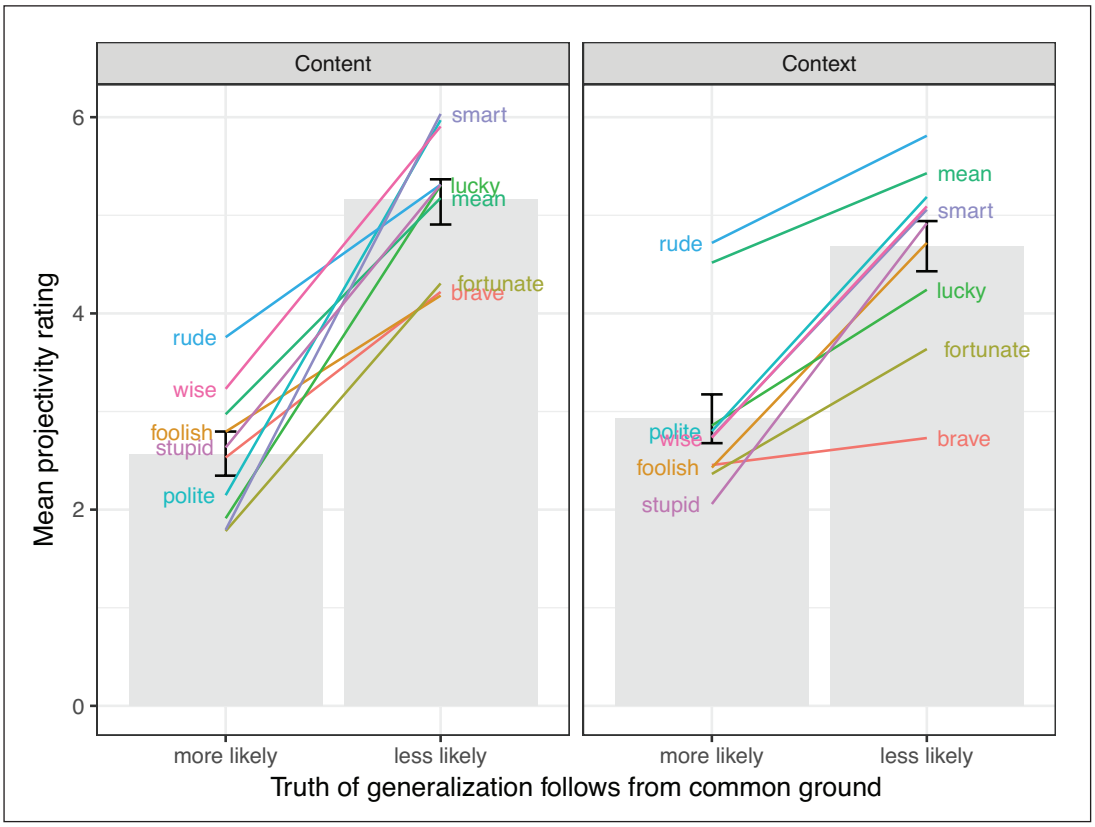

Figure 5: Mean projectivity ratings, with bootstrapped 95% confidence intervals, for prejacents of NEASs in the Content condition (left panel) and the Context condition (right panel) by how likely the truth of the generalization is taken to follow from the common ground. Adjective means in the two conditions overlaid.

findings provide empirical support for the prediction that at-issueness influences projection with EASs: when the truth of the generalization is less likely to follow from the common ground, the generalization is more at-issue, in which case the prejacent is more not at-issue and hence more projective.

We fitted ordinal mixed-effects regression models to the target data in the Content and Context conditions (668 and 672 data points, respectively), using the `clmm` function of the `ordinal` package (Christensen 2013) in *R* (R Core Team 2016; version 3.2.0). The models predicted projectivity ratings on the 7-point Likert scale from the fixed effect of the discourse status of the generalization (with 'more likely' as the reference level). The models included the maximal random effects structure justified by the data and the theoretical assumptions: random by-participant intercepts (capturing differences in projectivity between participants) and random by-item intercepts (capturing differences in projectivity between context/adjective/*to*-infinitive combinations) as well as random slopes for the discourse status of the generalization by participant (capturing that the effect of the discourse status may vary across participants). We obtained *p*-values by comparing models via likelihood ratio tests.

There was a significant main effect of the discourse status of the generalization such that the prejacent of items in which the truth of the generalization is less likely to follow from the common ground received higher projectivity ratings in the Content ($\beta = 3.29$, $SE = 0.34$, $z = 9.6$, $LR(1) = 73$, $p < .001$) and Context ($\beta = 2.1$, $SE = 0.37$, $z = 5.73$, $LR(1) = 29.68$, $p < .001$) conditions. These findings suggest that the discourse status of the generalization influences the projectivity of the prejacent, as predicted by the analysis developed in Section 2: prejacents of NEASs in which the truth of the generalization is less likely to follow from the common ground, i.e., is more at-issue, are more projective than prejacents of NEASs in which the truth of the generalization is more likely to follow from the common ground, i.e., is more not at-issue. Our findings also suggest that readers attend both to information from the EAS and to information from the context in determining the extent to which the truth of the generalization follows from the common ground and, therefore, the projectivity of the prejacent.

3.2 Experiment 2: At-issueness of the prejacent

This experiment tested the prediction in (20b), that the more the generalization is at-issue, the less the prejacent is at-issue. The at-issueness diagnostic used relies on the assumption that at-issue and not-at-issue content differ in the extent to which it is up for debate and can be directly assented/dissented with. For previous uses of diagnostics that rely on this assumption see, e.g., Amaral et al. (2007), Xue & Onea (2011), Tonhauser (2012), Murray (2014), AnderBois et al. (2015), Destruel et al. (2015), Syrett & Koev (2015) and Tonhauser et al. (2018). The diagnostic we used in Experiment 2 is the same as in Tonhauser et al.'s (2018) Exps. 2. The 3-turn dialogue in (25) illustrates how the diagnostic was set up on the basis of the appositive content associated with nominal appositives. The speaker of the first turn, Debby, utters an indicative sentence with the target expression, here a nominal appositive, and thereby commits herself to various utterance contents, including the appositive content that Martha's new car is a BMW. The speaker of the second turn, Harry, utters the question *Are you sure?*, thereby challenging some content of Debby's utterance. In the third turn, Debby utters an indicative sentence in which the content to be diagnosed for at-issueness, here the appositive content of her first utterance, realizes the content of the complement of *sure*, thereby identifying it as the content that Debby took the Harry to be challenging.

(25) At-issueness diagnostic from Tonhauser et al. (2018) Exps. 2
Debby: Martha's new car, a BMW, was expensive.
Harry: Are you sure?
Debby: Yes, I am sure that Martha's new car is a BMW.

To assess whether the relevant content is up for debate, i.e., at-issue, participants were asked whether Debby's utterance answered Harry's question, with 'yes' and 'no' as response options. A 'yes' response was taken to indicate that the relevant content was at-issue: the content was targeted by Harry's question and, therefore, Debby answered Harry's question. A 'no' response, in turn, was taken to indicate that the relevant content was not at-issue: the content was not targeted by Harry's question and, therefore, Debby did not answer Harry's question. We assume that the more 'yes' responses a content receives, the more at-issue it is.

3.2.1 Methods

3.2.1.1 Participants

75 participants with US IP addresses and at least 97% of HITs approved were recruited on Amazon's Mechanical Turk platform (18–66, mean: 35). They were paid 35 cents.

3.2.1.2 Materials

Stimuli consisted of 3-turn discourses between Debby and Harry, as shown in (26) and (27). In the target stimuli, the first turn of each discourse consisted of a past tense EAS that realized one of the 10 evaluative adjectives explored in Experiment 1. The second turn of the target stimuli consisted of Harry asking *Are you sure?*. In the third turn, Debby uttered *Yes, I am sure that*…, with the prejacent of the EAS realized as the content of the clausal complement of *sure*. There were 6 stimuli for each of the 10 evaluative adjectives, for a total of 60 target stimuli. To manipulate the at-issueness of the generalization, in 3 of the stimuli for each adjective the truth of the generalization was more likely to follow from the common ground, as in (26); in the other 3, it was less likely to follow from the common ground, as in (27). Given the 'Non-redundancy principle for at-issue content' in (18), this means that the generalization of (26) is less at-issue than that of (27). The full set of stimuli is given in Supplementary file 5.

(26) Truth of the generalization is more likely to follow from the common ground
Debby: Jane was stupid to post her social security number on Facebook.
Harry: Are you sure?
Debby: Yes, I am sure that Jane posted her social security number on Facebook.

(27) Truth of the generalization is less likely to follow from the common ground
Debby: Jane was stupid to dance like that.
Harry: Are you sure?
Debby: Yes, I am sure that Jane danced like that.

Participants were asked whether Debby answered Harry's question: we assume that a 'yes' response meant that Harry's question targeted the prejacent of Debby's utterance, i.e., the prejacent was at-issue, and that a 'no' response meant that Harry's question did not target the prejacent, i.e., the prejacent was not at-issue. We expected the prejacent of EASs for which the truth of the generalization was more likely to follow from the common ground, i.e., was more not at-issue, to be more at-issue than the prejacent of EASs for which the truth of the generalization was less likely to follow from the common ground.

The 60 target stimuli were distributed across 6 lists so that each of the 10 adjectives occurred once per list. Each list had 5 target stimuli for which the generalization was more likely to follow from the common ground and 5 target stimuli for which the generalization was less likely to follow from the common ground. To assess whether participants were attending to the task, each list also included the same two control stimuli. We also included on each list the same four stimuli that assessed the at-issueness of other projective content (see Supplementary file 5). In sum, each of the 6 lists consisted of 16 stimuli.

3.2.1.3 Procedure

Participants were told to imagine that they are at a party and that, upon walking into the kitchen, they overhear a short dialogue between Debby, the party host, and another guest, Harry. They were randomly assigned to a list and presented with the 16 stimuli, one after the other, in random order. They gave their ratings to the question of whether Debby answered Harry's question with two radio buttons labeled 'yes' and 'no', as in Figure 6.

After rating the 16 stimuli, participants filled out a brief questionnaire about their age, their native language(s) and, if English is a native language, whether it is American English, as opposed to, e.g., Indian or Australian English. Participants were told that they would be paid no matter how they responded to these questions, in order to encourage them to answer truthfully.

3.2.1.4 Data exclusion

The data from 2 participants who did not self-identify as native speakers of American English were excluded. 5 participants answered 'no' to at least one of the two control stimuli. Their data were excluded, leaving data from 68 participants (ages: 18–66; mean: 35).

3.2.2 Results

Each of the 60 target stimuli received between 10 and 14 ratings (mean: 11.3). Figure 7 shows the proportion of 'yes' responses, indicating at-issue- ness, in the two conditions: as

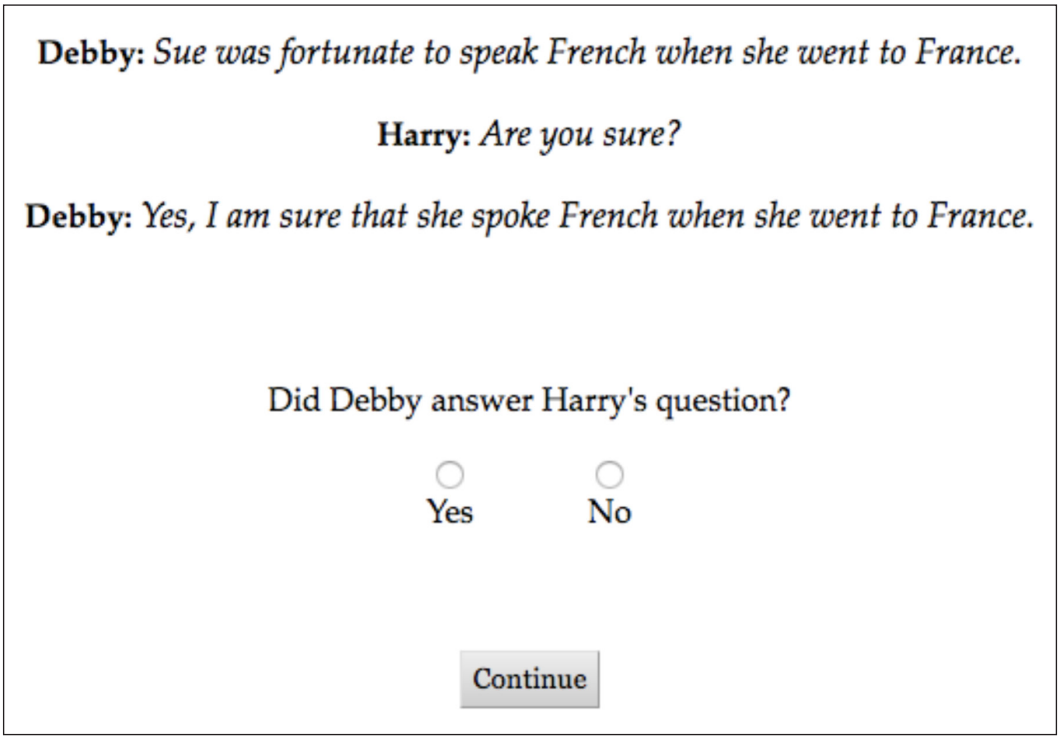

Figure 6: A sample trial in Experiment 2.

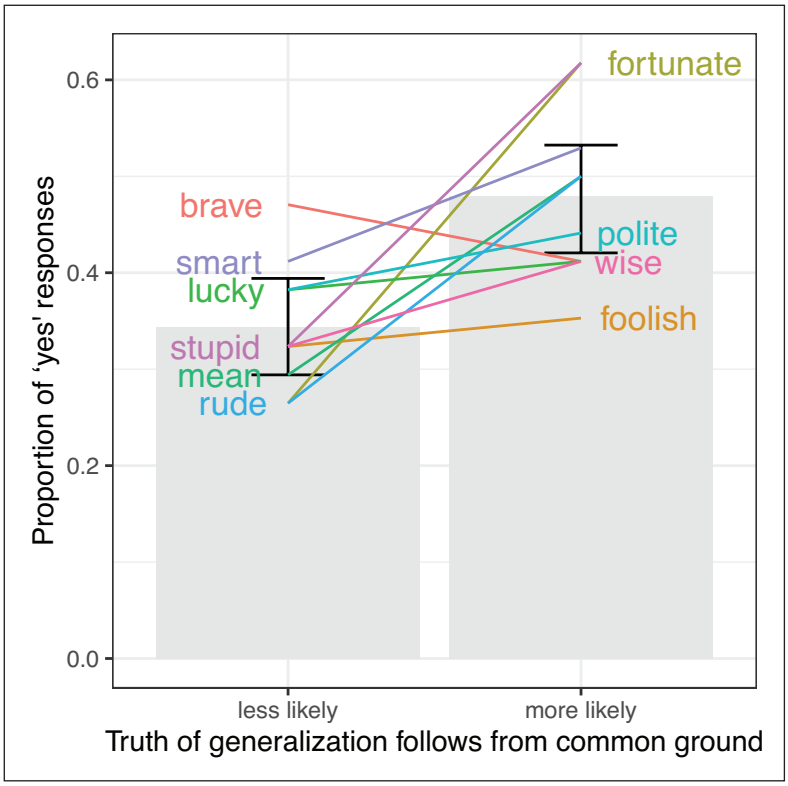

Figure 7: Proportion of 'yes' responses, indicating at-issueness of the prejacent, by condition. Error bars indicate bootstrapped 95% confidence intervals. Adjective means in the two conditions overlaid.

expected, the prejacent of EASs for which the truth of the generalization is more likely to follow from the common ground received more 'yes' responses than the prejacents of EASs for which the truth of the generalization is less likely to follow from the common ground. This finding supports the prediction of our analysis that the prejacent of EASs for which the generalization is less at-issue are more at-issue than the prejacent of EASs for which the generalization is more at-issue. As shown by the overlaid adjective means in Figure 7, the effect of the at-issueness of the generalization on the at-issueness of the prejacent was observed to varying degrees for all of the adjectives except *brave*.

To statistically evaluate the effect of the at-issueness of the generalization on the at-issueness of the prejacent, we fitted a Bayesian binomial mixed effects model with weakly informative priors using the R package `brms` (Bürkner 2017).[20] The model predicted the log odds of 'yes' over 'no' ratings from a fixed effect of how likely the truth of the generalization follows from the common ground (with 'less likely' as the reference level). We included the maximal random effects structure justified by the design: random intercepts for item (capturing random differences in at-issueness between items) and participant (capturing random differences in at-issueness between participants) as well as random by-participant slopes for the at-issueness of the generalization (capturing that the effect of the at-issueness of the generalization may vary by participant). Four chains converged after 2000 iterations each (warmup $= 1000$, $\hat{R} = 1$) for all estimated parameters).

[20] We fit a Bayesian binomial mixed effects model rather than a frequentist mixed effects model because this allowed us to fit a model with a full random structure that would not converge with frequentist methods (Nicenboim & Vasishth 2016). When we fit regular binomial mixed effects models, predicting response from a fixed effect of the discourse status of the generalization, the models only converged if we included either random by-item intercepts or random by-participant intercepts, but not both. Qualitatively, the results were identical to those obtained via the Bayesian method.

In order to evaluate the evidence for an effect of the at-issueness of the generalization, we report 95% credible intervals and the posterior probability $P\ (\beta > 0)$ that the at-issueness coefficient β is greater than zero. A 95% credible interval (CI) demarcates the range of values that comprise 95% of probability mass of our posterior beliefs such that no value inside the CI has a lower probability than any point outside it (Jaynes & Kempthorne 1976; Morey et al. 2016). There is substantial evidence for an effect if zero is (by a reasonably clear margin) not included in the 95% CI and $P(\beta > 0)$ is close to zero or one. Posterior probabilities tell us the probability that the parameter has a certain value, given the data and model (these probabilities are not frequentist *p*-values). In order to present statistics as close to widely used frequentist practices, and following Nicenboim & Vasishth (2016), we defined an inferential criterion that seems familiar (95%), but the strength of evidence should not be taken as having clear cut-off points (such as in a null-hypothesis significance testing framework).

The model provided evidence for the predicted effect of the at-issueness of the generalization on the at-issueness of the prejacent: the prejacent of EASs for which the truth of the generalizations was more likely to follow from the common ground were more likely to receive a 'yes' (at-issue) rating than the prejacent of EAS for which the truth of the generalization was less likely to follow from the common ground (posterior mean $\beta = 1.29$, 95% CI = [0.69,1.87], $P(\beta > 0) = 1$). This finding suggests that the prejacent of EASs is more at-issue when the generalization is less at-issue than when the generalization is more at-issue, as predicted by the analysis in Section 2.

3.3 Summary and discussion

The question-based analysis of EASs developed in Section 2 predicts that the more the generalization is at-issue, the more the prejacent is projective and not-at-issue. In this section, we provided experimental evidence for these two predictions: Experiment 1 showed that the prejacent is more projective when the generalization is more at-issue than when the generalization is less at-issue, and Experiment 2 showed that the prejacent is more not-at-issue when the generalization is more at-issue than when the generalization is less at-issue.

The analysis makes at least two additional predictions, to be investigated in future work. The first concerns the 'Non-redundancy principle for at-issue content' in (18). Our experiments manipulated the extent to which the truth of the generalization follows from the common ground and we hypothesized, by this principle, that this manipulation led to differences in the at-issueness of the generalization. Future research should experimentally investigate the principle itself, by manipulating the extent to which the truth of an utterance content follows from the common ground and measuring the at-issueness of the content. A second prediction concerns the question-based analysis itself. While the analysis is formulated with a gradient notion of at-issueness, both of our experiments involved only two levels of at-issueness for the generalization of EASs. Future research should investigate the prediction that the projectivity and not-at-issueness of the prejacent is sensitive to more fine-grained distinctions in the at-issueness of the generalization.

Finally, our investigation points to an open question about the relationship between the at-issueness of utterance content and what follows from the common ground about the content. The 'Non-redundancy principle for at-issue content' concerns the extent to which the truth of an utterance content follows from the common ground. In our experiments, this was manipulated such that the truth of the generalization was more likely to follow from the common ground in one condition and less likely in the other. The 'less likely' condition was implemented differently in the two experiments: the truth of the generalization was less likely to follow from the common ground by virtue of, in Experiment 1, its

falsity being likely to follow from the common ground and, in Experiment 2, the common ground being neutral with respect to the truth of the generalization. The 'Non-redundancy principle for at-issue content' does not distinguish between these two implementations: in either one, the truth of the generalization is less likely to follow from the common ground than in the 'more likely' condition. Future research should consider how the extent to which the falsity of utterance content follows from the common ground relates to its at-issueness.

4 Interspeaker variation

Before concluding the paper, we briefly return to the interspeaker variation we noted in Section 1: not all native speakers of American English would produce NEASs in which the prejacent does not project; many speakers prefer variants with *enough*. What might this variation be due to?

One hypothesis is that speakers differ in their lexical entries for evaluative adjectives. The acceptability rating study presented in Supplementary file 1 found that about 20–30% of the 94 self-reported native speakers of American English judged NEASs in which the prejacent does not project to be acceptable. Let's assume that such speakers have a lexical entry for evaluative adjectives according to which the prejacent is not lexically specified as a presupposition (as in the analysis developed in Section 2); this would allow them to produce NEASs in which the prejacent projects as well as ones in which it doesn't project. And let's assume that the remaining speakers (about 70–80%) are ones that would not produce NEASs in which the prejacent does not project. We can hypothesize that they have a lexical entry according to which the prejacent is lexically specified as a presupposition, thereby resulting in them preferring to produce EASs in which the prejacent is not-at-issue and therefore projects. When interpreting NEASs, speakers in this second group should consistently assign projecting interpretations to prejacents of NEASs in which the truth of the generalization is less likely to follow from the common ground: not only is the projecting interpretation the lexically specified one, but it is also supported by the common ground. A similar hypothesis can be found in Karttunen et al. (2014: 243), who suggested that there are about 3 times as many speakers who prefer giving interpretations to NEASs in which the prejacent projects, using the lexical entry in (5a), than speakers who prefer giving interpretations to NEASs in which the prejacent doesn't project, using the lexical entry in (5b). Thus, both hypotheses lead to the expectation that about 70–80% of native speakers of American English consistently assign projecting interpretations to NEASs in which the truth of the generalization is less likely to follow from the common ground.

A post-hoc analysis of the findings of Experiment 1 suggests that both hypotheses should be rejected. Recall that each of the 134 participants in Experiment 1 rated the projectivity of the prejacent of NEASs on a 7-point Likert scale: they rated the projectivity of the prejacent of 5 NEASs for which the truth of the generalization is more likely to follow from the common ground and of 5 NEASs for which the truth of the generalization is less likely to follow from the common ground. To explore the aforementioned hypotheses, we calculated each participants' mean projectivity ratings for these two types of NEAS: these mean projectivity ratings are an indication of how projective the 134 participants rated the prejacent of the two types of NEAS. Under both hypotheses, we expect 70–80% of these participants to assign highly projective interpretations to the prejacent of NEASs for which the truth of the generalization is less likely to follow from the common ground.

The histogram in Figure 8 shows the 134 participants' mean projectivity ratings for the two types of NEASs. The left panel reveals that, of the 134 participants, only 23 (17%) had mean projectivity ratings of at least 6.5 for the NEASs in which the truth of the generalization is less likely to follow from the common ground; when considering mean projectivity

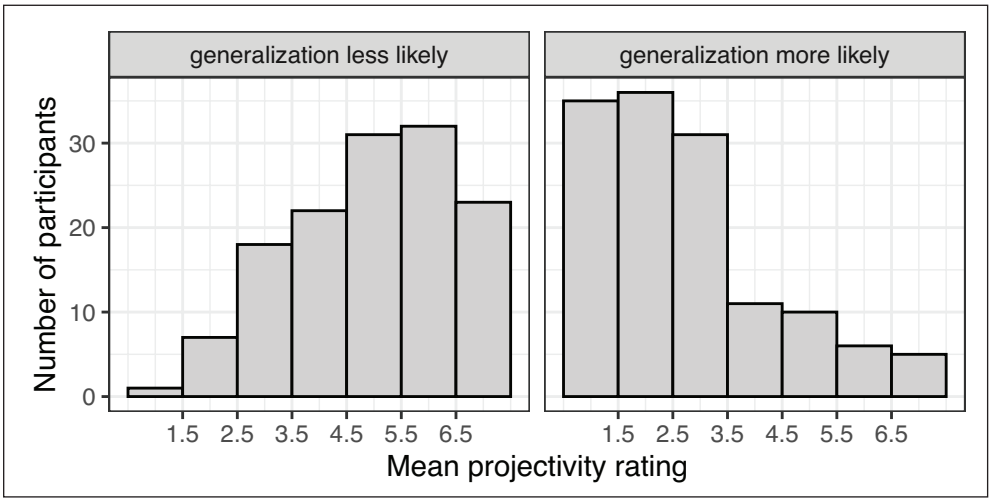

Figure 8: Histogram of participants' mean projectivity ratings in Experiment 1 by whether the truth of the generalization is less likely (left panel) or more likely (right panel) to follow from the common ground.

ratings of at least 5.5, this number still only rises to 55 (41%) participants. Thus, we do not find that a majority of the 134 participants assigned highly projective interpretations to prejacents of NEASs for which the truth of the generalization is less likely to follow from the common ground. This observation calls into question the hypothesis that a majority of native speakers of American English have a lexical entry for evaluative adjectives according to which the prejacent is a lexically-specified presupposition.

We leave further explorations of the question of how to capture the variability observed among native speakers of American English to future research.

5 Conclusions

Over the past decades, several different types of projection analyses have been proposed: on some analyses, projection is derived from the lexical specification of content as a presupposition (e.g., Heim 1983; van der Sandt 1992) or from conventional specification as a conventional implicature (e.g., Potts 2005; Murray 2014); on other analyses, projection is derived from the lexical specification of alternatives, together with pragmatic reasoning (e.g., Abusch 2002; 2010; Romoli 2015); on yet other analyses, projection is derived from discourse status (e.g., Abrusán 2011; 2013; 2016; Simons et al. 2010; 2017; Beaver et al. 2017). Which type of analysis is empirically adequate for any given projective content depends on empirical properties of that content (for discussion see, e.g., Kadmon 2001; Potts 2005; Tonhauser et al. 2013; 2018; 2019). In this paper, we argued for an analysis that derives the projection of the prejacent and the generalization of EASs from their discourse status based on two properties: first, the prejacent of EASs is not highly projective, in line with other content whose projection has been derived from its discourse status; second, there is an interaction between the projection of the prejacent and the generalization that falls out from the independently-motivated assumption that exactly one utterance content is at-issue. While it is possible to capture the projection of the prejacent and the generalization as well as their interaction from lexical specification of content as presupposed, we argued in Section 2.4 that our question-based analysis is more parsimonious, less stipulative and less idiosyncratic than Karttunen et al.'s (2014) analysis. Thus, the research presented here shows once again that the observation that content is projective provides insufficient grounds for settling on a particular analysis; rather, additional properties of the content must be considered, such as how projective it is, whether it is

associated with a Strong Contextual Felicity constraint (Tonhauser et al. 2013), and which aspects of discourse its projection depends on. Rather than making the default assumption that content projects because it is a lexically specified presupposition, future research must consider the growing landscape of analyses of projective content and properties of such content that identify which analysis is empirically adequate.

There are at least two questions about EASs that should be addressed in future work. The first concerns the interaction between the projection of the prejacent and the generalization. As discussed in Section 1, this interaction has not yet been observed for other projective content. Why does it arise with EASs, but not, for instance, for utterances of sentence with factive or change-of-state predicates? We can only speculate here that it may have to do with the prejacent and the generalization of EASs being independent of one another in that neither is a precondition for the truth of the other. To illustrate, consider (1) *Feynman was stupid to dance on the table*. For the generalization to be true, it does not matter whether Feynman danced on the table; likewise, for the prejacent to be true, it does not matter whether the degree to which Feynman dancing on the table is stupid was higher than the contextual standard of stupid. In being independent of one another, the lexical entailments of EASs differ from the lexical entailments of sentences with factive or change-of-state predicates, like *Sam knows that the meeting was canceled* or *Sam stopped going to church*: in order for Sam to know that the meeting was canceled, the content of the complement must be true, and in order for Sam to not go to church anymore, the prestate content, that they previously went to church, must be true. Future research needs to establish whether this is the reason for the interaction between the projection of the prejacent and the generalization.

A second question that future research should address is how information structure constrains the projection of the prejacent and the generalization. As mentioned in Section 2.1, previous research has established that information structure, in particular prosodically marked focus, provides a cue to the questions addressed by utterances of sentences with other projective content. We hypothesize that prosodically marked focus will also provide a cue to the question addressed by EASs and, hence, to which content projects.

Additional Files

The additional files for this article can be found as follows:

- **Supplementary file 1.** Acceptability rating experiment. DOI: https://doi.org/10.5334/gjgl.701.s1
- **Supplementary file 2.** Corpus-based web study. DOI: https://doi.org/10.5334/gjgl.701.s1
- **Supplementary file 3.** Norming study for Experiment 1. DOI: https://doi.org/10.5334/gjgl.701.s1
- **Supplementary file 4.** Materials of Experiment 1. DOI: https://doi.org/10.5334/gjgl.701.s1
- **Supplementary file 5.** Materials of Experiment 2. DOI: https://doi.org/10.5334/gjgl.701.s1

Acknowledgements

We are very grateful to Mandy Simons for critical feedback throughout the course of this project. For helpful comments we also thank David Beaver, Ashwini Deo, Lauri Karttunen, Stanley Peters, Craige Roberts, Stephanie Solt, the anonymous reviewers and the associ-

ate editor Min-Joo Kim as well as the audiences at the following venues: the 6th meeting of the DFG Network *Questions in Discourse* at the University of Göttingen, the University of Stuttgart, the 16th Sklarska Poreba workshop, the Redrawing Pragmasemantic Borders workshop at the University of Groningen, the University of Cologne, Oxford University, the University of Tübingen, the University of Göttingen, the *Experimental Pragmatics* 2015 conference in Chicago, The Ohio State University and Cornell University. We thank Willy Cheung for his assistance with the corpus study and Timo Roettger for consultation on the Bayesian regression analysis. We gratefully acknowledge financial support from the *National Science Foundation* (grants BCS-0952571 and BCS-1452674), the Alexander von Humboldt Foundation and the Ohio State University College of Arts and Sciences Targeted Investment in Excellence initiative.

Competing Interests

The authors have no competing interests to declare.

References

Abbott, Barbara. 2000. Presuppositions as non-assertions. *Journal of Pragmatics* 32. 1419–1437. DOI: https://doi.org/10.1016/S0378-2166(99)00108-3

Abrusán, Márta. 2011. Predicting the presuppositions of soft triggers. *Linguistics & Philosophy* 34. 491–535. DOI: https://doi.org/10.1007/s10988-012-9108-y

Abrusán, Márta. 2013. A note on quasi-presuppositions and focus. *Journal of Semantics* 30. 257–265. DOI: https://doi.org/10.1093/jos/ffs010

Abrusán, Márta. 2016. Presupposition cancellation: Explaining the 'soft-hard' trigger distinction. *Natural Language Semantics* 24. 165–202. DOI: https://doi.org/10.1007/s11050-016-9122-7

Abusch, Dorit. 2002. Lexical alternatives as a source of pragmatic presupposition. *Semantics and Linguistic Theory* 12. 1–19. DOI: https://doi.org/10.3765/salt.v12i0.2867

Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27. 37–80. DOI: https://doi.org/10.1093/jos/ffp009

Amaral, Patrícia Matos, Craige Roberts & E. Allyn Smith. 2007. Review of 'The Logic of Conventional Implicature' by Chris Potts. *Linguistics & Philosophy* 30. 707–749. DOI: https://doi.org/10.1007/s10988-008-9025-2

Anand, Pranav & Valentine Hacquard. 2014. Factivity, belief and discourse. In Luka Crnič & Uli Sauerland (eds.), *The art and craft of semantics: A festschrift for Irene Heim*, 69–90. MIT Working Papers in Linguistics. https://semanticsarchive.net/Archive/jZiNmM4N/.

AnderBois, Scott, Adrian Brasoveanu & Robert Henderson. 2015. At-issue proposals and appositive impositions in discourse. *Journal of Semantics* 32. 93–138. DOI: https://doi.org/10.1093/jos/fft014

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics & Philosophy* 25. 1–36. DOI: https://doi.org/10.1023/A:1014346114955

Beaver, David. 2001. *Presupposition and assertion in dynamic semantics*. Stanford, CA: CSLI Publications.

Beaver, David. 2010. Have you noticed that your belly button lint colour is related to the colour of your clothing? In Rainer Bäuerle, Uwe Reyle & Ede Zimmermann (eds.), *Presuppositions and discourse: Essays offered to Hans Kamp*, 65–99. Oxford: Elsevier. DOI: https://doi.org/10.1163/9789004253162

Beaver, David & Brady Clark. 2008. *Sense and sensitivity: How focus determines meaning*. Oxford: Wiley-Blackwell. DOI: https://doi.org/10.1002/9781444304176

Beaver, David, Craige Roberts, Mandy Simons & Judith Tonhauser. 2017. Questions Under Discussion: Where information structure meets projective content. *Annual Review of Linguistics* 3. 265–284. DOI: https://doi.org/10.1146/annurev-linguistics-011516-033952

Bürkner, Paul-Christian. 2017. brms: An R package for bayesian multilevel models using Stan. *Journal of Statistical Software* 80(1). 1–28. DOI: https://doi.org/10.18637/jss.v080.i01

Christensen, Rune Haubo Bojesen. 2013. "ordinal": Regression models for ordinal data. R package. Version ordinal_2015.6-28. https://cran.rproject.org/web/packages/ordinal/ordinal.pdf.

Cummins, Chris & Hannah Rohde. 2015. Evoking context with contrastive stress: Effects on pragmatic enrichment. *Frontiers in Psychology* 6. Article 1779. DOI: https://doi.org/10.3389/fpsyg.2015.01779

de Marneffe, Marie, Mandy Simons & Judith Tonhauser. 2019. The CommitmentBank: Investigating projection in naturally occurring discourse. *Sinn und Bedeutung* 23. 107–124. DOI: https://doi.org/10.18148/sub/2019.v23i2.601

Destruel, Emilie, Edgar Onea, Daniel Velleman, Dylan Bumford & David Beaver. 2015. A cross-linguistic study of the non-at-issueness of exhaustive inferences. In Floran Schwarz (ed.), *Experimental approaches to presupposition*, 135–156. Springer. DOI: https://doi.org/10.1007/978-3-319-07980-6

Djärv, Kajsa & Hezekiah Akiva Bacovcin. 2017. Prosodic effects on factive presupposition projection. *Semantics and Linguistic Theory* 27. 116–133. DOI: https://doi.org/10.3765/salt.v27i0.4134

Gazdar, Gerald. 1979. A solution to the projection problem. In Choon-Kyu Oh & David Dinneen (eds.), *Syntax and semantics 11: Presupposition*, 57–89. New York: Academic Press. DOI: https://doi.org/10.1163/9789004368880_003

Groenendijk, Jeroen. 1999. The logic of interrogation. *Semantics and Linguistic Theory* 9. 109–126. DOI: https://doi.org/10.3765/salt.v9i0.2835

Heim, Irene. 1982. *The semantics of definite and indefinite noun phrases*: University of Massachusetts, Amherst dissertation.

Heim, Irene. 1983. On the projection problem for presuppositions. *West Coast Conference on Formal Linguistics* 2. 114–125.

Jaynes, Edwin & Oscar Kempthorne. 1976. Confidence intervals vs. Bayesian intervals. In William Leonard Harper & Clifford Alan Hooker (eds.), *Foundations of probability theory, statistical inference, and statistical theories of science*, 175–257. Dordrecht: Springer Netherlands. DOI: https://doi.org/10.1007/978-94-010-1853-1

Kadmon, Nirit. 2001. *Formal pragmatics*. Oxford: Blackwell.

Karttunen, Lauri. 1971. Some observations on factivity. *Papers in Linguistics* 4. 55–69. DOI: https://doi.org/10.1080/08351817109370248

Karttunen, Lauri. 2013. You will be lucky to break even. In Tracy Holloway King & Valeria de Paiva (eds.), *From quirky case to representing space: Papers in honor of Annie Zaenen*, 167–180. Stanford: CSLI Publications.

Karttunen, Lauri, Stanley Peters, Annie Zaenen & Cleo Condoravdi. 2014. The chameleon-like nature of evaluative adjectives. *Empirical Issues in Syntax and Semantics* 10. 233–250.

Kertz, Laura. 2010. The argument structure of evaluative adjectives. A case of pseudo-raising. In Norbert Hornstein & Maria Polinsky (eds.), *Movement theory of control*, 269–298. Amsterdam/Philadelphia: John Benjamins. DOI: https://doi.org/10.1017/S0022226711000077

Kiparsky, Paul & Carol Kiparsky. 1970. Fact. In Manfred Bierwisch & Karl Erich Heidolph (eds.), *Progress in linguistics*, 143–173. The Hague: Mouton.

Mahler, Taylor. 2019. Does at-issueness predict projection? It's complicated! *North East Linguistic Society* 49. 245–255.

Morey, Richard D., Rink Hoekstra, Jeffrey N. Rouder, Michael D. Lee & Eric J. Wagenmakers. 2016. The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin and Review* 23(1). 103–123. DOI: https://doi.org/10.3758/s13423-015-0947-8

Murray, Sarah. 2014. Varieties of update. *Semantics & Pragmatics* 7(2). 1–53. DOI: https://doi.org/10.3765/sp

Nicenboim, Bruno & Shravan Vasishth. 2016. Statistical methods for linguistic research: Foundational Ideas, Part II. *Linguistics and Language Compass* 10(11). 591–613. DOI: https://doi.org/10.1111/lnc3.12207

Norrick, Neal R. 1978. *Factive adjectives and the theory of factivity*. Tübingen: Niemeyer.

Oshima, David Y. 2009. Between being wise and acting wise: A hidden conditional in some constructions with propensity adjectives. *Journal of Linguistics* 45. 363–393. DOI: https://doi.org/10.1017/S0022226709005684

Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199273829.001.0001

R Core Team. 2016. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.Rproject.org/.

Roberts, Craige. 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. DOI: https://doi.org/10.3765/sp.5.6

Romoli, Jacopo. 2015. The presuppositions of soft triggers are obligatory scalar implicatures. *Journal of Semantics* 32. 173–291. DOI: https://doi.org/10.1093/jos/fft017

Schlenker, Philippe. 2010. Local contexts and local meanings. *Philosophical Studies* 151. 115–142. DOI: https://doi.org/10.1007/s11098-010-9586-0

Simons, Mandy. 2001. On the conversational basis of some presuppositions. *Semantics and Linguistics Theory* 11. 431–448. DOI: https://doi.org/10.3765/salt.v11i0.3099

Simons, Mandy. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua* 117. 1034–1056. DOI: https://doi.org/10.1016/j.lingua.2006.05.006

Simons, Mandy, David Beaver, Craige Roberts & Judith Tonhauser. 2017. The Best Question: Explaining the projection behavior of factive verbs. *Discourse Processes* 3. 187–206. DOI: https://doi.org/10.1080/0163853X.2016.1150660

Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts. 2010. What projects and why. *Semantics and Linguistic Theory* 11. 309–327. DOI: https://doi.org/10.3765/salt.v20i0.2584

Smith, E. Allyn & Kathleen Currie Hall. 2011. Projection diversity: Experimental evidence. *2011 ESSLLI Workshop on Projective Content,* 156–170.

Stalnaker, Robert. 1999. Assertion. In Robert Stalnaker (ed.), *Context and content: Essays on intentionality in speech and thought*, 78–95. Oxford: Oxford University Press. Originally published Syntax and Semantics 9 (1978). DOI: https://doi.org/10.1093/0198237073.001.0001

Stevens, Jon Scott, Marie-Catherine de Marneffe, Shari R. Speer & Judith Tonhauser. 2017. Rational use of prosody predicts projectivity in manner adverb utterances. *Annual Meeting of the Cognitive Science Society* 39. 1144–1149.

Syrett, Kristen & Todor Koev. 2015. Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics* 32. 525–577. DOI: https://doi.org/10.1093/jos/ffu007

Tonhauser, Judith. 2012. Diagnosing (not-)at-issue content. *Semantics of Under-represented Languages in the Americas* 6. 239–254.

Tonhauser, Judith. 2016. Prosodic cues to speaker commitment. *Semantics and Linguistic Theory* 26. 934–960. DOI: https://doi.org/10.3765/salt.v26i0.3788

Tonhauser, Judith. in press. Projection variability in Paraguayan Guaraní. *Natural Language and Linguistic Theory*. DOI: https://doi.org/10.1007/s11049-019-09462-x

Tonhauser, Judith, David Beaver, Craige Roberts & Mandy Simons. 2013. Toward a taxonomy of projective content. *Language* 89. 66–109. DOI: https://doi.org/10.1353/lan.2013.0001

Tonhauser, Judith, David Beaver & Judith Degen. 2018. How projective is projective content? Gradience in projectivity and at-issueness. *Journal of Semantics* 35. 495–542. DOI: https://doi.org/10.1093/jos/ffy007

Tonhauser, Judith, Marie-Catherine de Marneffe, Shari R. Speer & Jon Stevens. 2019. On the information structure sensitivity of projective content. *Sinn und Bedeutung* 23. 363–389. DOI: https://doi.org/10.18148/sub/2019.v23i2.619

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9. 333–377. DOI: https://doi.org/10.1093/jos/9.4.333

von Fintel, Kai. 2008. What is presupposition accommodation again? *Philosophical Perspectives* 22. 137–170. DOI: https://doi.org/10.1111/j.1520-8583.2008.00144.x

Wilkinson, Robert. 1970. Factive complements and action complements. *Annual Meeting of the Chicago Linguistic Society* 6. 425–444.

Xue, Jingyang & Edgar Onea. 2011. Correlation between projective meaning and at-issueness: An empirical study. *2011 ESSLLI Workshop on Projective Content,* 171–184.

How to cite this article: Tonhauser, Judith, Marie-Catherine de Marneffe and Judith Degen. 2020. Evaluative adjective sentences: A question-based analysis of projection. *Glossa: a journal of general linguistics* 5(1): 87.1–33. DOI: https://doi.org/10.5334/gjgl.701

Submitted: 23 May 2018 **Accepted:** 22 May 2020 **Published:** 01 September 2020

Copyright: © 2020 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See http://creativecommons.org/licenses/by/4.0/.

]u[*Glossa: a journal of general linguistics* is a peer-reviewed open access journal published by Ubiquity Press. **OPEN ACCESS**