Kleene'ing up your attitude: Against universal presupposition projection*

Maik Thalmann
University of Göttingen

Andrea Matticchio
UMass Amherst

Version: November 4, 2025

It is often assumed that presuppositions in the scope of attitude predicates, a case of presupposition projection that is empirically unexplored, must be satisfied at every world that the attitude predicate quantifies over. This paper presents experimental evidence that such cases of presupposition projection can instead be modeled within a trivalent semantic framework, specifically using the Strong Kleene logic. The experimental results pose challenges for standard accounts based on a Universal Projection hypothesis, which fail to predict key aspects of the data. Instead, the findings align with theories in which presupposition failures can participate in compositional interpretation. Further, we find what we call Strawsonian squeamishness, increased variance in participants' response when faced with presupposition failure.

Keywords: presupposition projection, Strong Kleene, attitude predicates, local accommodation, experimental semantics

1 Introduction

To give an account of presupposition projection is to be able to predict what the presupposition of complex sentences is from the presupposition of its parts. A long tradition of research has addressed this problem as part of the compositional derivation of a sentence's meaning, such that at every application of a logical operator to an interpretable linguistic form, the meaning of the operator determines what the presupposition of the resulting expression is. An example is given in (1) for English *not*, where presuppositions are characterized in terms of definedness conditions.

(1) $[\![\textbf{not } A]\!]$ is defined only if $[\![A]\!]$ is defined. If defined, $[\![\textbf{not } A]\!]$ is true if and only if $[\![A]\!]$ is false.

A common intuition is that logical operators with similar meanings should lead to similar patterns of presupposition projection. Descriptively, operators have often been divided into groups according to their behavior with respect to projection, as in the case of holes, filters, and plugs for Karttunen (1973). The discussion about what the exact grouping should be, and why the groupings are the way they are, has been oscillating between trying to be more explanatory by giving unified accounts, and accepting to stipulate the projective behavior on an operator-by-operator basis.

*For feedback on this project we want to thank Clemens Mayr, Benjamin Spector, Danny Fox, and Thomas Weskott. For helpful comments and discussion, we want to express our gratitude to the members of the colloquium for experimental linguistics (LEGEBLING) in Göttingen, the audiences at the Bounded Meaning workshop in Amherst, USA in 2024, the Workshop on Speech Act Related Operators 2024 at ZAS Berlin, the Gaps and imprecision in natural language semantics workshop at ZAS Berlin, and the 24th Amsterdam Colloquium.

This project was carried out within the RTG 2636 "Form-Meaning Mismatches," funded by the *Deutsche Forschungsgemeinschaft*, project number: 429844083.

In this paper, we pursue the hypothesis that the behavior of a logical operator with respect to presupposition projection does not have to be stipulated lexically, but follows from the logical contribution of the operator and a single logic that can deal with truth-value gaps. Our focus will be on presupposition projection out of attitude predicates, which we will argue are quantificational operators. We will try and show that a Strong Kleene trivalent logic is a good approximation of such an underlying natural language logic for these cases, just as has been proposed for other quantificational environments (George 2010, Fox 2013). Our hypothesis can thus be stated as in (2).

(2) **Generalized Strong Kleene hypothesis:**
The logic that predicts presupposition projection is Strong Kleene for all domains of quantification.

The paper is structured as follows. Section 2 presents the prediction of a generalized Strong Kleene logic for presupposition projection out of the scope of quantifiers over individuals and that of attitude predicates, comparing it to the alternative approach of generalized universal projection that has been the undisputed standard for attitude predicates. Section 3 offers a review of previous experimental research focusing on two salient aspects: what factors have been proposed as drivers of variation in the data, and what methodologies have been adopted to investigate presupposition projection. Additionally, this section presents a novel method to diagnose the effects of presupposition failures in experimental data by focusing on standard deviations. Section 4 presents our experimental investigation of presupposition projection out of the scope of the attitude predicate *be certain*, which supports our hypothesis in (2) and is in conflict with approaches positing universal projection. Section 5 concludes.

2 Presupposition projection and quantificational environments

2.1 Quantifiers over individuals

Research on presupposition projection out of the scope of quantifiers has been affected by a data problem: there is no broad consensus about what empirical picture a theory of projection should derive for quantifiers over individuals. Consider the sentence in (3), where the presupposition trigger *stop* takes a universally quantified subject. In this case, it is generally agreed upon that for the sentence to be true, it must be true that all of the students used to smoke.

(3) All of the students λx [x stopped smoking.]
$\rightsquigarrow$ All of the students used to smoke.

A different inference might be drawn in cases like (4a) and (4b). In particular, some authors report that such sentences may be judged as true even in contexts where the presupposition is not satisfied universally, i.e., where not all of the students used to smoke (see Karttunen & Peters 1979, Beaver 2001 a.o.). We will turn to the experimental findings in this regard shortly.

(4) a. Not all of the students λx [x stopped smoking.]
$\not\rightsquigarrow$ All of the students used to smoke.
b. Some of the students λx [x stopped smoking.]
$\not\rightsquigarrow$ All of the students used to smoke.

One possible way to approach these cases was spelled out by Heim (1983). According to this proposal, the semantic computation has presuppositions project with universal quantificational force by default. All the sentences above, provided that real quantificational expressions are involved, should then presuppose that all the students used to smoke. We will refer to this hypothesis as the Universal Projection hypothesis, given in (5) for quantifying determiners.[1]

[1] Here and throughout we will center our discussion on universal projection on the work by Heim. We do not mean to suggest, however, that Heim's is the only approach that predicts, and is designed to predict, universal presupposition projection from

(5) **Universal Projection hypothesis (for quantifying determiners):**
Given a quantifying determiner $\mathcal{Q}$ and properties A, B_C such that for all x, $B_C(x)$ is defined only if $C(x)$ is true, the formula $\mathcal{Q}(A)(B_C)$ is defined only if for all x such that $A(x)$ is true, $C(x)$ true.

On top of this default rule for presupposition projection, other mechanisms may apply that lead to weaker presuppositions, potentially addressing the intuitions reported in (4). On the one hand, some indefinite expressions are not really quantificational according to the view of Heim (1983) and much subsequent work. On the other hand, natural language semantics might be equipped with alternative ways of interpreting certain presuppositional expressions, like local accommodation, which we come back to in Section 3. However, these strategies are given as possibilities to derive potential readings that these sentences might have according to some; the actual desiderata for a theory of presupposition projection from quantified expressions remain unclear.

A different approach to the same cases is adopted by Fox (2013), building on work by George (2010). Fox argues that a system based on the trivalent Strong Kleene logic allows for an adequate model of presupposition projection out of quantifiers over individuals, like in the cases above. Importantly, Strong Kleene does not exactly predict presupposition projection, but rather truth conditions for a quantified formula, given that all predicates are assumed to be total functions to a set of three truth-values {T, #, F}. The predictions of Strong Kleene for quantified formulas are given schematically in (6) and (7).

(6) **Strong Kleene projection out of a universal quantifier:**
The truth-value of a formula $(\forall x \in \mathcal{D})\ P(x)$ is:
   a. T if $P(x) = \text{T}$ for all $x \in \mathcal{D}$;
   b. F if there is an $x \in \mathcal{D}$ such that $P(x) = \text{F}$;
   c. # otherwise (if there is an $x \in \mathcal{D}$ such that $P(x) = \#$ and no $x \in \mathcal{D}$ such that $P(x) = \text{F}$).

(7) **Strong Kleene projection out of an existential quantifier:**
The truth-value of a formula $(\exists x \in \mathcal{D})\ P(x)$ is:
   a. T if there is an $x \in \mathcal{D}$ such that $P(x) = \text{T}$;
   b. F if $P(x) = \text{F}$ for all $x \in \mathcal{D}$;
   c. # otherwise (if there is an $x \in \mathcal{D}$ such that $P(x) = \#$ and no $x \in \mathcal{D}$ such that $P(x) = \text{T}$).

According to this view, presupposition projection cannot simply be called existential or universal. Rather, the presupposition of the quantifier's nuclear scope has to be satisfied to different degrees when the entire sentence is true and when it is false. In particular, for universal quantification in (6), the predicate needs to be true of all elements in the domain of quantification for the sentence to be true, but it only needs to be false of a single element for the sentence to be false overall. Because of this asymmetry, a different disjunct of the presupposition is verified depending on whether the sentence is asserted or negated. For a negated sentence, the disjunct with existential quantification associated with the falsity conditions must be true, while a sentence without negation makes the disjunct with universal quantification true. This is in line with the intuitions reported in (3) and (4). The valid inferences in the case of universal quantification are schematized in (8), where A_B is a predicate such that $A_B(x) \neq \#$ only if $B(x)$ is true, similarly to the notation in (5).

(8) a. If $(\forall x \in \mathcal{D})\ A_B(x)$ is true, it follows that $(\forall x \in \mathcal{D})\ B(x)$ is true;
   b. If $\neg(\forall x \in \mathcal{D})\ A_B(x)$ is true, it only follows that $(\exists x \in \mathcal{D})\ B(x)$ is true.

The schemas given in (8) predict that it is logically valid to infer that the presupposition holds universally in (3), under *all of the students*, but not in (4a), under *not all of the students*. However, similar to the approach adopted by Heim (1983), Fox (2013) assumes that mechanisms of presupposition weakening

quantified environments. While this has recently come to be criticized for quantifiers over individuals, there are still various approaches that adhere to the universal predictions (**schlenker2008presupposition**, Schlenker 2009; see also Mandelkern 2024 for similar predictions in the modal domain).

and presupposition strengthening can apply in certain cases, and additional differences can arise from the different nature of some presupposition triggers. Some of these options will be discussed in Section 3.

The Universal Projection hypothesis and Strong Kleene are two theoretical possibilities to model presupposition projection out of quantified environments. Their predictions have been discussed in the literature with respect to quantifiers over individuals, like *all of the students*. However, these predictions have generally not come without qualification, such as the potential availability of additional mechanisms of presupposition weakening and strengthening, and because of the frequent uncertainty about what the empirical desiderata, or good predictions, of a theory of presupposition projection are. In the next section, we will move on from quantifiers over individuals and focus on the predictions of these two approaches in the domain of attitude reports.

2.2 What happened to attitudes?

The research on presupposition projection out of quantificational environments has centered on quantifiers over individuals, neglecting modal/intensional operators. A contributing factor to this narrow focus is certainly that the debate about whether and which modal operators should be analyzed as quantifiers is not settled.[2] Furthermore, presupposition projection out of modal environments has been neglected from the empirical point of view, to our knowledge. Seminal work by Heim (1992), grounded in dynamic semantics, postulates that predicates like *think* should presuppose that their complement be defined at every accessible world they consider. This predicts the validity of the inference in (9).

(9) Aditi thinks that Skye stopped smoking.
⇝ Aditi thinks that Skye used to smoke.

Throughout this paper, we will assume that at least some attitude predicates can receive adequate treatment with a quantificational analysis along the lines of Hintikka (1969), leaving to future research how to give an account of our findings in other frameworks, if one is to be preferred. If a quantificational analysis of *think* is adopted, such that the verb introduces universal quantification over a set of doxastically accessible worlds, definedness and truth conditions following Heim (1992) can be given as in (10), where $\textsc{dox}_a^w$ is the set of doxastically accessible worlds for an individual a at a world w.

(10) $[\![\textbf{a thinks p}]\!]^w$ is defined only if $[\![\textbf{p}]\!]^{w'}$ is defined for all $w' \in \textsc{dox}_a^w$. If defined, $[\![\textbf{a thinks p}]\!]^w$ is true if and only if $[\![\textbf{p}]\!]^{w'}$ is true for all $w' \in \textsc{dox}_a^w$.

We will say that a presupposition that projects according to the rule in (10) projects universally out of the attitude predicate. This schema for presupposition projection is entirely parallel to the one given in (5) for quantifiers over individuals, replacing the restrictor of a quantificational determiner with a set of accessible worlds that the attitude verb quantifies over. In Heim (1992), every attitude predicate is assumed to behave in a way similar to (10) regarding presupposition projection, the difference between predicates being what set of accessible worlds is quantified over.[3]

A qualification is due at this point. According to the interpretation in (10), *think* does not behave like a hole for presupposition projection, i.e., a presupposition in its complement does not become, unaltered, a presupposition of the complex sentence. Other work, like that by Geurts (1999), instead assumes that attitude predicates do behave like holes for presupposition projection on the basis of inferences like (11).

(11) Aditi thinks that Skye stopped smoking.

[2] See for example Kratzer (2006) and Moulton (2009, 2015) for the idea that some attitude predicates may not be quantificational. Additionally, see Lassiter (2017), Santorio & Romoli (2017), and Koev (2019), among others, for the view that some attitude and epistemic predicates should come with a degree argument that expresses a probability.

[3] According to Heim, some predicates may consider different sets of accessible worlds for their definedness and their truth conditions, as in the case of preferential predicates like *want*. Definedness conditions are nonetheless expressed in terms of universal quantification. In this paper, we will focus on a predicate, *be certain*, that does not have a preferential meaning.

⇝ Skye used to smoke.

Although we agree that the inference in (11) is a reasonable one in out-of-the-blue contexts, we side with much previous work (Heim 1992, Sudo 2014, Thalmann & Matticchio (to appear)) in assuming that it is also cancelable, unlike the inference in (9). A difference in the possibility to be canceled between the two inferences is highlighted by a contrast like (12). A report of Aditi's thinking that Skye stopped smoking is felicitous if Aditi believes that Skye used to smoke even if they actually did not. It is not felicitous, instead, if Skye indeed used to smoke but Aditi does not know.

(12) a. Aditi mistakenly believes that Skye used to smoke, and she thinks that they stopped.
b. # Aditi doesn't know that Skye used to smoke, and she thinks that they stopped.

Regardless of the analysis that the inference in (11) may receive, it has to be acknowledged that the inference in (9) is not cancelable and is tied to the presupposition of the sentence via the meaning of the presupposing lexical item. Our experiment will only test the uncancellable inference in (9) and leave the truth of the inference in (11) unaddressed. So even though participants may draw inferences like (9) additionally, these will not impact the experimental judgments.

Among those approaches that do not treat attitude predicates as holes for presupposition projection, the claim that projection out of attitudes has the profile in (10) has not been questioned as a desideratum. More recent work, like that by Schlenker (2009) or Mandelkern (2024), is still designed to meet the intuition that a presupposition triggered in the scope of an attitude predicate has to be satisfied universally with respect to a set of worlds. We would like to compare this hypothesis with the predictions that a modal extension of a Strong Kleene logic makes given a quantificational analysis of attitude predicates.

To do so, we consider the attitude predicate *be certain*, instead of the more prototypical *think* or *believe*. We consider *be certain* to be one of the best candidates to approximate a universal quantifier in the domain of attitude predicates. Just like a universal quantifier, *be certain* validates the equivalence in (13) with conjunction, as illustrated by (14). The equivalence in (13) characterizes universal quantifiers and distinguishes them, for example, from existential quantifiers.

(13) $\forall x\, P(x) \wedge \forall x\, Q(x) \eqq \forall x\, (P(x) \wedge Q(x))$

(14) *Equivalently:*
a. Aditi is certain that Skye smokes, and she's certain that they drink.
b. Aditi is certain that Skye both smokes and drinks.

Since an approach based on Strong Kleene places a certain amount of emphasis on a difference between asserted and negated sentences, the behavior of the quantifier under negation plays a critical role. Unlike *think* or *believe*, the predicate *be certain* is not a neg-raiser: when embedded under negation, it does not report that the embedded proposition is false, and instead licenses the indirect scalar implicature that it might be true. As is familiar, the opposite behavior is observed with the neg-raising predicate *think*.[4]

(15) a. Aditi is not certain that Skye smokes.
⇝̸ Aditi is certain that Skye doesn't smoke.
⇝ Aditi considers it possible that Skye smokes.
b. Aditi doesn't think that Skye smokes.
⇝ Aditi thinks that Skye doesn't smoke.
⇝̸ Aditi considers it possible that Skye smokes.

[4] As an additional test for neg-raising (Gajewski 2007), it is possible to observe that strong negative polarity items, like *in weeks*, are not licensed in the complement of *not be certain*, but they are in the complement of *not think*:

(i) a. * Aditi is not certain that Skye has smoked in weeks.
b. Aditi doesn't think that Skye has smoked in weeks.

A similar indirect scalar implicature is observed with the universal quantifier over individuals *all* under negation (Chierchia 2004, 2013, a.o.), exemplified in (16), which mirrors the existential inference we see with *be certain* in (15a). That these inferences are implicatures can be seen from their cancellability in contexts like (17).[5] We remain agnostic as to how indirect scalar implicatures are derived.

(16) Skye did not smoke all of the cigarettes.
⇝ Skye smoked some of the cigarettes.

(17) a. A: Did Skye smoke all of the cigarettes?
B: They didn't smoke all of the cigarettes; in fact/maybe, they didn't smoke any.
b. A: Is Aditi certain that Skye smokes?
B: She isn't certain that Skye smokes; in fact/maybe, she's certain that they don't.

The predictions of a Strong Kleene logic for a modal predicate that can be analyzed as a universal quantifier are given in (18). We want to stress that our hypothesis is not that these predictions are correct for predicates like *think* and *believe*, because neg-raising might affect the falsity and third-value-ness conditions of these predicates. Instead, they are predictions for presupposition projection out of the complement of *be certain*, which we hope to have shown behaves like a universal quantifier. We illustrate these predictions by giving the truth conditions for the example in (19), which embeds the presupposition trigger *stop*.

(18) **Strong Kleene projection out of a universal modal:**
The truth-value of a formula $(\forall\, w \in \mathcal{W})\ p(w)$ is:
a. T if $p(w) = \text{T}$ for all $w \in \mathcal{W}$;
b. F if there is an $w \in \mathcal{W}$ such that $p(w) = \text{F}$;
c. # otherwise (if there is an $w \in \mathcal{W}$ such that $p(w) = \#$ and no $w \in \mathcal{W}$ such that $p(w) = \text{F}$).

(19) Aditi is certain that Skye stopped smoking. *(Strong Kleene)*
a. T if Aditi is certain that Skye used to smoke and stopped.
b. F if Aditi considers it possible that Skye used to smoke and did not stop.
c. # if Aditi considers it possible that Skye has never smoked in the past, but she is certain that if Skye used to smoke, they stopped.

The predictions of a Strong Kleene logic for presupposition projection out of the complement of *be certain* differ from the predictions of a Universal Projection hypothesis for the truth-values F and #. In particular, Universal Projection requires that the presupposition of the complement be universally met in the accessible worlds for the attitude report to be false. A weaker, merely existential condition is imposed by Strong Kleene, as stated in (19b). For the sake of completeness, we report in (20) the truth conditions predicted by the Universal Projection hypothesis for the same example. The difference between these two sets of predictions will be the focus of our experiment, presented in Section 4.

(20) Aditi is certain that Skye stopped smoking. *(Universal Projection)*
a. T if Aditi is certain that Skye used to smoke and stopped.
b. F if Aditi is certain that Skye used to smoke and considers it possible that they did not stop.
c. # if Aditi considers it possible Skye has never smoked in the past.

[5] Our assessment differs from the one by Uegaki (2020: p. 918), who suggests that this indirect scalar implicature is a presupposition. Yet, presuppositions cannot be canceled in the same way, as illustrated with the example below:

(ii) A: Did Skye stop smoking?
B: # They didn't stop smoking; in fact/maybe, they've never smoked.

3 Previous work and some methodological points

3.1 A variety of judgments

We will now present an overview of certain factors that the previous literature has proposed as drivers of variation in the judgments about the presupposition of quantified sentences. As we have tried to make clear above, one main issue is that there is no agreement about which presuppositions a good theory should predict for the core cases.

As we discussed in the previous section, proposals about presupposition projection out of quantificational environments often come equipped with mechanisms to weaken the predicted presupposition in certain cases. One notable strategy to weaken the presupposition of a sentence is the application of local accommodation, an operation that turns presupposition failures into the classical truth-value F. Local accommodation can be written as an operator that turns trivalent propositions into bivalent ones (see also Beaver & Krahmer 2001), as in (21), but its syntactic reality is not the focus of this paper.

$$(21) \quad [\![\mathcal{A}]\!] = \lambda p_{\langle s,t\rangle} \,.\, \lambda w \,.\, \begin{cases} \mathrm{T} & \text{if } p(w) = \mathrm{T} \\ \mathrm{F} & \text{if } p(w) = \# \text{ or } p(w) = \mathrm{F} \end{cases}$$

According to Fox (2013), local accommodation can apply and weaken the presupposition of certain sentences if they contain a soft presupposition trigger, being instead (largely) unavailable if the presupposition trigger is hard. This being the case, we now attempt to clarify the distinction between soft and hard presupposition triggers. Although different diagnostics have been proposed to differentiate between these two classes, what we believe is the most widely adopted one uses contexts of explicit ignorance and conditional constructions (**abrusan2011soft**, Abusch 2002, 2010, Romoli 2015, Abrusán 2016). In a context where the speaker asserts ignorance with respect to its presupposition, a soft trigger embedded in a conditional antecedent does not lead to infelicity, whereas a hard trigger does. We illustrate this contrast in (22) with the soft trigger *stop* and the hard trigger *too*.[6]

(22) a. I don't know if Skye has ever smoked, but if they stopped smoking, Aditi will be surprised.
b. # I don't know if Taro smokes, but if Skye smokes too, Aditi will be surprised.

According to Fox's approach, a sentence containing a soft presupposition trigger can sometimes receive a classical truth-value even if the presupposition is unsatisfied in the context provided the presupposition is understood as being locally accommodated. In particular, a sentence can be judged true instead of truth-value-less if local accommodation happens within the scope of negation. On the assumption that local accommodation is not an obligatory operation, it is possible to expect variability in the judgments of sentences containing soft triggers in presupposition failure scenarios. That is, if participants do not consistently make use of local accommodation, allowing for presupposition failure to remain in some cases but not in others, this should driver larger amounts of variability. Hard triggers, on the other hand, are anticipated to show less variability instead.

In addition to the differential availability of local accommodation for soft and hard triggers, some judgments provided by Charlow (2009) complicates matters further. According to Charlow's intuitions, soft and hard presupposition triggers are associated with different projection patterns altogether when embedded in quantificational environments.[7] These judgments can be illustrated with the examples in (23). The presuppositions of hard triggers, like *also*, are claimed to project universally in all cases, including

[6] There are other differences between soft and hard triggers that have been experimentally verified. For instance, soft triggers are more felicitous than hard triggers in response to questions that make the presupposition at-issue. Below, we show this with experimental material from Chen, Thalmann & Antomo (2022), with the soft trigger *win* and the hard trigger *again*:

(iii) A: Did the duck participate in the competition?
B: The duck won the competition.

(iv) A: Did the panda score a goal for the first time yesterday?
B: # The panda scored a goal again yesterday.

when the trigger is embedded under an existential quantifier. Instead, this is not necessarily the case if the presupposition trigger is soft, as in the case of *stop*. The absence of universal projection in (23a) cannot be the result of local accommodation as defined in (21).

(23) *According to Charlow (2009):*

a. Some of these 10 students stopped smoking.
   ⇝̸ All of these 10 students used to smoke.

b. Some of these 10 students also used to drink.
   ⇝ All of these 10 students used to smoke.

These judgments, although provided without much empirical support, muddy the waters since they suggest that the hardness of a presupposition trigger can affect projection out of a quantificational environment in a way that is orthogonal to local accommodation: this approach effectively mandates different projection algorithms for the two classes of triggers. Although with a different focus, the experiment that will be presented in Section 4 will shed some light on this issue.

In addition to the differences between presupposition triggers, variation in the judgments may come from the speakers that provide them. In an experiment investigating inter-speaker variation in the distribution of universal projection, Sudo et al. (2012) tested three quantified environments embedding the presupposition trigger *both*. The critical items used in their study are presented in (24).

(24) a. Some of these three triangles have the same color as both of the circles in their own cell.

b. None of these three circles have the same color as both of the squares in their own cell.

c. Do any of these three squares have the same color as both of the triangles in their own cell?

As is typical in the covered box paradigm, participants were presented with an overt and a covered picture for each critical item and were asked to choose the one the sentence referred to. In each of the critical cases, the overt image failed to satisfy the hypothetical universal presupposition of the sentence: for example, in the case of (24a), at least one triangle had only one circle in its cell, rather than two, as universal projection of the presupposition of *both* would require. Under the logic of the covered box paradigm, this setup is designed to differentiate participants who derive a universal inference (expected to choose the covered image) from those who do not (expected to choose the overt one).

Sudo et al. (2012) find that participants' choices, while not uniform, give rise to an implicational hierarchy about the different embedding quantifiers. A participant who chooses the covered box for *some* also has a universal presupposition with the negative existential quantifier *none* and the existential polar question. If a participant instead does not derive a universal inference for *some*, they may still derive one for the negative existential, the existential polar question, or both.[8]

While the study aims at capturing speaker-level variability, it faces some limitations in its design. Most notably, only a single item was used for each quantifier condition. This makes it difficult to evaluate the reliability of participant-level inferences, as item-level variation is neither estimable nor controlled for. That is, without sufficiently many items per condition, measurement error cannot be estimated or mitigated. This concern is further underscored by the fact that only 32 percent of the participants responded correctly to the set of unambiguous filler items. However, even granting reliability, the limited number of critical items complicates the generalization of the findings beyond the specific linguistic material tested (cf. Clark 1973).

Given these previous results and our discussion thereof, we will not design an experiment to look for inter-speaker variation, and we will instead try to see if a stronger result can be obtained, one where repeated measures over items and over participants can reveal a patter that conforms to either the Universal

[7] Although Fox (2013) cites Charlow (2009) in support of the idea that some judgment variability may come from the soft-hard presupposition distinction, the predictions of the two approaches remain quite different. For example, Fox does not predict the presupposition to project universally in (23b).

[8] Note that the result that some people did not derive a universal presupposition for (24a) is in conflict with the intuitions presented by Charlow (2009), on the assumption that *both* is a hard presupposition trigger.

Projection or the Strong Kleene hypothesis. On the other hand, we will introduce both soft and hard presupposition triggers in our items to see if differences can be observed due to the possibility of applying local accommodation or to idiosyncrasies in the projection pattern of trigger classes.

3.2 A variety of tasks

Two more points must be made before we introduce our experiment design: one is related to the goal of our study, which is to test two hypotheses against each other, and the other is about the experimental task that is best suited to do so. We will avoid giving our experiment an exploratory flavor, and we believe that using a truth-value judgment task is the best way to compare the predictions of the Universal Projection and Strong Kleene hypotheses. We will clarify these two points by briefly reviewing some past literature that has adopted different approaches.

One of the first experimental investigations of presupposition projection out of quantificational environments is due to Chemla (2009). This work, which has a fundamentally exploratory character, focuses on the variability of judgments that may arise if different quantificational elements are compared. The clearest result obtained by Chemla is that participants tend to draw universal inferences about the presupposition of sentences like (25), where the presupposition trigger is embedded under the negative quantifiers *none of these 10 students*. This result is certainly compatible with both the Universal Projection hypothesis and the predictions of Strong Kleene. For our study, we would like to move beyond exploration and find cases that can discriminate between the two approaches.

(25) None of these 10 students knows that their father is going to receive a congratulation letter.
⇝ The father of each of these 10 students is going to receive a congratulation letter.

An important feature of the experimental investigation by Chemla (2009) is the use of an inference task, namely a task where participants are presented with a sentence and asked to judge the validity of drawing a certain inference from that sentence (with a response of *yes* or *no*). Chemla spends some words to justify treating presuppositions as inferences and investigating them through the observation of the inferential process associated with a sentence that carries a presupposition: upon reading a sentence out of the blue, a speaker has to infer a context that satisfies, i.e., entails, its presupposition. This process is also known as global accommodation (see von Fintel 2008 among others). However, more recent literature has started converging on the view that for a given sentence, global accommodation does not always return an inference identical to its presupposition, as long as the context that is accommodated satisfies it (Fox 2013, Mandelkern & Rothschild 2018, Thalmann & Matticchio (to appear)). The process of inferring a context can be affected by other pragmatic and extra-linguistic factors that can lead a speaker to accommodate the most likely context, as opposed to the minimal amount of information necessary to accept a sentence as felicitous. In other words, there may be non-identity between the presuppositions of a sentence and the context that is accommodated in order to make it acceptable, even though the accommodation strictly speaking only needs to cover the presuppositions.

Most notably, the discussion that Chemla (2009) offers of some control items about scalar implicatures points towards this very same direction. Consider the stimuli in (26a), which only differ with respect to the choice of the determiner, *none* or *most*. Neither of these sentences has (26b) as a valid inference on any known approach to natural language literal meaning and scalar implicatures. Nonetheless, participants more easily accepted (26b) after being presented with the *most* version of (26a) that its *none* version. This can result from participants finding the context described by (26b) as more likely or plausible if they know that most of the students failed all of their exams. This behavior shows that factors that are not linguistic or logical can contribute to a participant's drawing an inference, and it is hard or impossible to make participants only draw inferences related to the linguistic domain, not the mention the hypothesis we are investigating.

(26) a. *Stimuli:*
i. Among these 20 students, none failed all of their exams.

ii. Among these 20 students, most failed all of their exams.

b. *Inference to be judged:*
Each of these 20 students failed some of their exams.

Another and more recent experimental investigation was pursued by Wang & Buccola (2025), who focus on projection of presuppositions out of the scope of quantifiers over individuals seeking to compare the predictions of various approaches, including the Universal Projection and Strong Kleene hypotheses. The critical material in their experiment consists of sentences containing presupposition triggers in the scope of three different quantifiers, illustrated in (27), all presented at the same time together with a single common context. Their context manipulated the extent to which the presuppositional requirement of the sentence was met, considering uncontroversial presupposition failure, existential, and universal satisfaction.

(27) zhe san ge nanhai zhong, {meige / youde / meiyou} ren houhoi ziji chidao.
the three CL boy among each some none person regret SELF arrive_late
'Among the three boys, {each/some/none} regretted being late.' (Wang & Buccola 2025: ex. 3)

Participants were tasked with judging the acceptability of Mandarin Chinese context-sentence pairs on a 7-point Likert scale labeled at the endpoints to indicate naturalness. Assuming that the tested sentences were syntactically acceptable, we would in principle want to distinguish sentences by their truth-value, with judgments differentiating between true, false, and undefined items, given that this is where different approaches come apart. In this design, a classically true sentence is expected to fall at the very top of the naturalness scale, but with false and undefined items it is unclear whether we expect them to occupy similar or distinct regions on the scale. If there is an expectation of difference, there is no straightforward way to determine which should be more natural than the other. As far as we can see, no fillers were included to aid interpretation.

The results suggest that participants made a three-way distinction between the quantifiers. All quantifiers were judged as unacceptable in the control context that uncontroversially triggered presupposition failure. Their baseline quantifier *each* was completely natural only in the universal contexts, and unnatural in the other. The quantifier *some* was judged to be natural in both the universal and existential contexts. The quantifier *none*, on the other hand, was more complicated, leading to full naturalness in the universal context, but only partial naturalness, different from all other conditions, in the existential context.

Although Wang & Buccola (2025) argue that their results are compatible with a theory of presupposition projection based on the Strong Kleene logic, the impossibility to distinguish between false and undefined sentences makes their design ill-suited for disentangling the predictions of Strong Kleene from those of a competing theory. Consider for example one of the sentences with universal quantification and the context provided in (28): the Universal Projection and the Strong Kleene hypothesis do not make different predictions about when this sentence is true in the context, which it is not, but rather about whether the sentence is false or undefined. In a design where false and undefined sentences are not clearly distinguished, evaluating the diverging predictions of the two approaches is simply not possible.

(28) CONTEXT: *One boy was late without regretting it, and the other two were not late.*
Each of the three boys regretted being late.

This distinction is instead what our experimental design will focus on and test, but in the realm of attitude reports with the predicate *be certain*. A (ternary) truth-value judgment task is best suited to assess this distinction, together with the introduction of negation as a manipulation, as discussed in Section 4. This choice allows us to directly contrast the Universal Projection and the Strong Kleene hypothesis instead of running an exploratory study and assessing the compatibility of our results with different theories.

3.3 Squeamishness and the variance hypothesis

It is generally assumed that presuppositions are conditions that a context must meet for an utterance to be felicitous. What we neglected to consider in the discussion so far is how interlocutors and experimental

participants react in response to encountering utterances that are asserted in a context that fails to meet its presuppositional requirements. While one option would be to assume that presupposition failure triggers a similar reaction to encountering an ungrammatical sentence, such as strong and immediate intuitions of unacceptability, Strawson (1964) considers a different option. For him, an interlocutor confronted with presupposition failure experiences what he calls squeamishness, a brief moment of hesitation and uncertainty as to how the ongoing conversation is to proceed. In this sense, true and false sentences, which do not elicit squeamishness, should be easier to judge than undefined ones.[9]

In an exploration of whether squeamishness might be detectable experimentally, Thalmann (n.d.) finds support for what he calls the variance hypothesis, the assumption that—at least in truth-value experiments—Strawsonian squeamishness finds its expression in standard deviation effects. In his experiments on the presuppositional import of the definite determiner *the* where participants were tasked with giving truth-value judgments, he finds that conditions exhibiting presupposition failure come with judgment distributions that are more spread out than conditions that participants judge as true or false. In other words, participants unconsciously responded to truth-value gaps with more erratic, less clustered response behavior, which was detectable statistically as inflated standard deviation scores. Finding squeamish rating behavior for the gappy conditions in his experiment, Thalmann argues that inflated indices of dispersion may be used to detect intuitions of presupposition failure in addition to the standard truth-value judgments. Below, we illustrate two sample conditions that lead to squeamish responses in his second experiment: first, a violation of the so-called homogeneity inference associated with definite plurals in (29a), and second, a violation of the uniqueness presupposition for singular definites in (29b).

(29) a. Die Dreiecke stehen links neben dem Quadrat.
the triangles stand left next.to the square
'The triangles stand left of the square.'

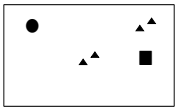

b. Die Wolke steht links neben dem Schild.
the cloud stands left next.to the shield
'The cloud stands left of the shield.'

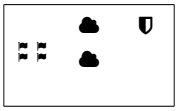

With the focus of our experiment on the truth-value of attitude reports and the projective behavior of presupposition triggers in their scope, we will attempt a replication of Thalmann's findings for presupposition projection from attitudes. That is, we will test whether the impact that presupposition failure had on participants' response behavior generalizes beyond the empirical domain under investigation in Thalmann (n.d.). At the same time, we will be able to use the variance hypothesis to test whether the predictions of the Strong Kleene hypothesis are in line with speakers' intuitions.

4 Experiment: Quantifiers over worlds

4.1 Preliminary considerations on the task

In order to see whether our generalized Strong Kleene hypothesis in (2) aligns with speakers' intuitions or whether the more standard Universal Projection hypothesis for presupposition projection from attitude predicates should remain unchallenged, we designed an experiment that compares their contrasting predictions. Since the differences between the accounts amount to differences in truth-value, we designed the

[9] Being a total logic, Strong Kleene is not itself equipped with a way to make squeamishness the consequence of encountering truth-value gaps, which formally do not exist. This is different from binary logics where presupposition failure results in semantic uninterpretability, which may be identified as the trigger for squeamishness. For Strong Kleene, by contrast, Thalmann (n.d.) argues that the independently needed pragmatic principle called Stalnaker's Bridge could be held responsible for squeamishness. Since Stalnaker's Bridge denies the status of being assertable to utterances that are assigned # rather than T or F, squeamishness in this logic may be considered the result of encountering an unassertable utterance. The source of squeamishness in a Strong-Kleene-based approach would be pragmatic rather than semantic.

(v) **Stalnaker's Bridge** (von Fintel 2008, Fox 2013)
A declarative sentence S is assertable given a context set C only if for all $w \in C$, $[\![S]\!](w) \neq \#$.

experiment as a truth-value judgment task but with a continuous slider, shown in (30), to allow for fine-grained distinctions. For the non-continuous version of this task, and in particular the labels we chose, see Križ & Chemla (2015), who established that this method allows for the detection of truth-value gaps. Thalmann (n.d.) used the same experimental paradigm with a continuous slider and found a replication of the results about homogeneity violations for all the conditions that his experiment shared with Križ & Chemla's, which suggests that the modifications of the method did not affect its original sensitivity to truth-value gaps.

(30) (completely false) (neither completely true nor completely false) (completely true)

To test projection, we embedded presupposition triggers under the attitude predicate *be certain* and manipulated the context in such a way as to tease apart the predictions between the two accounts. Because the detection of presupposition failure relies on an accurate understanding of belief states, this manipulation was carried out redundantly both in terms of verbal contexts and visual materials, which we hoped would make the experimental task of evaluating attitude reports less demanding.

On the assumption that negation is a hole for presuppositions, as in (1), we additionally manipulated the presence of negation in the embedding clause: the judgment of a sentence that is neither true nor false should not be affected by negation. As we discussed in Section 2.2, the prototypical attitude predicate *think* is a neg-raiser, which means that the falsity conditions of a *think* report might not be the Boolean complement of its truth conditions. For this reason, *be certain* was used for all of our items.

Lastly, to see whether speakers differentiate between trigger types both in terms of local accommodation specifically and differences in projection generally (cf. Charlow 2009), we included the soft trigger *stop* and hard trigger *again* as presupposition triggers, see (22) again for the diagnostic we relied on for determination of trigger status.

4.2 Design and materials

The experiment was set up as in (31) and carried out in German via PCIbex (Zehr & Schwarz 2022).[10] All together, the experiment included 48 experimental items: 24 with *again* and 24 with *stop*.

(31) 2×4 design; 48 items

a. NEGATION: without vs. with matrix negation (within items and participants)

b. SCENARIO: `true` vs. `false` vs. `undefined` vs. `critical` (within items and participants)

c. Presupposition TRIGGER as a pseudo-factor: *again* vs. *stop* (between items, 24 each)

Participants were asked to judge the truth of an attitude report in a context (the SCENARIO manipulation) using the slider in (30). For the pseudo-factor, each presupposition trigger had two lexical realizations: *aufhören* and *aufgeben* for 'stop', and *wieder* and *nochmal* for 'again'. Each experimental trial consisted of a SCENARIO portion followed by a critical utterance, which for each item only varied according to the NEGATION and TRIGGER manipulations.

In each item, a character expresses their doxastic state about a past and a present situation regarding another character via direct quotation. In an image presented alongside that quotation designed to redundantly encode that same doxastic state, past and present situations are represented by emojis that were connected by an arrow. The partitioned cartoon thought bubble containing different emoji states represents what the character considers possible.[11] In (32), we list a sample item (translated from German) with *again*; the images corresponding to the different scenarios are shown in Figure 1. In (33), we show an item with *stop*, the accompanying images being shown in Figure 2.[12] The German versions of

[10] A demo version of the experiment can be found at `https://farm.pcibex.net/r/paIsAk/`.

[11] See **pearson2013dese** for the use of cartoon speech/thought bubbles in their experimental investigation of dream reports.

[12] The `true` and `false` scenarios also control for the inference compatible with treating *be certain* like a hole, (11): if participants judge the `true` scenario as true and the `false` scenario as false, the potential additional inference should be independent of the experimental judgments.

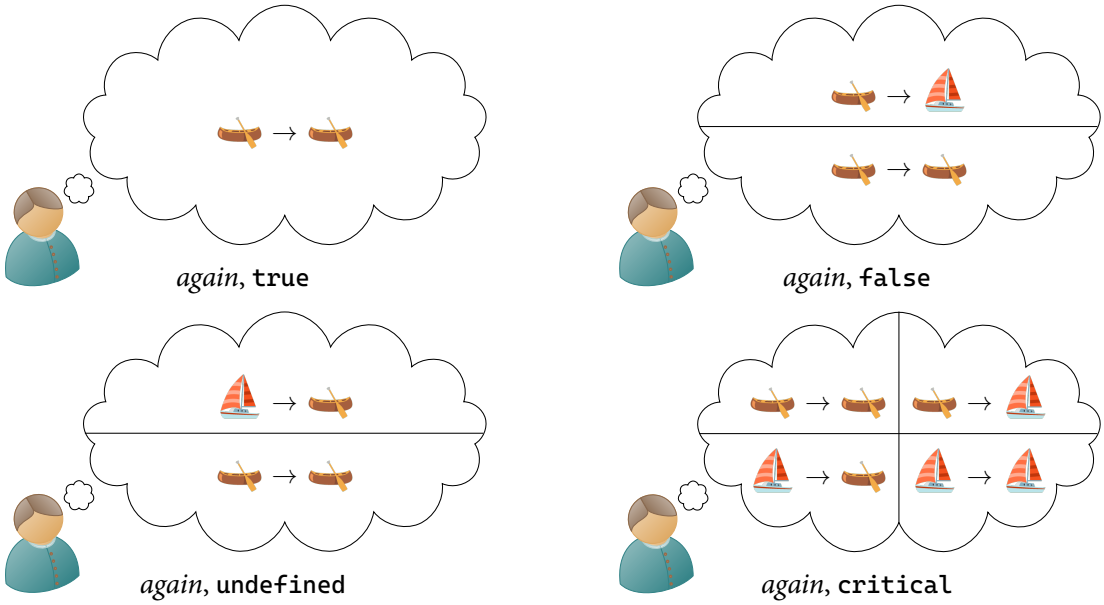

Figure 1: Visual scenarios for "Peter (not) is certain that Jan canoed again".

these two sample items are given in (45) and (46) in Section A. Participants were made aware that the text and the picture conveyed the same information about the context.

(32) Peter is (not) certain that Jan canoed again.

`true` Peter: "I'm certain that Jan canoed last time, and I'm certain that Jan canoed this time."

`false` Peter: "I'm certain that Jan canoed last time, but I have no idea if this time, Jan canoed or not."

`undefined` Peter: "I have no idea if last time, Jan canoed or not, but I am certain that Jan canoed this time."

`critical` Peter: "I have no idea if last time, Jan canoed or not, and I have no idea if this time, Jan canoed or not."

(33) Markus is (not) certain that Sonja stopped drinking wine.

`true` Markus: "I'm certain that Sonja drank wine in the past, and I'm certain that Sonja doesn't drink wine now."

`false` Markus: "I'm certain that Sonja drank wine in the past, but I have no idea if Sonja drinks wine now."

`undefined` Markus: "I have no idea if Sonja drank wine in the past, but I am certain that Sonja doesn't drink wine now."

`critical` Markus: "I have no idea if Sonja drank wine in the past, and I have no idea if Sonja drinks wine now."

Due to the lexical differences between the two kinds of presupposition trigger, the visual material is slightly different for the two. For the *stop* items, the two potential components were a symbol and its crossed-out version, the latter indicating that the situation did not obtain. With the *again* items, we either had a symbol that represented the action mentioned in the item or a different, unrelated symbol. This latter symbol was used to indicate that an alternative to the mentioned action was carried out instead of the mentioned action. Since the controls were rated as expected for both *stop* and *again* items, as we will show, we do not believe that this difference between the two triggers in the pictorial representation of the context had any adverse effects.

As is well known, participants tend to dislike contexts in which an indirect scalar implicature is violated (see **chierchiaetal2012scalarimplicatures**; for experimental evidence, see **bill2016scalar**, Thalmann &

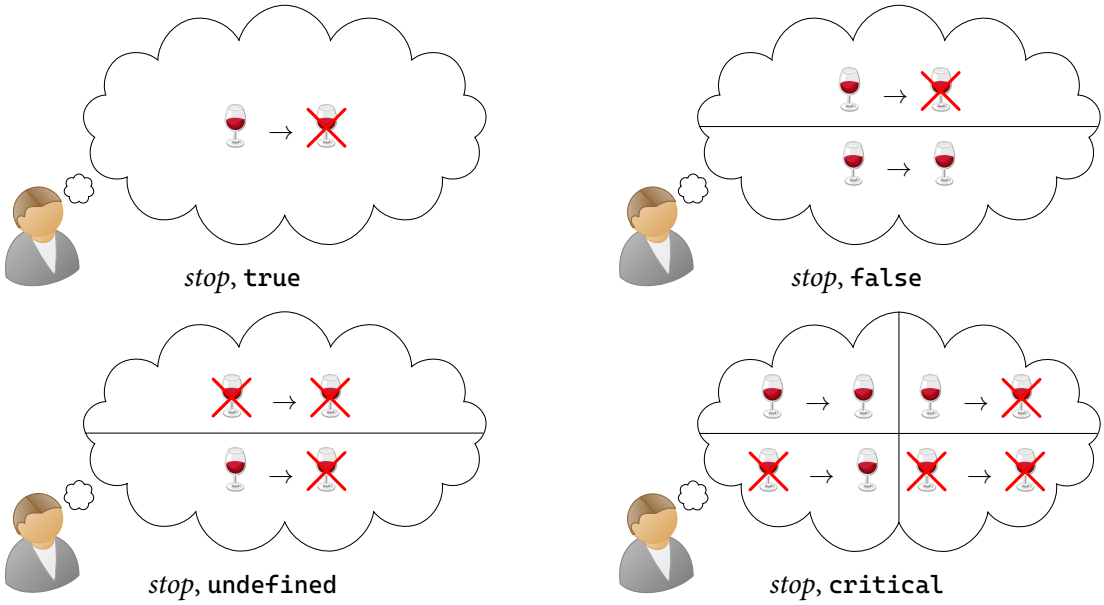

Figure 2: Visual scenarios for "Markus is (not) certain that Sonja stopped drinking wine".

Panizza 2019). Since the negated experimental items all trigger one such implicature, which amounts to a possibility statement along the lines discussed in Section 2.2, we made sure that this implicature is satisfied in all of our scenarios. Consider for example the item in (32) and the pictures in Figure 1: in each of the SCENARIO manipulations, Peter considers it possible that Jan canoed last time and this time. This way, the ratings in the condition with matrix negation are not affected by the presence of this additional inference.

4.3 Participants and procedure

We tested 34 participants (mean age 23.7 ± 6.2, 27 self-identified as female). All participants reported being native speakers of German. Participants received 7 euros in exchange for their participation in the experiment, which lasted around 30 minutes.

At the beginning of the experiment, participants were informed that they would be seeing utterances by a speaker and then what was supposed to be a summary of that utterance. The experimental task consisted of evaluating the extent to which the report was true or not using the scale in (30). Before the beginning of the trials, participants were shown the characters whose beliefs would be summarized and whose faces would be included in the visual materials, reported in Figure 3.

In the next step, participants were familiarized with the meaning of the arrows and the clouds in the visual material with the help of seven warm-up trials, which included instructions on how the visual context corresponded to the verbal context. Just like with the experimental trials, the warm-up items showed the image at the top, above the verbal context, which in turn was displayed above the bold-faced critical item, whose truth-value was to be judged using the scale at the bottom of the screen. Of the warm-up items, none involved presupposition violations, and none included the presupposition triggers participants would see in the actual experiment. An example of a warm-up item is given in (34), translated from German.

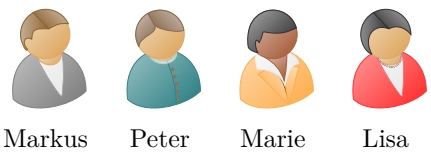

Figure 3: Attitude holder illustration from the instruction portion of the experiment.

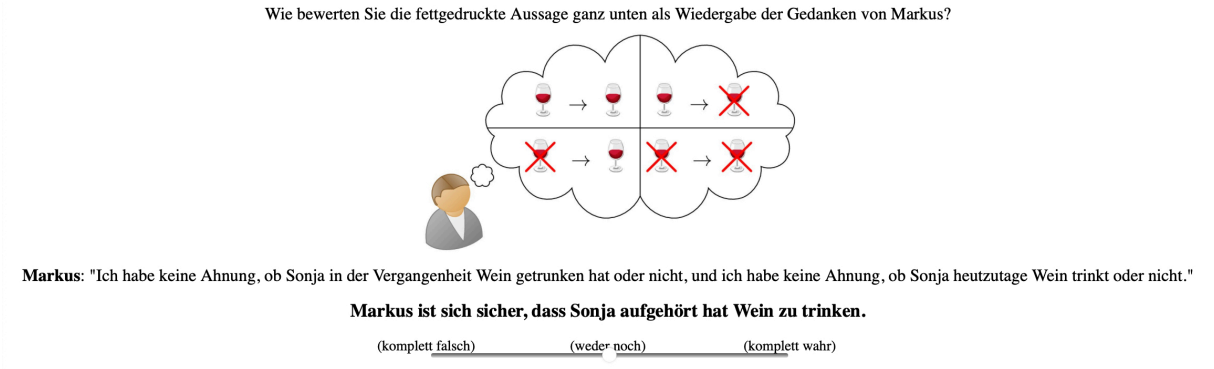

Figure 4: Screenshot of the experimental procedure.

(34) **Susanne:** "I'm certain that Markus didn't do archery in the past, and I'm certain that Markus does archery nowadays."
Susanne is certain that Markus started doing archery.

Upon completing the warm-up trials, participants moved on to the main part of the experiment, where they were shown the critical items in random order. Each trial each also included a repetition of the experimental task at the top of the screen, phrased as in (35). Figure 4 shows what participants saw with the item in (33) in the `critical` scenario in the original German. After half of the experimental trials, participants were given the option for a break.

(35) How do you evaluate the bolded statement at the bottom as a report of [attitude holder]'s thoughts?

Because we wanted to control for the possibility that participants would rate conditions differently depending on whether they relied on the verbal context or the visual material, we included a short question at the end of the experiment asking participants to indicate whether they had relied on the verbal context, the visual material, or both when making their judgments. As a response, 24 said they had relied on the verbal material, 2 on the visual context, and 8 on both. The results were not affected by this factor, so we do not include it in the analysis report.

4.4 Hypotheses

The `true` and `false` scenarios served as control conditions. The `undefined` scenario similarly served as a control, but for an uncontroversially gappy judgment. Contrary to the `true` and `false` scenarios, the `undefined` scenario was expected not to be affected by the presence of matrix negation, leading to a scale-medial overall judgment in both cases.

In the `critical` condition, the doxastic state of the character allows the possibility that the presupposition is not met, and also that the presupposition is met and the embedded proposition is false. If presuppositions project universally out of attitude predicates, as in Heim (1992), the possibility that the presupposition is not met should make the sentence undefined. If, on the other hand, truth-values are calculated via Strong Kleene as in (19), the possibility that the embedded proposition is false should make the sentence false, and its negation should be true. Additionally, if local accommodation is available in this context with soft triggers, as hypothesized by Fox (2013), the results for *stop*, but not for *again*, should display a larger tendency to depart from judgments of intermediate truth, even in the `undefined` scenario. These predictions are summarized in Figure 5.

To the extent that his intuitions generalize to quantifiers over worlds, Charlow (2009) predicts that hard triggers should project universally out of attitude contexts, while soft triggers should not (see again (23)). If this is the case, we expect that the *again* items should be judged as undefined in the `critical` scenario, as expected under the Universal Projection hypothesis, while the *stop* items, without negation, should be judged false in the same scenarios, as in the Strong Kleene or the Local Accommodation hypothesis in Figure 5. Hence, each of the two triggers should display its own projection profile.

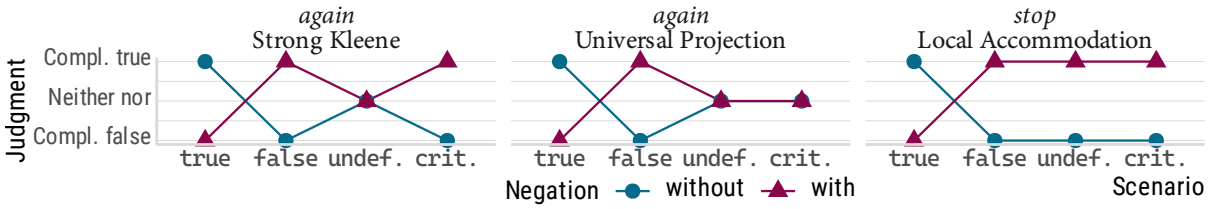

Figure 5: Predictions for the two approaches and the possibility for interpretations with local accommodation for *stop*.

Lastly, the variance hypothesis from Thalmann (n.d.) predicts that conditions judged as involving presupposition failure should elicit larger standard deviations than true or false conditions because of squeamishness (Strawson 1964). For our case, we expect that presupposition failure conditions—those where negation does not impact mean judgments—should come with larger standard deviations compared to classically true or false conditions.

4.5 Data analysis

As we established earlier, one goal of the present experiment is to further explore whether the variance hypothesis holds. But establishing this with the models that are standardly used in experimental semantics is not possible: a run-of-the-mill linear mixed model assumes homoscedasticity, i.e., that all conditions have comparable amounts of spread around the condition means. If we want to investigate whether the variance hypothesis makes the correct predictions for our data, we simply cannot make this assumption. To avoid having to do so, we adopt a so-called distributional model architecture for our data analysis, which allows us to model not only by-condition means but also by-condition standard deviations, which in turn make it possible to evaluate the variance hypothesis for presupposition failure in attitude contexts. It is important to note, however, that the location (means) and scale (standard deviation) portions of these models are independent model parameters and can be modeled using different predictors. Owing to the complexity of the model, we will have to make use of this freedom during the model specification to avoid fit issues.

With this in mind, we fit a Bayesian distributional linear mixed model using the `brm` function from the `brms` package (Bürkner 2021) in R 4.5.1 (Team 2023) with 8 chains with 10,000 iterations each, of which 5,000 were warm-up samples that do not represent actual posterior samples. The dependent variable ranged between 1 and 100 but was scaled to range between −2 (completely false) and 2 (completely) true prior to the analysis to aid interpretability. The model formula is shown in (36). For the location estimates, we included all manipulations from (31) as well as their interactions as fixed effects. The random effects structure was maximal in the sense of Barr et al. (2013). As for modeling the standard deviation portion of the model in (36b), we went with a slightly reduced set of coefficients involving no random effects structure to mitigate fit-issues. As (36b) shows, standard deviations are modeled on the log scale owing to the fact that it is not possible for a condition to display negative dispersion. All predictors were sum-coded.

(36) a. $Y \sim \textsc{trigger} * \textsc{scenario} * \textsc{negation} +$ (location)
$(1 + \textsc{scenario} * \textsc{negation} * \textsc{trigger} \mid \textsc{participant}) +$
$(1 + \textsc{scenario} * \textsc{negation} \mid \textsc{item})$

b. $\log \sigma \sim \textsc{negation} * \textsc{scenario} + \textsc{trigger}$ (scale)

The priors for the model parameters (except those for random effects, which were not modified from the default) were $N(0, x)$, with x ranging between .25 (for the log standard deviation) and 1 (for the location slopes). For the intercept and the slopes in the location sub-model, we also added lower and upper bounds to the priors to reflect the coding of the bounded dependent variable at −2 ('completely false') and 2 ('completely true').

In addition to the model predictions, we will use a Bayes Factor analysis for Regions of Practical Equivalence (ROPE) (for an introduction, see Kruschke 2018) for hypothesis testing. Rather than compare

models with or without a certain parameter, here we generalize this procedure to test for contrasts between parameter settings. A (log) Bayes Factor of 0 indicates that the data provide no preference between the competing hypotheses (e.g., difference vs. equivalence), meaning the posterior odds are equal to the prior odds. A positive (log) Bayes Factor indicates evidence in favor of the alternative hypothesis (e.g., a difference), whereas a negative value indicates evidence in favor of the null or equivalence hypothesis. For example, a log Bayes Factor of 2 would indicate that data are $e^2 = 7.39$ times as likely under the alternative hypothesis.

Because we use weakly informative priors centered on the null (i.e., positing no effect for any slope or intercept parameters), the Bayes Factors are somewhat biased toward the equivalence hypothesis. Bayes Factors evaluate how much the data update prior beliefs: they reflect the extent to which the posterior probabilities differ from the priors. Given our non-difference priors, it is harder to obtain strong evidence in favor of equivalence than to support a difference. With this caveat in mind, we will not rely on Bayes Factors alone. We will also examine posterior marginal effects—model-adjusted condition means of the relevant parameters—which can be interpreted for both location and scale estimates in our distributional model. Since these marginal means are not point estimates but full posterior distributions, comparing them involves assessing overlaps to judge whether the observed difference is meaningful in theoretical terms. To facilitate these decisions, we employ two types of credibility intervals, detailed below. Fortunately, our results are reasonably clear, so these decisions are not ambiguous.

The Bayes Factors were computed using the `bayestestR` package (Makowski, Ben-Shachar & Lüdecke 2019); for the posterior marginal effects we used the `emmeans` package (Lenth 2024).

4.6 Results

We excluded 4.35% of the responses due to the submission times. First, all observations where submission times exceeded 60 s were excluded, and then all observations whose submission times exceeded the new mean ±3 SD.

The descriptive results are shown in Figure 6. Note in particular the almost uniform distribution of the raw ratings along the scale in the `undefined` scenarios. Compared to the other conditions, this highlights the need for a model that does not assume homoscedasticity, like the distributional models we use here. Rather than interpret the raw results directly, we will discuss the posterior marginal means, starting with the location estimates, which are based on the distributional linear mixed model. The coefficients for this model are shown in Figures 7 and 8; the summary is in the appendix in Table 1. Importantly, Figure 7 suggests that the manipulation of the type of presupposition trigger did not have a meaningful effect, given that neither the main effect nor any of the interactions led to reliable rating differences.

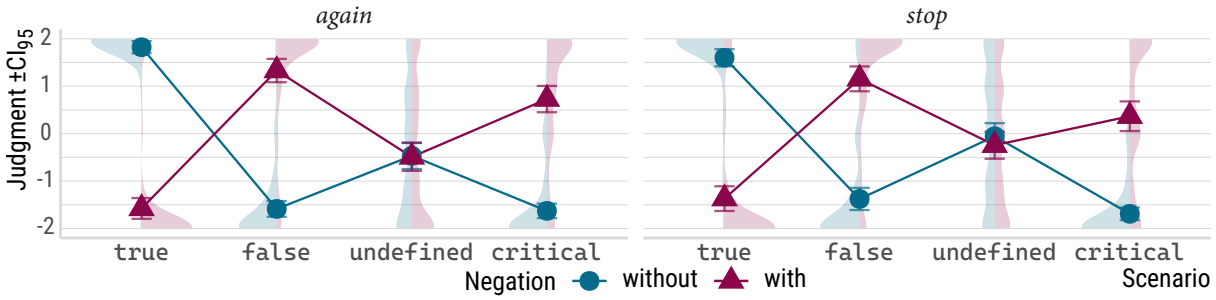

Figure 6: Results for the experiment. The shaded areas indicate the raw rating distributions; the (standard) confidence intervals are calculated on the whole data set and involve no by-participant or by-item clustering.

In Figure 9, we see the posterior marginal estimates for the location portion of the model. The `true` and `false` conditions for both triggers show the expected judgments and inverting effects of negation. For the `undefined` controls, we find scale-medial judgments and, as predicted for a truth-value gap, no effect of negation. By contrast, the `critical` scenario is affected by the negation manipulation, which elicits a

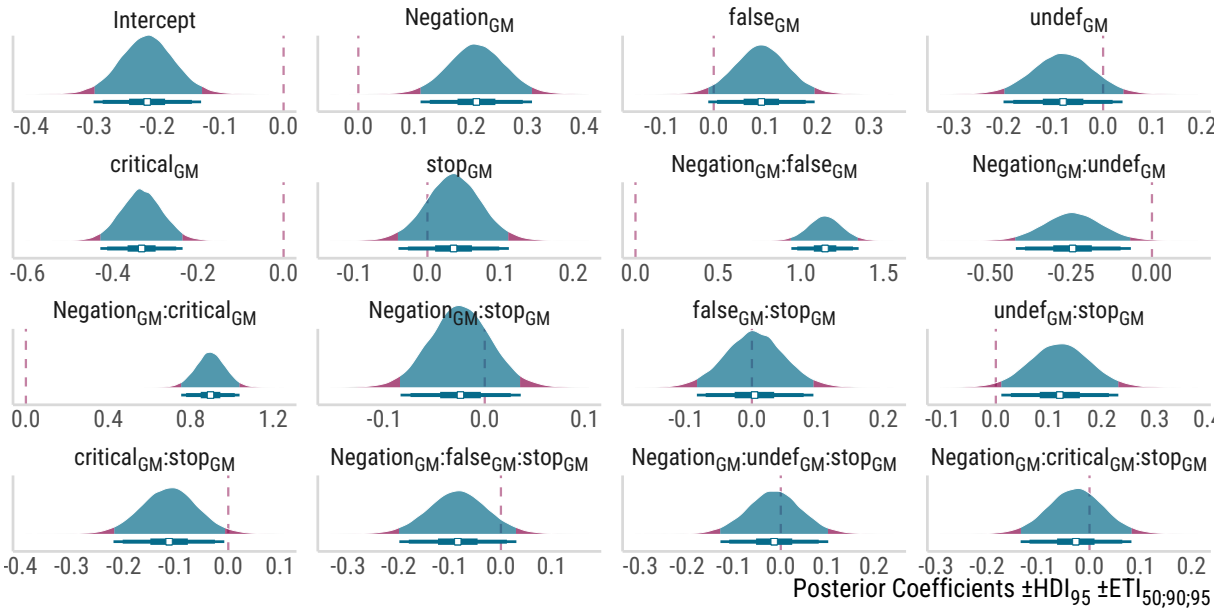

Figure 7: Model coefficients: location parameters. Highest Density intervals (HDIs) show as distributions the most probable values of a parameter, containing a set percentage of the posterior with highest density. Equal-Tailed Intervals (ETIs) are the horizontal bars and show the central mass of the posterior with equal probability in each tail.

change in truth-value judgments that is comparable to the `false` scenarios. The TRIGGER manipulation had little effect, with all of the contrasts we just described present for both the hard and the soft trigger levels of the pseudo-factor.

In Figure 10, we show the outcome of the Bayes Factors for the ROPEs. Importantly, these indicate that NEGATION has an effect for the `true`, `false`, and `critical` scenarios, but not for `undefined`. While the `critical` scenario aligns with the `false` one when no negation is present, the ROPEs do indicate a contrast between the two scenarios with matrix negation.

Turning from the location estimates to the scale parameters, we see in Figure 11 that the `true` and `false` scenarios show small standard deviations, which increase with negation. The `undefined` scenario, by contrast, shows large standard deviations regardless of the presence of matrix negation in the critical items. The `critical` scenario patterns with `true` and `false` without negation, but dispersion increases when negation is present. Figure 12 shows the Bayes Factors for the ROPEs for the scale parameters. These substantiate the results from the posterior marginal estimates.

4.7 General discussion

In our experiment, we tested projection patterns for soft and hard presupposition triggers from the attitude predicate *be certain*, manipulating both the context via belief states and matrix negation. The main goal was to determine what the experimentally founded desiderata for a logic of presupposition projection from attitude predicates are. To achieve this, we employed a continuous truth-value judgment task using the truth-value labels suggested by Križ & Chemla (2015), as first done by Thalmann (n.d.).

As a prerequisite for the following discussion, we need to establish that the experiment was successful at detecting truth-values in attitude reports, which are commonly regarded as a complicated domain for experiments, and which have so far not been investigated experimentally from this perspective. For the present case, all of our six control conditions, the `true`, `false`, and `undefined` scenario for both levels of NEGATION, show the expected results. Where both accounts predict truth/falsity, conditions are judged at the extremes of the continuous truth-value scale and show the predicted judgment inversion when the items are negated. The presupposition failure scenario, by contrast, elicited scale-medial condition means and the absence of an effect from matrix negation; this is expected from an operator that behaves like a hole

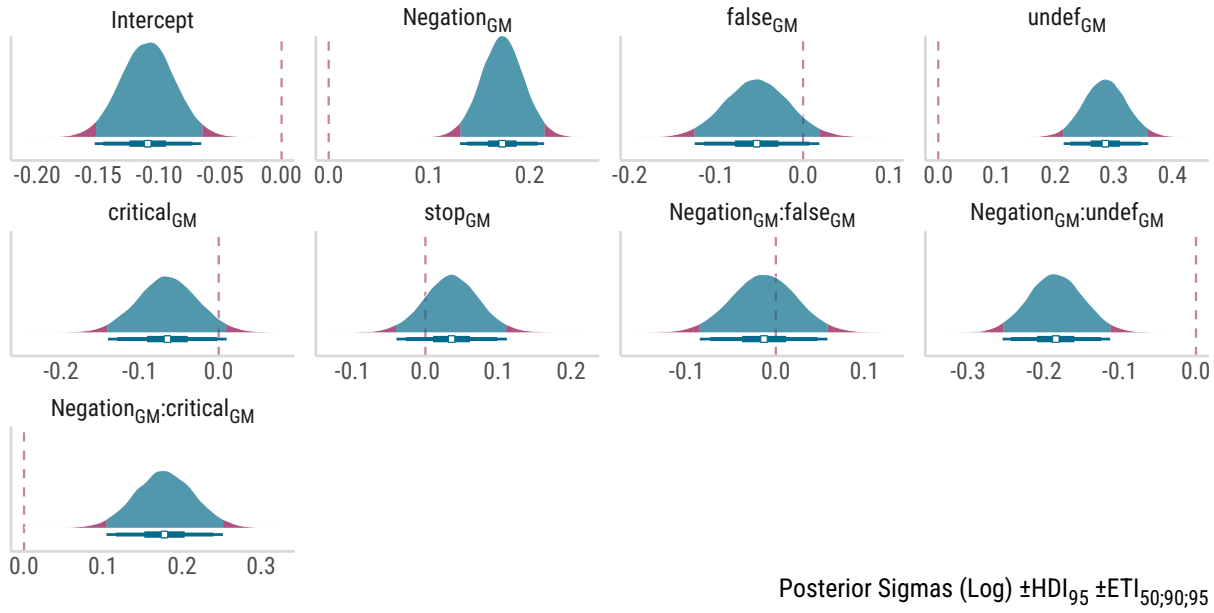

Figure 8: Model coefficients: scale parameters (modeled on the log scale).

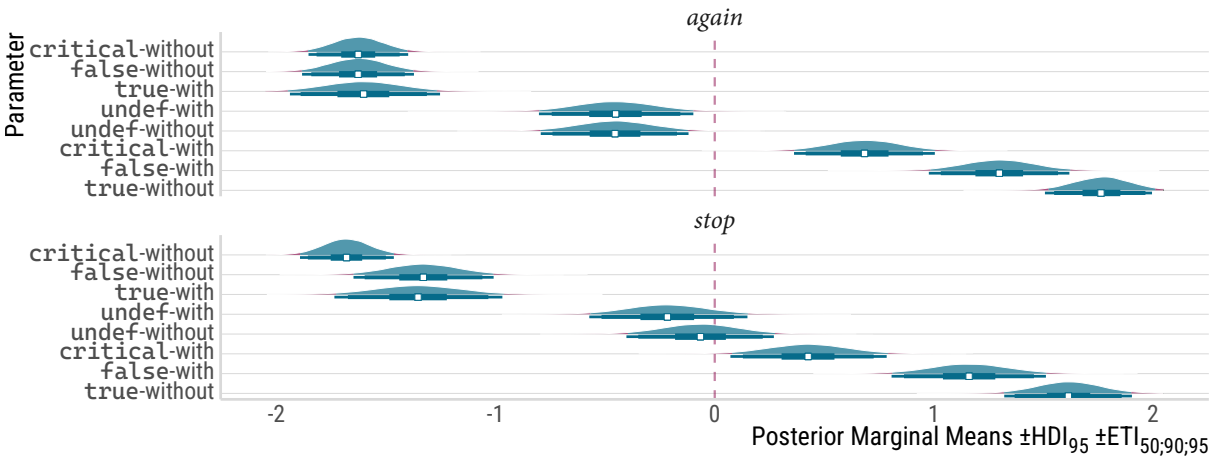

Figure 9: Posterior marginal means for the location parameters.

for presuppositions, as defined in (1). From this, we conclude that our continuous truth-value judgment task with doubly-encoded context manipulations is well-suited for the investigation of presuppositions with attitude reports.

As a further validation of the method for attitude predicates, we will next turn to an expected overlap with an experiment that made use of the same paradigm. Recall that Thalmann (n.d.) finds that intuitions of presupposition failure in this paradigm have measurable effects on participants' response behavior, albeit with the presuppositions of the definite article. The variance hypothesis Thalmann advances posits that the characteristic sensation of Strawsonian squeamishness following presupposition failure is detectable experimentally via increases in standard deviations (Thalmann n.d.). In line with this prediction, we find that our presupposition failure controls cause almost erratic-seeming response behavior in our participants.[13] But not only did the `undefined` scenario elicit large amounts of variance around the condition mean; the addition of negation, which in other scenarios led to increases in variability, had no

[13] That the `undefined` scenario shows a widely dispersed rating distribution is expected under the variance hypothesis, but the distribution we found is an extreme form of excessive dispersion, even under the assumption of squeamishness. The uncertainty is so substantial that the response behavior in this scenario approaches a uniform distribution across the provided scale. Interestingly, in Thalmann (n.d.), the raw ratings for the truth-value gap conditions looked much more normal than they

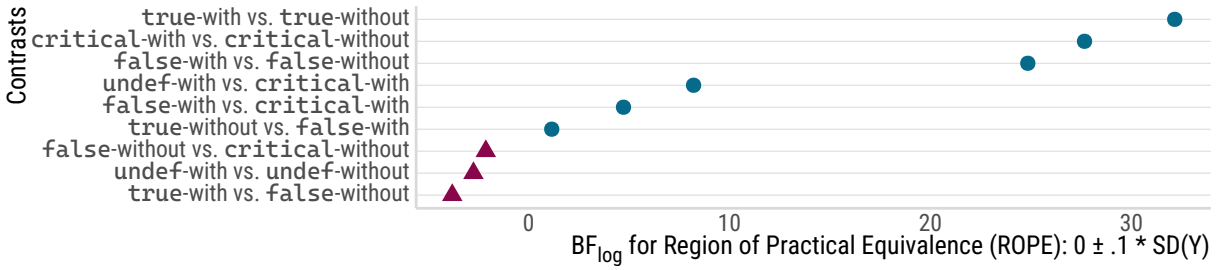

Figure 10: Bayes Factors for the location comparisons. Red triangles indicate that the posterior distribution of the difference between the two conditions is contained in the ROPE, i.e., that the data do not provide evidence for a difference. Blue circles indicate that the posterior distribution of the difference between the two conditions is outside the ROPE, i.e., that the data provide evidence for a difference.

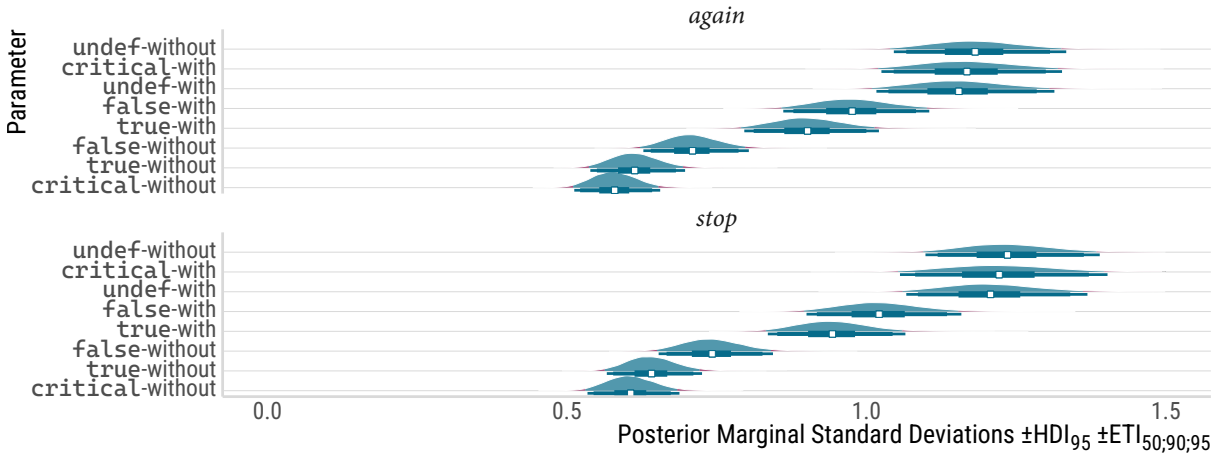

Figure 11: Posterior marginal means for the back-transformed scale parameters.

effect on the standard deviations for the `undefined` scenario. Since our scale estimates mirror the truth-value-driven effect found in Thalmann (n.d.), we have further evidence that the method was successful, in addition to finding additional support for the variance hypothesis and the diagnosis of presuppositions via squeamishness operationalized through standard deviations.

The critical condition was included in our experiment to tease apart the predictions of two approaches to presupposition projection from quantified environments. The standard assumption is that presuppositions project universally from attitude predicates (Heim 1992). By contrast, Strong Kleene logic predicts a kind of asymmetric projection pattern for universal quantifiers. For a sentence to be true, the presupposition must be met for all elements of the domain; for an intuition of falsity, one element in the domain yielding false is sufficient (see (6)). While Strong Kleene has been hypothesized to derive projection for quantifiers over individuals in Fox (2013), it has not been considered as an option for attitude predicates before, to our knowledge. Our `critical` scenario describes a context with two crucial features: on the one hand, there are some doxastic alternatives where the presupposition is not met, and some where it is; on the other hand, the embedded proposition is false in some of the doxastic alternatives where the presupposition is met. The first feature led to a prediction of presupposition failure under the Universal Projection hypothesis, while the second feature led to a prediction of falsity under Strong Kleene. In our results for this scenario, we find not only judgments of falsity but also an effect of matrix negation, neither of which is expected with intuitions of presupposition failure, but which align with the predictions of Strong Kleene. Finally, we find that the presence of matrix negation leads to larger standard deviations

do here. Investigating the extent to which quantitative differences between squeamishness-indicating standard deviations is informative will have be left for future work.

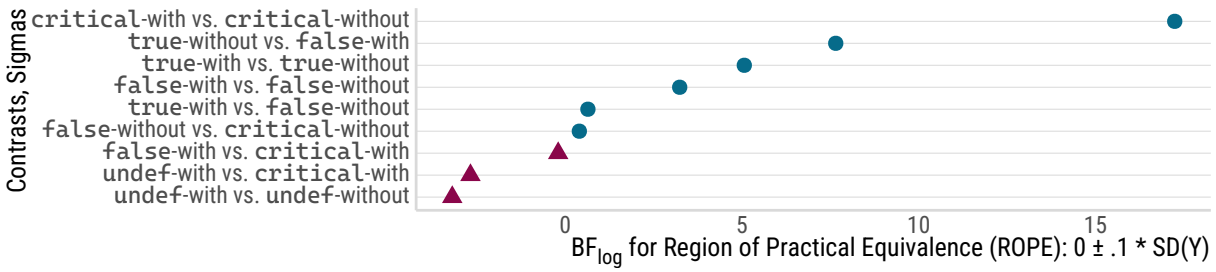

Figure 12: Bayes Factors for the scale comparisons.

in the **critical** scenario, which, in analogy with the false controls, is expected under the Strong Kleene but not under the Universal Projection hypothesis.

However, we also find two unexpected patterns, both in the presence of matrix negation: a contrast between the truth-value judgments for the **critical** and **false** scenarios and additionally, a larger amount of uncertainty in the **critical** scenario. We would like to suggest two ways of explaining these unexpected contrasts between **critical** and **false** scenarios with negation. They inevitably remain at the level of speculation.

The first option concerns the complexity of the SCENARIO manipulation. The **critical** scenario with negation is the most complex condition in the experiment: it involves the highest number of partitions in the character's thought bubble and requires extra computational cost because of negation—which we know inflates standard deviations (see again Figure 8). It is possible that the additional increases in uncertainty we observe in this condition are the result of participants' cognitive load due to this uniquely complex condition.

The second option we consider for the unexplained results in the negated **critical** scenario has to do with differences in information density between the **false** controls and the **critical** scenario. In our design, participants are fully informed about the character's doxastic state, but the report in the **critical** condition with negation might be underinformative, even if true. To see why, consider the item in (37) together with its relevant entailment and indirect scalar implicature.

(37) Peter is not certain that Jan canoed again.

a. *Entails:* Peter considers it possible that Jan did not canoe again.

b. *Implicates:* Peter considers it possible that Jan canoed again.

Taken together, the entailment and the implicature of the sentence address only two out of four of Peter's doxastic possibilities that feature in the **critical** scenario—see Figure 1 for the images and (38) for the verbal contexts. The **false** scenario with negation, which Strong Kleene predicts parity with, shows these two doxastic possibilities and none besides.

(38) **false** Peter: "I'm certain that Jan canoed last time, but I have no idea if this time, Jan canoed or not."

critical Peter: "I have no idea if last time, Jan canoed or not, and I have no idea if this time, Jan canoed or not."

From the perspective of an experimental participant, the items in the **critical** scenario might have felt like worse reports of the character's attitude simply via the comparisons with the **false** scenario.[14] On this view, the contrast between the two scenarios in terms of scale and location is caused by this underinformativity rather than differences in truth-value intuitions.

[14] In a helpful discussion about this general point, Clemens Mayr suggested that there might be a competition with embedding a polar interrogative, exemplified in (vi), which is in general a felicitous option when *certain* is under negation (Mayr 2019). To explore this option, we would need a theory of embedded interrogatives grounded in trivalent logic.

(vi) Peter is not certain whether Jan canoed again.

Despite the fact that the explanations we entertained above do not offer a conclusive answer to the unexpected patterns, we believe that the overall interpretation of the `critical` scenario as not triggering intuitions of presupposition failure still stands. We found the predicted contrast between the negated and the non-negated versions, a difference that is hard to explain on the assumption of presupposition failure and the absence of such a contrast in the `undefined` scenario. Similarly, the effect of negation on the standard deviations is hard to capture with the background that our undefined controls showed no such impact. In sum, we believe that the results for the `critical` scenario are compatible with the Strong Kleene hypothesis, but hard to reconcile with the Universal Projection hypothesis.

Having dealt with the validity of the experimental method and the results for the `critical` scenario, we can now turn to the results for hard and soft presupposition triggers, which should give rise to different projection profiles according to Charlow (2009). This hypothesis, which was made on the basis of introspective judgments on sentences with quantifiers over individuals, would predict that the presupposition of a hard trigger always projects universally.

According to our results, the two presupposition triggers *again* and *stop* are associated with the same projection profile, such that no effect of negation is detected in the `undefined` scenario, but an effect of negation is revealed in the `critical` scenario. As we discussed above, this profile is tough to explain as a result of universal projection, which means that the presupposition of neither of our triggers projects universally. Although it has been suggested by Abrusán (2016: p. 192) that *stop* might not be a soft trigger—an isolated claim, to our knowledge,—*again*, alongside additive particles *too* and *also*, is a prototypical hard presupposition trigger (Abusch 2002, 2010, Romoli 2015). The behavior of *too* and *again* as hard triggers is also experimentally observed by Thalmann & Matticchio ((to appear)). The results of our experiment do not provide any support for a difference between soft and hard triggers along the lines of Charlow's hypothesis: the presupposition of *again* does not projection universally, just like the presupposition triggered by *stop*.

Soft and hard presupposition triggers are also taken to differ with respect to their ability for local accommodation (Abusch 2010, Thalmann & Matticchio (to appear); but see Göbel & Schwarz 2023). In (22), repeated below, we saw that the soft trigger *stop* is acceptable in explicit ignorance contexts, while the hard trigger *again* leads to the intuitions of infelicity indicative of presupposition failure.

(22) a. I don't know whether Peter ever smoked, but if he stopped smoking, Jan will be surprised.
b. # I don't know whether Peter ever smoked, but if he smoked again, Jan will be surprised.

However, given our experiment, the predictions of Strong Kleene for presupposition projection seem to be correct for both soft and hard triggers, without the need for the assumption of local accommodation. After all, we found that there was no tendency for the soft trigger *stop* to be interpreted as merely false in the `undefined` scenario, which would be expected if local accommodation were available. Instead, we found the impact of the TRIGGER manipulation to be minor overall, with no reliable differences between the two levels of the pseudo-factor. These results suggest that participants do not make use of local accommodation, but it is unclear what this unavailability is due to, especially against the background of contrasts like (22). Below, we consider two possibilities, one task-related, the other in connection with potential syntactic and semantic restrictions on local accommodation.

At first blush, locally accommodating a presupposition in our `undefined` items seems pragmatically desirable. On the assumption of Stalnaker's Bridge (see Footnote 9), according to which an utterance is only assertable if it receives a classical truth-value in the context, we expect a certain pressure for an interpretation that does not lead to a violation of that principle. That is, we expect participants to make use of local accommodation to avoid presupposition failure when the option arises, namely when a soft trigger occurs in a context that does not support its presupposition. This is illustrated below for (39) where local accommodation returns the truth-value T instead of # in the presence of presupposition failure, with $\mathcal{A}$ defined as in (21).

(39) Markus is NOT certain that $\mathcal{A}$ [λw . Sonja stopped drinking wine in w]
a. F if Markus is certain that Sonja used to drink wine and stopped.

b. T otherwise.

However, maybe the nature of our task interfered with the pragmatic pressure we normally expect. In a conversation, presupposition failure is generally taken to engender a breakdown. In our trivalent truth-value-judgment task, instead, an alternative perspective on presupposition failure opens up. Since the provided scale features a response option that is non-penalized with respect to the true and false options, it seems possible that participants felt no pressure to locally accommodate soft presuppositions in the `undefined` scenario simply because the rules of the experiment did not offer any advantage to these interpretations. This appears even more plausible on the assumption that local accommodation is a last resort operation. From this perspective, the absence of local accommodation in our items is not surprising, but it suggests the question of whether local accommodation would be detectable in a task with quantifiers over individuals. After all, Fox (2013) argues that local accommodation is necessary to derive the projection patterns for presuppositions embedded under nominal quantifiers. If the analytical option we have been entertaining here is on the right track, we would expect that in a minimal variant of our experiment with quantifiers over individuals, we would likewise find no evidence for local accommodation. We leave exploring this question for future work.

Another, rather more radical speculation about the absence of judgments indicating local accommodation with our items is that local accommodation is simply not available in the syntactic position indicated in (39), right above the complement clause. Our results are compatible with a view where local accommodation is not freely available, as often assumed, but restricted to certain syntactically or semantically determined environments. An important observation in the literature is that local accommodation is reportedly possible in the antecedent, but not in the consequent of a conditional statement (cf. Singh 2014). A minimal contrast that illustrates this observation is shown in (40), where the antecedent allows for local accommodation, while the consequent appears to be much more marked, despite the use of the soft trigger *stop* in both cases.

(40) Context: *Peter is in a pool of people who routinely take part in market research. Currently, the agency is working on an anti-smoking ad.*

a. I don't know whether Peter ever smoked, but if he stopped smoking, the ad will have served its purpose.

b. # I don't know whether Peter ever smoked, but if the ad have served its purpose, he stopped smoking.

It is then natural to expect that local accommodation is not available in unembedded clauses. If it were, it would be hard to explain what could block it from applying in the consequent of a conditional, which is syntactically a main clause. The literature often mentions cases of "local accommodation" under negation (Heim 1983, Bruno 2017), but these cases seem to us to be a separate phenomenon, specific to negative operators and more similar to meta-linguistic negation (Horn 1985, 1989; see also Beaver & Krahmer 2001).[15] Crucially, this type of phenomenon is not sensitive to the cut between soft and hard presupposition triggers, as the example in (41) is perfectly acceptable with contrastive intonation on *again* (see also Jackendoff 1972, Büring 2003, Beaver & Clark 2008). This interpretation of negation, which presumably requires specific contextual support, is evidently not adopted by our participants during the experiment.

(41) Jan didn't canoe AGAIN yesterday; he canoed for the first time yesterday!

One could instead wonder why local accommodation does not apply to the complement of *be certain*, as suggested in (39). Under the hypothesis that attitude predicates and conditionals share a common semantic core as modal universal quantifiers, it would be surprising if local accommodation could apply in

[15] In support of this distinction is the observation by Homer (2008) for French that a soft presupposition that is locally accommodated in a conditional antecedent can stop anti-licensing negative polarity items, whereas the same does not happen under negation.

the complement of *be certain*, given that it cannot apply to the consequent of a conditional. The impossibility to apply local accommodation in the scope of a modal operator is parallel between the two cases, as illustrated in (42).

(42) a. Conditional report, C being the set of worlds compatible with the context:
$*\,(\forall\, w' \in C)\,\big([\lambda w\,.\,\textit{if...}\ \textit{in}\ w](w') \rightarrow \mathcal{A}\,[\lambda w\,.\,\textit{then...}\ \textit{in}\ w](w')\big)$

b. Attitude report, $\textsc{dox}_{\mathrm{M}}^{w_0}$ being the set of Markus's doxastic alternatives in w_0:
$*\,\forall\, w'\,\big([\lambda w\,.\,w \in \textsc{dox}_{\mathrm{M}}^{w_0}](w') \rightarrow \mathcal{A}\,[\lambda w\,.\,\textit{that ...}\ \textit{in}\ w]\big)$

Although it remains at the level of a speculation, this option could add to our understanding of distributional properties of local accommodation. The facts revealed by our experiment remain the same, though: the results of our task can be explained by a Strong Kleene hypothesis for presupposition projection without the aid of a presupposition weakening mechanism like local accommodation.

5 Conclusions

We observed that empirical research is needed on the topic of presupposition projection out of the scope of quantifiers. In particular, the case of attitude predicates has not been explored in previous research into the topic. Instead, the standard response has been to assume on the basis of introspective data alone that presuppositions project universally from the complement of attitude predicates (Heim 1992). In the pursuit of a common logic driving projective behavior from quantified environments, we conjectured that Strong Kleene logic, considered by George (2010) and Fox (2013) for projection from quantifiers over individuals, could be generalized to explain presupposition projection for modal quantifiers as well.

Our experimental results for the attitude predicate *be certain* conform to the prediction of Strong Kleene, much like what has been argued for quantifiers like *every*—the one difference being that in our data, we found no reason to invoke additional mechanisms like local accommodation (contra what is assumed in Fox 2013). These results suggest that it may be possible to unify the mechanisms assumed to be responsible for projection across different kinds of quantifiers, along the lines of our original hypothesis in (2). In addition, our results conflict with the idea that soft and hard triggers have distinct projection profiles as entertained in Charlow (2009), at least in the domain of quantifiers over individuals. Still, more research is needed into other attitude predicates, as well as non-doxastic and non-universal modal quantifiers, to see whether what looks like a promising candidate logic for presupposition projection now can be maintained in the face of a more complete empirical, experimentally grounded picture.

A helpful example is the case of neg-raising predicates, like *think*, discussed in Section 2.2. On top of a fundamental Strong-Kleene-like logic for presupposition projection, additional semantic or pragmatic complexity should be taken into consideration. We can hypothesize that the truth conditions of a *think*-report must take neg-raising into account. The results can be schematized as in (43), the intuitions being illustrated in (44).

(43) **Projection out of a neg-raiser:**
The truth-value of $\textsc{think}_{\mathcal{W}}\ p$ is: *($\mathcal{W}$ being the set of doxastic alternatives)*

a. T if $p(w) = \mathrm{T}$ for all $w \in \mathcal{W}$;

b. F if $p(w) = \mathrm{F}$ for all $w \in \mathcal{W}$;

c. # otherwise (if there is $w_1, w_2 \in \mathcal{W}$ such that $p(w_1) = \mathrm{T}$ and $p(w_2) = \mathrm{F}$ (excluded middle violation), or if there is a $w \in \mathcal{W}$ such that $p(w) = \#$)

(44) Aditi thinks that Skye stopped smoking.

a. T if Aditi thinks that Skye used to smoke and stopped

b. F if Aditi thinks that Skye used to smoke and didn't stop

c. # if Aditi thinks that Skye might have used to smoke and stopped but also that Skye might have used to smoke and not stopped; or if Aditi thinks that Skye might have not used to smoke

These predictions, specific to *think*, match the intuitions reported by Heim (1992) and the Universal Projection hypothesis, but they are explained as the result of neg-raising. They do not undermine the hypothesis that Strong Kleene is a good candidate to model the natural language logic for presuppositions and, more in general, the hypothesis that a logic that manipulates non-classical truth-values is needed to model presupposition projection.

Finally, instead of treating the dispersion of the experimental data as a nuisance variable indicative of mere noise or trying to explain it as the result of inter-speaker variation, we adopted the variance hypothesis for presupposition failure from Thalmann (n.d.). Finding that participants' responses are more spread out in scenarios with presupposition failure, we argued that the dispersion of responses is a reflection of the uncertainty, or squeamishness, that unmet presuppositional requirements engender in speakers. In this way, we were able to use what is often discarded as noise to diagnose intuitions of penalized truth-values. Future work will hopefully determine whether inflated (and negation-invariant) standard deviations are a reliable indicator of presupposition failure in other experimental paradigms, too.

References

Abrusán, Márta. 2016. Presupposition cancellation: explaining the 'soft–hard' trigger distinction. *Natural Language Semantics* 24(2). 165–202. https://doi.org/10.1007/s11050-016-9122-7.

Abusch, Dorit. 2002. Lexical alternatives as a source of pragmatic presuppositions. In Brendan Jackson (ed.), *Proceedings of Semantics and Linguistic Theory (SALT) 12*, 1–19. https://doi.org/10.3765/salt.v12i0.2867.

Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1). 37–80. https://doi.org/10.1093/jos/ffp009.

Barr, Dale J., Roger Levy, Christoph Scheepers & Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: keep it maximal. *Journal of Memory and Language* 68(3). 255–278. https://doi.org/10.1016/j.jml.2012.11.001.

Beaver, David I. 2001. *Presupposition and assertion in dynamic semantics*. Stanford: CSLI publications.

Beaver, David I. & Brady Z. Clark. 2008. *Sense and sensitivity: how focus determines meaning*. Malden, MA: Wiley-Blackwell.

Beaver, David I. & Emiel Krahmer. 2001. A partial account of presupposition projection. *Journal of Logic, Language and Information* 10(2). 147–182. https://doi.org/10.1023/a:1008371413822.

Bruno, Chris. 2017. Contrastive negation and the theory of alternatives. In *Proceedings of SALT 27*, 59–74.

Büring, Daniel. 2003. On D-trees, beans, and B-accents. *Linguistics and Philosophy* 26(5). 511–545. https://doi.org/10.1023/a:1025887707652.

Bürkner, Paul-Christian. 2021. Bayesian item response modeling in R with brms and stan. *Journal of Statistical Software* 100(5). 1–54. https://doi.org/10.18637/jss.v100.i05.

Charlow, Simon. 2009. "Strong" predicative presuppositional objects. In *Proceedings of ESSLLI 2009*.

Chemla, Emmanuel. 2009. Presuppositions of quantified sentences: experimental data. *Natural Language Semantics* 17(4). 299–340. https://doi.org/10.1007/s11050-009-9043-9.

Chen, Yuqiu, Maik Thalmann & Mailin Antomo. 2022. Presupposition triggers and (not-)at-issueness. Insights from language acquisition into the soft-hard distinction. *Journal of Pragmatics* 199. 21–46. https://doi.org/10.1016/j.pragma.2022.06.014.

Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Adriana Belletti (ed.), *Structures and beyond*, 39–103. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780195171976.003.0003.

Chierchia, Gennaro. 2013. *Logic in grammar: polarity, free choice, and intervention*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199697977.001.0001.

Clark, Herbert H. 1973. The language-as-fixed-effect fallacy: a critique of language statistics in psychological research. *Journal of verbal learning and verbal behavior* 12(4). 335–359. https://doi.org/10.1016/s0022-5371(73)80014-3.

von Fintel, Kai. 2008. What is presupposition accommodation, again? *Philosophical Perspectives* 22(1). 137–170. https://doi.org/10.1111/j.1520-8583.2008.00144.x.

Fox, Danny. 2013. Presupposition projection from quantificational sentences: Trivalence, local accommodation, and presupposition strengthening. In Ivano Caponigro & Carlo Cecchetto (eds.), *From grammar to meaning*, 201–232. Cambridge University Press. https://doi.org/10.1017/cbo9781139519328.011.

Gajewski, Jon R. 2007. Neg-raising and polarity. *Linguistics & Philosophy* 30(3). 289–328. https://doi.org/10.1007/s10988-007-9020-z.

George, B. 2010. A new case for an old logic: Reviving Strong Kleene approaches to presupposition projection. Unpublished ms. UCLA.

Geurts, Bart. 1999. *Presuppositions and pronouns*. Amsterdam: Elsevier.

Göbel, Alexander & Bernhard Schwarz. 2023. Comparing global and local accommodation: Rating and response time data. In Tyler Knowlton, Florian Schwarz & Anna Papafragou (eds.), *Proceedings of experiments in linguistic meaning (elm) 2*, 95–103. https://doi.org/10.3765/elm.2.5400.

Heim, Irene. 1983. On the projection problem for presuppositions. In Michael Barlow, Daniel P. Flickinger & Michael T. Wescoat (eds.), *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*, 114–125.

Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* (9). 183–221. https://doi.org/10.1093/jos/9.3.183.

Hintikka, Jaakko. 1969. Semantics for propositional attitudes. In *Models for modalities. selected essays*, 87–111. Dordrecht: Reidel. https://doi.org/10.1515/9783111546216-009.

Homer, Vincent. 2008. Disruption of NPI licensing: the case of presuppositions. *Proceedings of Semantics and Linguistic Theory (SALT) 18*. 429–446. https://doi.org/10.3765/salt.v18i0.2483.

Horn, Laurence R. 1985. Metalinguistic negation and pragmatic ambiguity. *Language* 61(1). 121–174. https://doi.org/10.2307/413423.

Horn, Laurence R. 1989. *A natural history of negation*. Chicago: Chicago University Press. https://doi.org/10.2307/2275088.

Jackendoff, Ray S. 1972. *Semantic interpretation in generative grammar*. MIT Press. https://doi.org/10.2307/413161.

Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4(2). 169–193.

Karttunen, Lauri & Stanley Peters. 1979. Conventional implicature. In David A. Dineen & Choon-Kyu Oh (eds.), *Presupposition*, 1–56. New York: Academic Press.

Koev, Todor. 2019. Strong beliefs, weak commitments. In M. Teresa Espinal, Elena Castroviejo, Manuel Leonetti, Louise McNally & Cristina Real-Puigdollers (eds.), *Proceedings of sinn und bedeutung 23*, vol. 2, 1–18.

Kratzer, Angelika. 2006. Decomposing attitude verbs. Talk given at the Hebrew University of Jerusalem.

Križ, Manuel & Emmanuel Chemla. 2015. Two methods to find truth-value gaps and their application to the projection problem of homogeneity. *Natural Language Semantics* 23(3). 205–248. https://doi.org/10.1007/s11050-015-9114-z.

Kruschke, John K. 2018. Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science* 1(2). 270–280. https://doi.org/10.1177/2515245918771304.

Lassiter, Daniel. 2017. *Graded modality*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198701347.001.0001.

Lenth, Russell V. 2024. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. R package version 1.10.0.

Makowski, Dominique, Mattan S. Ben-Shachar & Daniel Lüdecke. 2019. bayestestR: Describing effects and their uncertainty, existence and significance within the Bayesian framework. *Journal of Open Source Software* 4(40). 1541. https://doi.org/10.21105/joss.01541.

Mandelkern, Matthew. 2024. *Bounded meaning: the dynamics of interpretation*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780192870049.001.0001.

Mandelkern, Matthew & Daniel Rothschild. 2018. Independence day? *Journal of Semantics* 36(2). 193–210. https://doi.org/10.1093/jos/ffy013.
Mayr, Clemens. 2019. Triviality and interrogative embedding: context sensitivity, factivity, and neg-raising. *Natural Language Semantics* 27(3). 227–278. https://doi.org/10.1007/s11050-019-09153-8.
Moulton, Keir. 2009. *Natural selection and the syntax of clausal complementation*. UMass Amherst dissertation.
Moulton, Keir. 2015. Cps: Copies and compositionality. *Linguistic Inquiry* 46(2). 305–342. https://doi.org/10.1162/ling_a_00183.
Romoli, Jacopo. 2015. The presuppositions of soft triggers are obligatory scalar implicatures. *Journal of Semantics* 32(2). 173–219. https://doi.org/10.1093/jos/fft017.
Santorio, Paolo & Jacopo Romoli. 2017. Probability and implicatures: A unified account of the scalar effects of disjunction under modals. *Semantics and Pragmatics* 10(13). 1–61. https://doi.org/10.3765/sp.10.13.
Schlenker, Philippe. 2009. Local contexts. *Semantics & Pragmatics* 2. https://doi.org/10.3765/sp.2.3.
Singh, Raj. 2014. A note on local accommodation. In Luka Crnič & Uli Sauerland (eds.), *The art and craft of semantics: A Festschrift for Irene Heim vol. 2*, 155–173. Cambridge, MA: MITWPL.
Strawson, P. F. 1964. Identifying reference and truth-values. *Theoria* 30(2). 96–118. https://doi.org/10.1111/j.1755-2567.1964.tb00404.x.
Sudo, Yasutada. 2014. Presupposition satisfaction in attitude contexts and modal subordination. In Luka Crnič & Uli Sauerland (eds.), *The art and craft of semantics: A Festschrift for Irene Heim*, vol. 2, 175–199. Cambridge, MA: MIT Press.
Sudo, Yasutada, Jacopo Romoli, Martin Hackl & Danny Fox. 2012. Presupposition projection out of quantified sentences: Strengthening, local accommodation and inter-speaker variation. In Maria Aloni, Floris Roelofsen, Galit W. Sassoon, Katrin Schulz & Matthijs Westera (eds.), *Logic, language and meaning. 18th amsterdam colloquium*, 210–219. Berlin: Springer. https://doi.org/10.1007/978-3-642-31482-7_22.
Team, R Core. 2023. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria.
Thalmann, Maik. N.d. Not so squeamish, Your majesty: The presuppositions of singular definites between existence and uniqueness. Ms. University of Göttingen.
Thalmann, Maik & Andrea Matticchio. (to appear). How to be mistaken and still happy: Belief-Relative presuppositions and factivity illusions. In *Proceedings of the 48th penn linguistics conference*.
Thalmann, Maik & Daniele Panizza. 2019. Present to the eye, away from the mind. Dissociating online comprehension and offline judgments of indirect scalar inferences. In Megan M. Brown & Brady Dailey (eds.), *Proceedings of 43rd annual Boston University Conference on Language Development*, 667–678. Somerville, MA: Cascadilla Press.
Uegaki, Wataru. 2020. The existential/uniqueness presupposition of *wh*-complements projects from the answers. *Linguistics and Philosophy* 44(4). 911–951. https://doi.org/10.1007/s10988-020-09309-4.
Wang & Brian Buccola. 2025. Weakening the trivalent semantics of quantifiers: Evidence from Mandarin Chinese. In *West coast conference on formal linguistics (WCCFL) 41*, 600–607.
Zehr, Jérémy & Florian Schwarz. 2022. PennController for internet based experiments (IBEX). https://doi.org/10.17605/OSF.IO/MD832.

A German sample items

(45) Peter ist sich sicher, dass Jan nochmal Kanu gefahren ist.
Peter is REFL certain that Jan again canoe driven is

a. Ich bin mir sicher, dass Jan letztes Mal Kanu gefahren ist, und ich bin mir sicher, dass
I am REFL certain that Jan last time canoe driven is and I am REFL certain that
Jan dieses Mal Kanu gefahren ist.
Jan this time canoe driven is

b. Ich bin mir sicher, dass Jan letztes Mal Kanu gefahren ist, aber ich habe keine Ahnung,
I am REFL certain that Jan last time canoe driven is but I have no clue
ob Jan dieses Mal Kanu gefahren ist oder nicht.
whether Jan this time canoe driven is or not

c. Ich habe keine Ahnung, ob Jan letztes Mal Kanu gefahren ist oder nicht, aber ich bin
I have no clue whether Jan last time canoe driven is or not but I am
mir sicher, dass Jan dieses Mal Kanu gefahren ist.
REFL certain that Jan this time canoe driven is

d. Ich habe keine Ahnung, ob Jan letztes Mal Kanu gefahren ist oder nicht, und ich habe
I have no clue whether Jan last time canoe drive is or not and I have
keine Ahnung, ob Jan dieses Mal Kanu gefahren ist oder nicht.
no clue whether Jan this time canoe driven is or not

(46) Markus ist sich sicher, dass Sonja aufgehört hat Wein zu trinken.
Markus is REFL certain that Sonja stopped has wine to drink

a. Ich bin mir sicher, dass Sonja in der Vergangenheit Wein getrunken hat, und ich bin mir
I am REFL certain that Sonja in the past wine drunk has and I am REFL
sicher, dass Sonja heutzutage keinen Wein trinkt.
certain that Sonja nowadays no wine drinks

b. Ich bin mir sicher, dass Sonja in der Vergangenheit Wein getrunken hat, aber ich habe
I am REFL certain that Sonja in the past wine drunk has but I have
keine Ahnung, ob Sonja heutzutage Wein trinkt oder nicht.
no clue whether Sonja nowadays wine drinks or not

c. Ich habe keine Ahnung, ob Sonja in der Vergangenheit Wein getrunken hat oder
I have no clue whether Sonja in the past wine drunk has or
nicht, aber ich bin mir sicher, dass Sonja heutzutage keinen Wein trinkt.
not but I am REFL certain that Sonja nowadays no wine drinks

d. Ich habe keine Ahnung, ob Sonja in der Vergangenheit Wein getrunken hat oder
I have no clue whether Sonja in the past wine drunk has or
nicht, und ich habe keine Ahnung, ob Sonja heutzutage Wein trinkt oder nicht.
not and I have no clue whether Sonja nowadays wine drinks or not

B In-depth model summary

Parameter	Median (±ETI)	Mode (±HDI)	$\hat{R}$	Bulk ESS	Tail ESS
Intercept	-0.22 (-0.3, -0.13)	-0.21 (-0.3, -0.13)	1	18965	26013
sigma_Intercept	-0.11 (-0.15, -0.07)	-0.11 (-0.15, -0.06)	1	17678	25792
NEGATION	0.21 (0.11, 0.31)	0.2 (0.11, 0.31)	1	20912	23437
`false`	0.09 (-0.01, 0.19)	0.09 (-0.01, 0.19)	1	31200	29038
`undef`	-0.08 (-0.2, 0.04)	-0.09 (-0.2, 0.04)	1	32298	30232
`critical`	-0.33 (-0.43, -0.24)	-0.34 (-0.43, -0.24)	1	37265	31722
stop	0.04 (-0.04, 0.11)	0.04 (-0.04, 0.11)	1	26188	26990
NEGATION:`false`	1.15 (0.94, 1.35)	1.14 (0.94, 1.34)	1	20244	24526
NEGATION:`undef`	-0.25 (-0.42, -0.07)	-0.25 (-0.42, -0.07)	1	21952	27166
NEGATION:`critical`	0.9 (0.75, 1.04)	0.9 (0.75, 1.04)	1	26200	29092
NEGATION:*stop*	-0.02 (-0.08, 0.04)	-0.03 (-0.08, 0.04)	1	39426	30180
`false`:*stop*	0 (-0.08, 0.09)	0 (-0.08, 0.09)	1	43608	30566
`undef`:*stop*	0.12 (0.01, 0.23)	0.13 (0.01, 0.23)	1	39388	29582
`critical`:*stop*	-0.11 (-0.22, -0.01)	-0.11 (-0.22, -0.01)	1	32643	29595
NEGATION:`false`:*stop*	-0.08 (-0.2, 0.03)	-0.08 (-0.2, 0.03)	1	27816	29044
NEGATION:`undef`:*stop*	-0.01 (-0.13, 0.1)	-0.02 (-0.13, 0.1)	1	36678	29406
NEGATION:`critical`:*stop*	-0.03 (-0.13, 0.08)	-0.02 (-0.13, 0.08)	1	31857	29861
sigma_NEGATION	0.17 (0.13, 0.21)	0.17 (0.13, 0.21)	1	36975	31346
sigma_`false`	-0.05 (-0.13, 0.02)	-0.05 (-0.13, 0.02)	1	28481	30984
sigma_`undef`	0.28 (0.21, 0.36)	0.28 (0.22, 0.36)	1	24991	29953
sigma_`critical`	-0.07 (-0.14, 0.01)	-0.07 (-0.14, 0.01)	1	19420	29475
sigma_*stop*	0.02 (-0.02, 0.06)	0.02 (-0.02, 0.06)	1	38132	31818
sigma_NEGATION:`false`	-0.01 (-0.09, 0.06)	-0.02 (-0.09, 0.06)	1	33779	30658
sigma_NEGATION:`undef`	-0.18 (-0.25, -0.11)	-0.19 (-0.26, -0.11)	1	33522	31284
sigma_NEGATION:`critical`	0.18 (0.1, 0.25)	0.18 (0.1, 0.25)	1	28436	31138

Table 1: Model summary for the so-called fixed effects. Note that the standard deviation parameters, prefixed with *sigma_*, are on the log scale.