

# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

### Title

Negative strengthening and speaker adaptation with vague quantifiers

### Permalink

<https://escholarship.org/uc/item/9f68x8vs>

### Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

### Authors

Pecsok, Emily

Aparicio, Helena

### Publication Date

2026

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Negative strengthening and speaker adaptation with vague quantifiers

Emily Pecsok (epecsok@sas.upenn.edu)<sup>1</sup> & Helena Aparicio<sup>2</sup>

<sup>1</sup>Department of Linguistics, University of Pennsylvania

<sup>2</sup>Department of Linguistics, Cornell University

## Abstract

The interpretation of negated vague predicates, whose meanings depend on underspecified thresholds, is variable. Some positive relative adjectives exhibit asymmetric negative strengthening (*not large*  $\rightsquigarrow$  *small*), whereas negative adjectives generally do not (*not small*  $\not\rightsquigarrow$  *large*). However, the behavior of negated vague quantifiers (e.g., *many*, *few*) remains understudied. Previous work also shows that listeners adapt to speakers' thresholds for vague predicates through non-negated assertions, but it is unknown whether such adaptation extends to negated utterances. We investigate whether listeners adapt to speaker thresholds conveyed via negated assertions and whether negated vague quantifiers are interpreted as symmetrically weaker than their antonyms. We find that listeners track speaker thresholds despite negation and interpret negated quantifiers as symmetrically weaker than their antonyms, challenging asymmetric accounts of negative strengthening.

**Keywords:** vague quantifiers; negation; negative strengthening; metalinguistic disagreement; speaker adaptation

## Introduction

Despite its seemingly simple semantic contribution, negation can make determining an utterance's meaning considerably more difficult and complex. Negated propositions have the potential to be significantly harder to process than their non-negated counterparts (Clark & Chase, 1972; Wason, 1961), and they can trigger complex pragmatic inferences that are heavily dependent on the conversational context (Nieuwland & Kuperberg, 2008; Tian et al., 2010). This complexity is especially evident when negation targets a vague predicate—such as vague quantifiers (e.g., *many*, *few*) and relative adjectives (e.g., *large*, *small*)—whose interpretations rely on context- and speaker-dependent thresholds that determine whether the predicate is applicable. Because these thresholds vary across individuals, two speakers may disagree on the cutoff for a vague predicate's applicability, giving rise to 'metalinguistic disagreements,' where interlocutors agree on the facts of the world (e.g., an object's size) but disagree on the linguistic label that best describes those facts (e.g., if the object is *large* or *not large*).

To resolve threshold ambiguity and coordinate on the meaning of vague predicates in conversation, listeners can adapt their threshold interpretation to that displayed by the speaker. Prior work has shown that listeners successfully track speaker thresholds, including those expressed within metalinguistic disagreements (Pecsok & Aparicio, in prep; Schuster & De-

gen, 2020; Yildirim et al., 2016). However, all previous work on speaker adaptation in the domain of meaning has focused exclusively on non-negated assertions. Given that negated assertions have the potential to impose a higher cognitive load and incur greater processing costs than their positive counterparts, it remains an open question whether participants can successfully track speaker thresholds when expressed via negated assertions.

The additional potential difficulty with adapting to negated assertions raises the question of how negated vague quantifiers are interpreted. This question has primarily been addressed in the domain of relative adjectives (e.g., *large*, *small*), another type of vague predicate. Adjectives are analyzed as existing along a scale measuring some property that positive adjectives have more of and negative adjectives have less of (e.g., *large* is higher than *small* on the *size* scale, making *large* positive and *small* negative). Positive adjectives also tend to be perceived as more desirable than their negative counterparts. Generally speaking, the meaning of a negative relative adjective, such as *small*, when negated, falls into the extensional gap, a region on the scale where neither adjective is applicable (*neither large nor small*). However, the interpretation of positive adjectives when negated is debated. Previous work argues that negated positive adjectives (e.g., *not large*) strengthen to their antonyms (*not large*  $\rightsquigarrow$  *small*) in a process called asymmetric negative strengthening (Horn, 1989). Other prior work argues for symmetric negative strengthening, where both negated positive and negative forms strengthen only enough to fall into the extensional gap (Krifka, 2007). As all prior experimental work on negative strengthening has focused on relative adjectives, it remains an open question whether vague quantifiers (e.g., *many*, *few*) behave similarly. Here, we investigate how listeners interpret utterances containing negated vague quantifiers. We use an exposure-and-test paradigm experiment to examine (a) whether listeners track speaker-specific thresholds for vague quantifiers when expressed via negated assertions embedded in metalinguistic disagreements, and (b) whether negated positive quantifiers strengthen to their antonyms more than negated negative quantifiers. The results show that participants successfully track speaker thresholds for quantifier use following exposure to negated assertions in metalinguistic disagreements, and that both positive (e.g., *many*) and negative (e.g., *few*) quantifiers strengthen symmetrically toward their antonyms, suggesting that negative strengthening displays variability within and across lexical categories.

## Vagueness and metalinguistic disagreement

The meaning of vague predicates is underspecified and sensitive to both the speaker and the context. Partee (2004) models their meaning as follows:

- (1) a. Many of [the students]<sub>A</sub> [were late.]<sub>B</sub>  
 b.  $\frac{|A \cap B|}{|A|} \geq \theta$

(1-b) posits that *many* is a relation between two sets—the set of students and the set of individuals who were late—and a threshold variable  $\theta$  that denotes the minimum proportion for the quantifier to be applicable. This threshold, however, is underspecified and dependent on both the context and the speaker (Fernando & Kamp, 1996; Kennedy, 2007; Pepper & Prytulak, 1974; Ramotowska et al., 2024). The interspeaker variability of these thresholds can result in metalinguistic disagreements about a predicate’s usage.

In a metalinguistic disagreement, interlocutors agree on the facts of the situation, but disagree on how best to describe it (Barker, 2013; Horn, 1989). In (2), for example, S1 and S2 can both know the exact proportion of students who were late, but still disagree over whether or not the use of the vague quantifier *many* is appropriate for the context.

- (2) a. S1: Many of the students were late.  
 b. S2: I don’t think many of the students were late.

The disagreement in (2) is metalinguistic because the object of discussion is a linguistic parameter, i.e., the vague threshold. When S1 asserts that many of the students were late, they are saying that their threshold for what constitutes *many* is less than the proportion of students who were late. When S2 disagrees, they are saying that their threshold for what constitutes *many* exceeds the proportion of students who were late. Since the facts of the world are controlled, the disagreements are due to mismatches in thresholds. Thus, metalinguistic disagreements allow us to probe threshold representations.

## Tracking speaker opinions

Since these thresholds are underspecified, they remain dynamically negotiable within a local discourse context. Listeners can adapt their own meaning representations for different vague predicates in response to a speaker’s usage (Heim et al., 2020; Schuster & Degen, 2020; Xiang, Kramer, & Kennedy, 2020; Yildirim et al., 2016). Prior work has also demonstrated that listeners are able to track speaker thresholds expressed via non-negated assertions (Schuster & Degen, 2019, 2020; Yildirim et al., 2016). Within metalinguistic disagreements, listeners can adapt their own meaning representations for vague quantifiers (Pecsok & Aparicio, 2024) and imprecise predicates (Wu & Aparicio, 2025), as well as track speaker thresholds when presented as non-negated assertions (Pecsok & Aparicio, in prep).

All prior work on speaker adaptation in the domain of meaning has focused on listeners adapting to non-negated assertions; no work has examined whether listeners can adapt to a

speaker’s thresholds expressed via negated assertions. Negative sentences can incur more processing costs than their positive counterparts, potentially hindering a listener’s ability to adapt. The amount of additional processing cost associated with negated assertions varies by context; negation that is pragmatically licensed can have little to no effect on processing cost (Nieuwland & Kuperberg, 2008; Xiang, Kramer, & Nordmeyer, 2020), but generally there is an extra cognitive load associated with processing and verifying the truth of a negated assertion (Clark & Chase, 1972; Fischler et al., 1983; Kounios & Holcomb, 1992; Wason, 1961). Therefore, it remains unknown whether participants can track speaker thresholds with the added difficulty associated with processing negation. We address this gap in the literature by asking if participants adapt to negated assertions concerning vague quantifiers expressed as metalinguistic disagreements. Building on evidence that participants successfully adapt to speaker thresholds expressed via non-negated assertions within metalinguistic disagreements (Pecsok & Aparicio, in prep), the current study utilizes this established paradigm to determine whether negation hinders a listener’s ability to track this information. That is, we ask whether listeners can track the opinions expressed by the speaker via a metalinguistic challenge, similar to S2 in (2-b). Our results provide evidence for participants tracking a speaker’s threshold for vague predicates when expressed through negated assertions in metalinguistic disagreements, despite the potentially increased cognitive difficulty associated with negation.

## Negative strengthening

The meaning of negated vague predicates is contested. Most existing work focuses on the negation of relative adjectives. Similar to vague quantifiers, relative adjectives exist on the same scale as their antonyms (Bierwisch, 1989), but have different underspecified thresholds at which they become applicable, resulting in antonyms having an extensional gap on the scale where neither adjective is applicable. Negated relative adjectives are generally analyzed as distinct from their antonymic counterparts, but their exact meaning is debated.

Horn (1989, 1993) argues that the meanings of negated relative adjectives are asymmetric, where the negated positive form strengthens to imply the antonym, while the negated negative form only strengthens enough to remain in the extensional gap, as demonstrated in (3).

- (3) Asymmetric Negative Strengthening (Horn)  
 a. My apartment is not large.  $\rightsquigarrow$  My apartment is small.  
 b. My apartment is not small.  $\rightsquigarrow$  My apartment is neither small nor large.

Horn provides a pragmatic account of this phenomenon by assuming that the meaning of a negated adjective is the complement of the adjective itself (i.e., *not large* means *small* and the extensional gap and *not small* means *large* and the extensional gap). This leads to cases where both a negated

adjective and its antonym are applicable. Using Grice's Maxims of Quantity, Manner, and Relation, Horn claims a speaker will opt for a briefer term whenever possible. Therefore, the meaning of the more complex negated terms is restricted to the extensional gap, where neither the positive nor the negative adjective is applicable. However, the negated positive adjective strengthens to the antonym because a listener can draw other pragmatic inferences as to why a speaker would use the weaker term for face-saving reasons. Since positive adjectives often have more positive connotations than their negative counterparts, the speaker in (3-a) could opt for the weaker term *not large* to conceal the stronger meaning *small*. This alternative explanation for why a speaker would opt for the more complex and weaker term when negating positive relative adjectives allows for a strengthened interpretation toward the antonym. These social pressures do not exist when negating the negative adjective, making this strengthening asymmetric.

Krifka (2007), on the other hand, argues that the meanings of antonym pairs when negated are symmetric. More specifically, Krifka claims they both fall into the gap where neither adjective is applicable, but the negated positive form falls in the lower part of the gap and the negated negative form falls in the upper part of the gap, as demonstrated in (4).

- (4) Symmetric Negative Strengthening (Krifka)
- a. My apartment is not large.  $\rightsquigarrow$  My apartment is neither small nor large, but closer to small.
  - b. My apartment is not small.  $\rightsquigarrow$  My apartment is neither small nor large, but closer to large.

Similar to Horn, Krifka argues that between two expressions with overlapping meanings, a speaker will opt for a briefer expression to describe a situation, restricting the negated terms to the extensional gap.

Prior experimental work on negative strengthening has only focused on adjectives, and while evidence for symmetric negative strengthening has been found in a set of studies for relative adjectives (Giora et al., 2004), and in others for absolute adjectives (Alexandropoulou & Gotzner, 2024; Gotzner & Mazzarella, 2024), the majority support the theory of asymmetric negative strengthening for relative adjectives (see Colston, 1999; Gotzner & Mazzarella, 2021; Mazzarella & Gotzner, 2021, a.o.). However, Ruytenbeek et al. (2017) found that negative strengthening is a gradient process and the amount that an adjective strengthens depends on the adjective itself. Given previous experimental work, it is unlikely that Krifka's theory of symmetric negative strengthening is true for all vague predicates; however, the by-item variation found by Ruytenbeek et al. (2017) suggests that vague quantifiers could undergo symmetric negative strengthening. In particular whether their meanings when negated (e.g., *not many*) are distinct from the meanings of their non-negated counterparts (e.g., *few*) and whether they undergo asymmetric negative strengthening or maintain symmetric interpretations under negation. In what follows, we spell out our predictions for the research questions pertaining to the interpretation of negated quantifiers.

**Interpretations of negation vs antonyms** Given that negated vague predicates have been claimed to have a distinct meaning, despite potential overlap with their antonyms (Horn, 1989; Krifka, 2007), we anticipate that participants will interpret the negated vague quantifiers as distinct from their antonymic counterparts when presented as responses in metalinguistic disagreements. We also anticipate that there will be some level of strengthening, such that the negated quantifiers have distinct interpretations from each other, in line with both Horn and Krifka's analyses. Indeed, our results provide evidence for participants interpreting negated vague quantifiers as distinct from their antonyms when used as a response within a metalinguistic disagreement. Interpretations of both negated vague quantifiers were more likely to fall into the extensional gap, in contrast with their antonyms, but the meanings of the two negated quantifiers are distinct from each other, suggesting that the updated interpretations involved some amount of negative strengthening.

**Asymmetric vs symmetric negative strengthening** If the behavior of negated vague quantifiers follows that of negated relative adjectives as posited by Horn (1989, 1993), we would expect asymmetric strengthening, where the meaning of *not many* would strengthen to overlap with *few*, while the meaning of *not few* would not strengthen to overlap with *many*. If the behavior of negated vague quantifiers follows what was predicted by Krifka (2007), we would expect symmetric strengthening of the negated vague quantifiers. Our results provide evidence that the strengthening found when negating both positive and negative quantifiers is symmetric. That is, there is no significant difference in how *not many* is interpreted in relation to *few* compared to how *not few* is interpreted in relation to *many*. Neither vague quantifier was significantly weaker than the other in relation to its antonym.

## Experiment

The goal of our experiment is to determine if participants can adapt to a speaker's threshold for the vague quantifiers *many* and *few* when expressed as negated assertions within metalinguistic disagreements. Further, it seeks to address whether the meaning of the positive quantifier *many* strengthens to the antonym when negated more than the negative quantifier *few*.

## Materials and procedure

Our experiment followed a blocked design, similar to other work on speaker adaptation (Schuster & Degen, 2020; Yildirim et al., 2016). The first block exposed participants to the range of contexts in which a speaker would and would not use *many* or *few*, which we refer to as the speaker's bias. In the second block, participants were asked in different contexts whether the speaker would be likely to use *many*, *few*, or neither quantifier. The exposure block consisted of 20 trials. In each trial, participants saw a series of images containing 50 circles of different sizes. Each image had blue and/or yellow circles on a gray background, ranging from having 0-100% of either color. To the left of the image with circles was a

| Condition | Dialogue                                                                                                                       |
|-----------|--------------------------------------------------------------------------------------------------------------------------------|
| Negation  | S1: { <i>Many/few</i> } of the circles are blue.<br>S2: I disagree, I don't think { <i>many/few</i> } of the circles are blue. |
| Antonym   | S1: { <i>Many/few</i> } of the circles are blue.<br>S2: I disagree, I think { <i>few/many</i> } of the circles are blue.       |

Table 1: Exposure block critical trial example dialogues. Colors correspond to dialogue pairs.

|                     | Many      |            | Few       |            | All      |          | None     |          |
|---------------------|-----------|------------|-----------|------------|----------|----------|----------|----------|
|                     | <i>n</i>  | <i>P</i>   | <i>n</i>  | <i>P</i>   | <i>n</i> | <i>P</i> | <i>n</i> | <i>P</i> |
| <i>Many</i> -biased | <b>10</b> | <b>50%</b> | 5         | 20%        | 5        | 100%     | —        | —        |
| <i>Few</i> -biased  | 5         | 80%        | <b>10</b> | <b>50%</b> | —        | —        | 5        | 0%       |

Table 2: Number of exposure trials (*n*) per quantifier with the proportion of target color circles (*P*) in Speaker 1 *many*-biased vs Speaker 1 *few*-biased conditions. Critical trials are bolded.

picture of the first speaker, and to the right was a picture of the second speaker. The images of the speakers appeared as their corresponding audio clip played. The auditory stimuli varied based on the trial type. Speaker 1 would describe the image and, in the filler trials, Speaker 2 would agree with the description, while in the critical trials, Speaker 2 would disagree with Speaker 1's description.

In the critical trials, we manipulated Speaker 1 Bias, which was whether Speaker 1 used *many* or *few* to describe 50% of the circles. We also manipulated Response Type, which was whether Speaker 2 disagreed by negating Speaker 1's utterance or by asserting the antonym. To avoid the infelicitous phrase *not few* in the negation groups, Speaker 2 said *I don't think {many/few} ...* to disagree with Speaker 1. We used a 2x2 between-subjects design with Speaker 1 Bias and Response Type as factors. Table 1 contains example dialogues in the critical trials, and Table 2 contains a summary of the conditions. Every trial featured one male and one female speaker. For each participant, the first speaker remained consistent, but the gender of the first speaker was counterbalanced across participants. See Figure 1 for an example critical trial in the exposure block.

The test block was modeled after Yildirim et al. (2016). In it, participants saw the same images with circles and the same Speaker 2 as in the exposure trials. Participants were presented with images ranging from having 20-80% of either color (in 10% increments) and were asked to indicate how likely Speaker 2 would be to use *many*, *few*, or neither quantifier to describe the image. They did this by distributing 100 points across three sliders. The sliders automatically adjusted to ensure exactly 100 points were used (see Figure 2 for an example trial). There were 56 critical trials (8 for each target color percentage) and 8 attention checks where participants were asked to indicate the likelihood that Speaker 2 would use *all* or *none* to describe images with 0% or 100% of the target color.

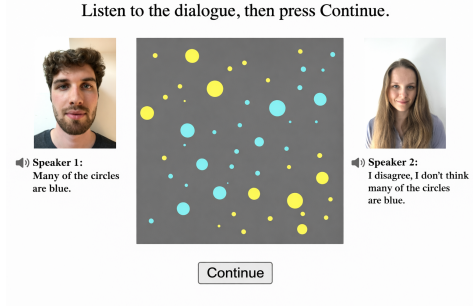


Figure 1: Exposure block critical trial example in the Speaker 1 *many*-Bias and Speaker 2 Negation Response Condition.

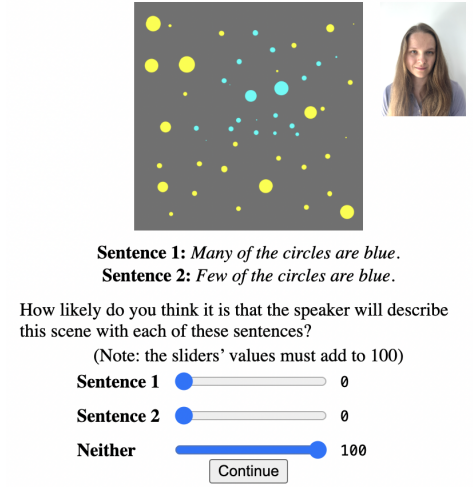


Figure 2: Test block example.

## Participants

The experiment was coded and run online using the PClbex platform (Zehr & Schwarz, 2018). Sixty-one participants were recruited via Prolific. Participants were born and currently located in the United States and self-described native English speakers between the ages of 18 and 25. They were paid \$3.40, which amounted to an hourly wage of approximately \$12.00. We excluded participants who failed more than 4 attention checks in the test block. Six participants were excluded based on the above criteria, leaving a total of 55 participants.

## Results

**Adaptation** First, we tested if participants adapt to speaker thresholds when expressed as negated assertions by assessing if responses differed across the two Speaker 1 bias conditions (see Figure 3). Following Yildirim et al. (2016) and Schuster and Degen (2020), we quantified the differences between the participants exposed to the Speaker 1 who preferred *many* vs the Speaker 1 who preferred *few* at 50% by fitting natural splines with 2 degrees of freedom to the response curves for *many* and *few* for each participant. We then calculated the area under the curve (AUC) for each spline and computed a




Figure 3: Test block results by Speaker 2 Response Type. Error bars correspond to 95% confidence intervals.

difference score by subtracting the AUC for *few* from the AUC for *many*. We disregarded the AUC for *neither* as there is no theoretical reason to expect a significant difference in *neither* responses across biases within response type.

We compared these difference scores across the two exposure groups using a Welch two-sample t-test for each of the two Speaker 2 response types: antonym (to test for adaptation to non-negated assertions)<sup>1</sup> and negation (to test for adaptation to negated assertions). We found participants exposed to the *many*-biased Speaker 1—meaning their Speaker 2 said either *not many* or *few* at 50% depending on the Response Type condition—gave significantly higher probability ratings for Speaker 2 using *few* over *many* compared to those exposed to the *few*-biased Speaker 1, both when Speaker 2 responded with the antonym ( $t(18.90) = 2.17, p = 0.043$ ) and when Speaker 2 responded with metalinguistic negation ( $t(23.47) = 2.41, p = 0.024$ ), demonstrating that listeners can track speaker thresholds presented as negated assertions.

**Interpretations of negation vs antonyms** The above results indicate not only that participants are able to adapt to preferences expressed via negated assertions, but also that the negated quantifiers are interpreted as distinct from each other, indicating that there is some level of strengthening. We now assess the strength of this effect by looking at whether participants interpreted negated vague quantifiers as different from their antonyms by testing if responses differed across the two response types for Speaker 2 (see Figure 4 for results when Speaker 1 preferred *many* and Figure 5 for results when Speaker 1 preferred *few*). Similar to the adaptation analysis, we quantified this by calculating the AUC for the spline-fitted response curves. However, unlike before, we included the results for *neither*, as we anticipate a difference across the two Speaker 2 response types. Participants may interpret the negation of Speaker 1’s assertion as Speaker 2 preferring the extensional gap of neither *many* nor *few*, while that same uncertainty does not exist when Speaker 2 responds with the antonym. Therefore, we expect higher rates of *neither* and

<sup>1</sup>We test for adaptation in the antonym condition to replicate the findings of Pecsok and Aparicio (in prep) that participants adapt to speaker thresholds expressed within non-negated assertions in a metalinguistic disagreement.

lower rates of *many* and *few* when Speaker 2 disagrees via metalinguistic negation compared with when they disagree using the antonym. For each participant, we computed a composite score:  $(AUC_{Many} + AUC_{Few}) - AUC_{Neither}$ , where lower scores indicate an increased reliance on *neither*.

Welch’s t-tests comparing the Speaker 2 Response types confirmed that the negation condition yielded significantly lower scores than the antonym condition, both when Speaker 1 preferred *many* ( $t(17.93) = 3.81, p = 0.001$ ) and when Speaker 1 preferred *few* ( $t(15.30) = 2.15, p = 0.047$ ). These results indicate that participants exposed to a Speaker 2 who replied with negation relied more heavily on *neither* as a response than participants exposed to a Speaker 2 who replied with the antonym, indicating that a quantifier’s negation is not interpreted as equivalent to its antonym (i.e., *not many* is not equivalent to *few* and *not few* is not equivalent to *many*).




Figure 4: Test block results for Speaker 1 *Many* Bias. Error bars correspond to 95% confidence intervals.




Figure 5: Test block results for Speaker 1 *Few* Bias. Error bars correspond to 95% confidence intervals.

**Symmetric negative strengthening** Despite the fact that the metalinguistically negated quantifiers are both interpreted as weaker than their antonyms, there is still some level of negative strengthening, as participants in the negation condition do not respond exclusively with *neither* at 50%. The question remains whether this strengthening is symmetric, that is, if one quantifier strengthened toward its antonym more than the other. We fit a linear mixed-effects model with fixed effects for Speaker 1 Bias, Speaker 2 Response Type, and their interaction. We analyzed ratings for the specific quantifiers (*many*

and *few*), excluding the *neither* category to ensure independence, while controlling for the percentage of the target color and including random intercepts for participants.

We found no significant interaction between Speaker 1 Bias and Speaker 2 Response Type ( $\beta = -3.80$ ,  $t = -0.83$ ,  $p = 0.408$ ), indicating that the difference in the interpretation of the negated assertion compared to its antonym was consistent across quantifier polarity. Put differently, the distance between the interpretation of *not many* and the interpretation of *few* was not significantly different from the distance between the interpretation of *not few* and the interpretation of *many*.

## Discussion

Our results show that participants exposed to metalinguistic disagreements can adapt to a speaker's threshold for vague quantifier usage presented in negated assertions. This could be the result of speaker adaptation being resilient to the increased cognitive load associated with processing negation, or it could be that there is no added processing cost because the negated assertion is pragmatically licensed and informative, as it is within a disagreement (Nieuwland & Kuperberg, 2008; Xiang, Kramer, & Nordmeyer, 2020). Regardless, these results align with previous findings supporting speaker adaptation for vague quantifiers when presented in non-negated assertions within a metalinguistic disagreement (Pecsok & Aparicio, in prep) and other previous work showing speaker adaptation to different types of vague expressions (Schuster & Degen, 2019, 2020; Yildirim et al., 2016), as they indicate speaker adaptation is possible after exposure to negated assertions.

We further conclude that participants interpret negated vague quantifiers as weaker than their antonyms. That is, the negated quantifiers *not many* and *not few* both elicited higher rates of *neither* responses compared to their antonyms *few* and *many*, indicating the interpretation of these negated quantifiers overlaps with the extensional gap of *neither many nor few*. These results are consistent with Horn (1989) and Krifka (2007)'s theories about the semantics of negated relative adjectives, which predict that such meanings should be distinct from their antonyms, even if the meanings overlap.

Finally, our results support the conclusion that the negative strengthening of the two vague quantifiers is symmetric. The interaction between quantifier and response type was not significant, suggesting no evidence that the difference between *not many* and *few* and the difference between *not few* and *many* were significantly different. These findings are in contrast with much of the experimental work supporting asymmetric interpretations of negated relative adjectives. There are several possible explanations for these results. A seemingly straightforward possibility is that Krifka (2007)'s theory of negated adjectives having symmetric interpretations is correct for all relative adjectives and similar vague predicates. Ultimately, this explanation seems unlikely given prior experimental work supporting asymmetric negative strengthening (Gotzner & Mazzarella, 2021; Mazzarella & Gotzner, 2021). Another alternative is that this symmetric negative strength-

ening, as described by Krifka (2007), is a property only of vague quantifiers and not of relative adjectives. Once again, this theory seems unlikely as there are no inherent properties across which vague quantifiers and relative adjectives differ that could suggest this variation in behavior across categories.

A more appealing option is that asymmetric negative strengthening is the result of the positive adjective conventionally having positive connotations that would make the speaker prefer to use the negated positive over the antonym to save face, as Horn (1989) described, but then the effect is not realized when those same connotations associated with an antonym pair are missing. While previous work has shown that asymmetric negative strengthening remains even when context is manipulated (Gotzner & Mazzarella, 2021), negative strengthening is a gradient effect, where, in general, adjectives with positive connotations strengthen more than those without (Ruytenbeek et al., 2017). Compared to the adjective *large*, which has more positive connotations and could strengthen toward *small* when negated more than *not small* could strengthen to *large*, *many* is not associated with the same positive connotations, and therefore would not be predicted to strengthen toward its antonym when negated any more than *few*. It is reasonable to predict that asymmetric negative strengthening is the result of the associations of the adjectives in the adjective pair, although further testing is required in future work. Regardless, our results do not support the hypothesis that asymmetric negative strengthening is an inherent property of all vague predicates, including vague quantifiers, and our results provide evidence that asymmetric negative strengthening could be subject to by-item variation, where antonym pairs without conventionalized positive and negative connotations default to symmetric negative strengthening.

Finally, future work will also need to determine whether the symmetric strengthening effects reported here for vague quantifiers remain when they are negated not as part of a metalinguistic disagreement, but rather as an answer to a question (e.g., *Do you think many of the circles are blue?*). Future work should also examine whether asymmetric strengthening emerges when the context induces a positive or negative association with one of the quantifiers.

## Conclusion

The present study investigated how listeners interpret and track negated vague predicates, specifically the vague quantifiers *many* and *few*, when presented within metalinguistic disagreements. Our results showed that despite increased cognitive demands associated with processing negated assertions, listeners successfully track and adapt to a speaker's individualized thresholds when expressed via negated assertions. Furthermore, we confirmed not only that negated vague quantifiers are interpreted as distinct from their antonymic counterparts, but also that their interpretations were equally different from their antonymic counterparts, indicating symmetric negative strengthening of the quantifiers *many* and *few* and supporting a theory of by-item variation in negative strengthening.

## References

- Alexandropoulou, S., & Gotzner, N. (2024). The interpretation of relative and absolute adjectives under negation. *Journal of Semantics*, 41(3-4), 373–399.
- Barker, C. (2013). Negotiating taste. *Inquiry*, 56(2-3), 240–257. <https://doi.org/10.1080/0020174X.2013.784482>
- Bierwisch, M. (1989). The semantics of gradation (M. Bierwisch & E. Lang, Eds.). *Dimensional adjectives: Grammatical structure and conceptual interpretation*, 71–261.
- Clark, H. H., & Chase, W. G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology*, 3(3), 472–517. [https://doi.org/10.1016/0010-0285\(72\)90019-9](https://doi.org/10.1016/0010-0285(72)90019-9)
- Colston, H. L. (1999). “Not good” is “bad,” but “not bad” is not “good”: An analysis of three accounts of negation asymmetry. *Discourse Processes*, 28(3), 237–256. <https://doi.org/10.1080/01638539909545083>
- Fernando, T., & Kamp, H. (1996). Expecting many. *Semantics and Linguistic Theory*, 6, 53–68. <https://doi.org/10.3765/salt.v6i0.2761>
- Fischler, I., Bloom, P. A., Childers, D. G., Roucos, S. E., & Perry Jr, N. W. (1983). Brain potentials related to stages of sentence verification. *Psychophysiology*, 20(4), 400–409.
- Giora, R., Balaban, N., Fein, O., & Alkabets, I. (2004). Negation as positivity in disguise. In *Figurative language comprehension* (pp. 233–258). Routledge.
- Gotzner, N., & Mazzarella, D. (2021). Face management and negative strengthening: The role of power relations, social distance, and gender. *Frontiers in Psychology*, 12, 602977. <https://doi.org/10.3389/fpsyg.2021.602977>
- Gotzner, N., & Mazzarella, D. (2024). Negative strengthening: The interplay of evaluative polarity and scale structure. *Journal of Semantics*, 41(1), 103–117.
- Heim, S., Peiseler, N., & Bekemeier, N. (2020). “Few” or “Many”? An Adaptation Level Theory Account for Flexibility in Quantifier Processing. *Frontiers in Psychology*, 11, 382. <https://doi.org/10.3389/fpsyg.2020.00382>
- Horn, L. R. (1989). *A natural history of negation*. University of Chicago Press.
- Horn, L. R. (1993). Economy and redundancy in a dualistic model of natural language. *Finnish Journal of Linguistics*, 6(3), 33–72.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1), 1–45. <https://doi.org/10.1007/s10988-006-9008-0>
- Kounios, J., & Holcomb, P. J. (1992). Structure and process in semantic memory: Evidence from event-related brain potentials and reaction times. *Journal of Experimental Psychology: General*, 121(4), 459–479.
- Krifka, M. (2007). Negated antonyms: Creating and filling the gap. *Presupposition and implicature in compositional semantics*, 163–177.
- Mazzarella, D., & Gotzner, N. (2021). The polarity asymmetry of negative strengthening: Dissociating adjectival polarity from face-threatening potential. *Glossa: a journal of general linguistics*, 6(1). <https://doi.org/10.5334/gjgl.1342>
- Nieuwland, M. S., & Kuperberg, G. R. (2008). When the truth is not too hard to handle: An event-related potential study on the pragmatics of negation. *Psychological Science*, 19(12), 1213–1218.
- Partee, B. H. (2004). Many quantifiers. In B. H. Partee (Ed.), *Compositionality in Formal Semantics* (1st ed., pp. 241–258). Wiley. <https://doi.org/10.1002/9780470751305.ch12>
- Pecsok, E., & Aparicio, H. (2024). How can they both be right?: Faultless disagreement and semantic adaptation. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 46, 3939–3945.
- Pecsok, E., & Aparicio, H. (in prep). Meaning adaptation under conflicting input.
- Pepper, S., & Prytulak, L. S. (1974). Sometimes frequently means seldom: Context effects in the interpretation of quantitative expressions. *Journal of Research in Personality*, 8(1), 95–101. [https://doi.org/10.1016/0092-6566\(74\)90049-X](https://doi.org/10.1016/0092-6566(74)90049-X)
- Ramotowska, S., Haaf, J., Van Maanen, L., & Szymanik, J. (2024). Most quantifiers have many meanings. *Psychonomic Bulletin & Review*, 31(6), 2692–2703. <https://doi.org/10.3758/s13423-024-02502-7>
- Ruytenbeek, N., Verheyen, S., & Spector, B. (2017). Asymmetric inference towards the antonym: Experiments into the polarity and morphology of negated adjectives. *Glossa: a journal of general linguistics*, 2(1), 1–27. <https://doi.org/10.5334/gjgl.151>
- Schuster, S., & Degen, J. (2019). Speaker-specific adaptation to variable use of uncertainty expressions. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 41, 2769–2775. <https://escholarship.org/uc/item/2bz715sm>
- Schuster, S., & Degen, J. (2020). I know what you’re probably going to say: Listener adaptation to variable use of uncertainty expressions. *Cognition*, 203, 1–31. <https://doi.org/10.1016/j.cognition.2020.104285>
- Tian, Y., Breheny, R., & Ferguson, H. J. (2010). Why we simulate negated information: A dynamic pragmatic account. *Quarterly Journal of Experimental Psychology*, 63(12), 2305–2312. <https://doi.org/10.1080/17470218.2010.525712>
- Wason, P. C. (1961). Response to affirmative and negative binary statements. *British Journal of Psychology*, 52(2), 133–142.
- Wu, Y., & Aparicio, H. (2025). Meaning adaptation in the discourse dynamics of imprecision. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47, 102–108. <https://escholarship.org/uc/item/6r9914v5>
- Xiang, M., Kramer, A., & Kennedy, C. (2020). Semantic adaptation in gradable adjective interpretation. *unpublished MS, U Chicago*.
- Xiang, M., Kramer, A., & Nordmeyer, A. E. (2020). An informativity-based account of negation complexity. *Journal of Experimental Psychology: Learning, Memory, and*

- Cognition*, 46(10), 1857–1867. <https://doi.org/10.1037/xlm0000851>
- Yildirim, I., Degen, J., Tanenhaus, M. K., & Jaeger, T. F. (2016). Talker-specificity and adaptation in quantifier interpretation. *Journal of Memory and Language*, 87, 128–143. <https://doi.org/10.1016/j.jml.2015.08.003>
- Zehr, J., & Schwarz, F. (2018). PennController for internet based experiments (ibex). <https://doi.org/10.17605/OSF.IO/MD832>