

## How does a speaker's intent to deceive affect scalar inference and lie judgments?

The question of whether false implicatures are lies has interested theoreticians at the semantics/pragmatics interface for several years,<sup>1,2</sup> with more recent work turning to experimental evidence to help clarify the picture.<sup>3,4</sup> Despite differing results, this experimental work has begun to highlight elements of context that explain variation in such judgments, with a particular eye toward intention to deceive. If context establishes that a speaker has a clear intention to deceive a hearer, participants will reliably rate a false implicature from that speaker as more of a lie than in context without such an intention established.<sup>5</sup>

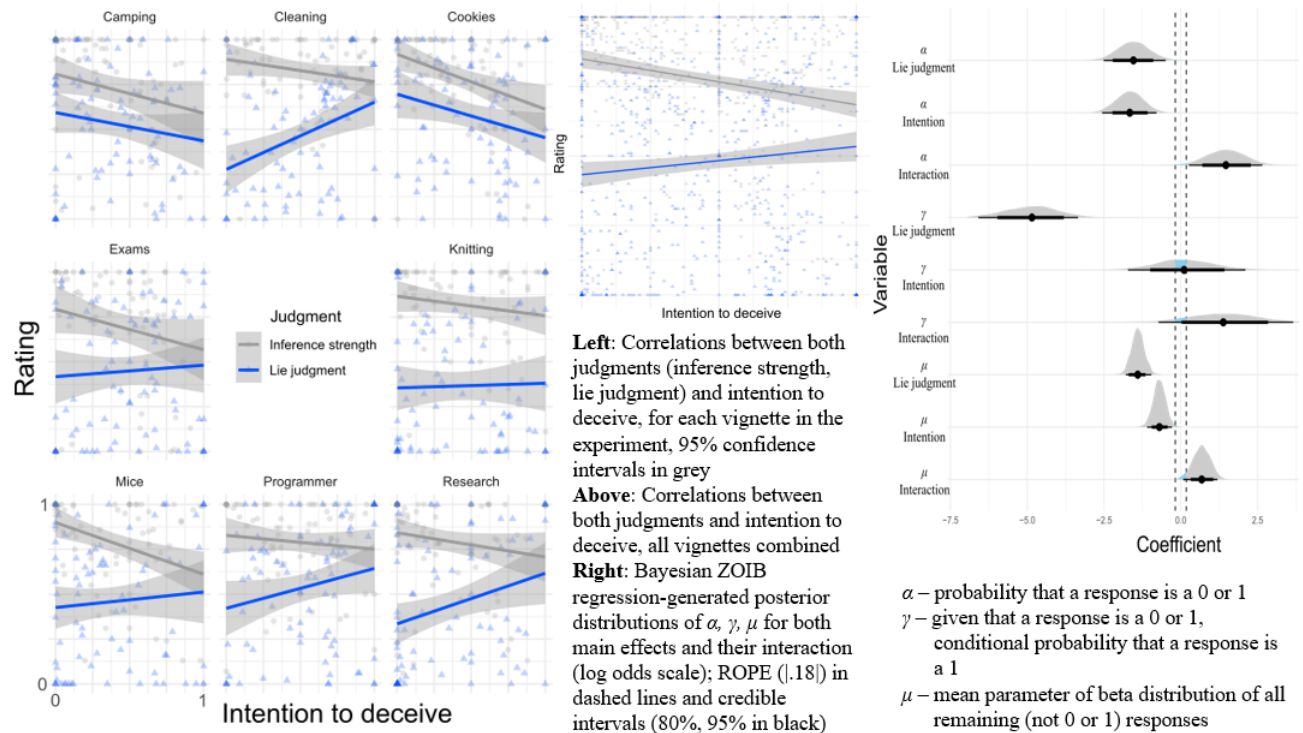
The present research adds to this strand of false implicature research by connecting to recent work investigating the effects of elements of context on the strength of scalar inference.<sup>6</sup> That study found that speaker competence (i.e., whether the speaker knows whether *all* is true when they use *some*) and prior probability (i.e., the a priori likelihood of *all* being true) significantly affect the strength of the *some but not all* inference made by participants. Inference strength is closely tied to lie judgments as well, as recent proposals argue that the level of commitment attributed to the speaker with respect to false implicated content modulates the degree to which that utterance will be considered a lie.<sup>3,7</sup>

The present work aims to investigate whether perceived intention to deceive significantly affects (a) strength of scalar inference drawn and (b) lie judgment of the utterance; in addition, this research will compare the magnitude with which intention to deceive affects (a) and (b). Eight vignettes were crafted, each of which led to a speaker delivering a line licensing a *some but not all* inference. For each vignette, two versions were created: one with a clear motivation for the speaker to try to deceive the hearer, and one without such a motivation explicitly provided; this intention to deceive is counterbalanced within-subjects across the vignettes.

Participants in the main experiment first saw a vignette without the critical utterance and made a sliding-scale judgment about the likelihood of the speaker intending to deceive the hearer in the situation. Following a comprehension question, the second judgment probed either inference strength or lie judgments, varying between subjects. For the former, the critical utterance was added to the vignette, and participants provided a sliding-scale scalar inference judgment. For the latter, the truth of the situation was revealed (i.e., that *all* is, in fact, true) and the critical utterance is added to the vignette, prompting a sliding-scale lie judgment.

Thus far, 240 native English-speaking participants were recruited via Prolific (avg. age = 37.4, sd = 13.7, 117F/120M/3 other); data collection is continuing to 320. Each participant sees 16 items total – 8 *some but not all* items and 8 fillers without scalar implicatures. For participants in the lie judgment condition, the revealed truth of the critical *some* items is balanced between critical cases where *all* is actually true and distractor cases where *some* is actually true or *none* is actually true. Target and filler items are counterbalanced across the two intention conditions, and trial order is randomized for all participants.

A Bayesian mixed effects model regresses sliding scale ratings against fixed effects of speaker's intention to deceive, the judgment being made (scalar inference strength vs. lie judgment), and their interaction, with random intercepts and slopes by item and random intercepts by participant. Zero-one-inflated-beta regression is used due to inflation at 0 and 1 (the sliding scale extremes).



Results (thus far) highlight the complexities in assessing the relationship between these judgments. There is a consistent negative effect of intention to deceive on inference strength whereby scalar inference gets weaker as intention to deceive increases. There is an inconsistent positive effect of intention to deceive on lie judgments whereby lie ratings get higher as intention to deceive increases.

These findings appear to complicate the commitment-based account, though they do not necessarily refute it. This analysis adds to the growing body of research investigating effects of context on the strength of scalar inferences. It also begins to quantify the preliminarily-documented finding that intention to deceive affects lie judgments of false implicatures. Lastly, it helps to clarify the relationship between context, commitment, and message interpretation, or at least helps to highlight the complexity in such a relationship.

<sup>1</sup> Meibauer, J. (2014). *Lying at the Semantics-Pragmatics Interface*. Berlin: De Gruyter Mouton.

<sup>2</sup> Saul, J. M. (2012). *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford University Press.

<sup>3</sup> Reins, L. M., & Wiegmann, A. (2021). Is lying bound to commitment? Empirically investigating deceptive presuppositions, implicatures, and actions. *Cognitive Science*, 45(2).

<sup>4</sup> Weissman, B., & Terkourafi, M. (2019). Are false implicatures lies? An empirical investigation. *Mind & Language*, 34(2), 221–246.

<sup>5</sup> Wiegmann, Alex. (2022). Lying with deceptive implicatures? Solving a puzzle about conflicting results. *Analysis*, 1–11.

<sup>6</sup> Tsvilodub, P., Van Tiel, B., & Franke, M. (2023). The role of relevance, competence, and priors for scalar inferences. *Experiments in Linguistic Meaning*, 2, 288.

<sup>7</sup> Meibauer, J. (2023). On commitment to untruthful implicatures. *Intercultural Pragmatics*, 20(1), 75–98.