Polarity modulates the processing cost of scalar implicatures

Chao Sun (Leibniz-Centre General Linguistics ZAS), Elizabeth Pankratz (University of Potsdam), Bob van Tiel (Donders Institute for Brain, Cognition and Behaviour), c.sun@leibniz-zas.de

A sentence containing a scalar word (1a) often gives rise to a scalar implicature (SI) (1b), which is the negation of its scalar alternative (1c). Competing models have been proposed to explain SI processing. According to the *literal-first hypothesis*, listeners access the literal interpretation of sentences with scalar words before deriving SIs [1]. Hence, computing SIs should be slow and cognitively costly. By contrast, the *parallel processing* model argues that information from all sources, including both the literal interpretation and the corresponding SI, is processed in parallel. Hence, SIs should not be delayed and may not be associated with a processing cost [2]. Studies that investigated these models of SI computation have mostly focused on the lexical scale <some, all> [3, 4]. Yet [5] found that results for <some, all> do not always generalise across all lexical scales. [5] tested the cognitive processing of seven scalar words by measuring response times in a truth-value judgement task and by measuring responses in a memory load task. It was found that the SIs of positive scalar words (e.g., 'some', 'or') were associated with a processing cost, but those of negative scalar words (e.g., 'low', 'scarce') were not. This asymmetry can be explained by assuming that the SIs of positive scalar words express negative propositions (e.g., 'some' implying 'not all') while the SIs of negative scalar words express positive propositions (e.g., 'scarce' implying 'present'). The apparent processing cost of SIs can then be attributed to the difficulty of processing negative information.

To evaluate this *polarity hypothesis* during online processing, we conducted a self-paced decision task adapted from [6]. It has been shown that this paradigm replicates several key findings from visual world eye-tracking studies. In our study, 48 participants viewed a scene consisting of four objects (Fig. 1) while reading sentences containing positive scalar adverbs ('usually', 'always'), negative scalar adverbs ('rarely', 'never'), or exact probability words ('0/25/50/75/100% chance'). Sentences were presented in a self-paced manner starting with the article 'the', followed by a specific object ('marble' or 'arrow'), a scalar adverb and verb (e.g., 'usually lands on'), and a colour word. At each stage, participants had to click on what they believed to be the object described by the sentence. We measured which objects participants clicked on, as well as response times. Here, we only report the former.

During the critical adverbial time window, the literal interpretations of 'always' and 'never' (as well as the exact probability words) were sufficient to determine the correct reference. By contrast, 'usually' and 'rarely' were referentially ambiguous on their literal interpretations, but the correct reference could be established if their SIs were computed. If the computation of SIs is costly, we expect fewer clicks on the correct object (the target) for 'usually' and 'rarely' than for 'always' and 'never'. By contrast, if SIs are processed at no cost, we expect to find no significant difference. Lastly, if the polarity hypothesis is correct, we expect fewer clicks on the target for 'usually' compared to 'always', but not for 'rarely' compared to 'never'.

Fig. 2 shows the proportions of clicks on targets and competitors for different adverbs. We constructed a mixed effects logistic regression model predicting clicks on targets from condition ('always' and 'never' vs. 'usually' and 'rarely') and polarity (positive or negative). No effects were found in the first two windows. In the critical adverbial window, we found a significant interaction between condition and scalarity ($p = .01$). Planned comparison within the levels of condition revealed that the proportion of target clicks was lower for 'usually' compared to 'always' ($p = .007$), but there was no difference between 'never' and 'rarely'. Comparison within the levels of polarity revealed that the proportion of target clicks was lower for 'usually' compared to 'rarely' ($p = .003$) with no difference between 'always and 'never'. Finally, in the colour window, nearly all clicks were on the target.

Our results confirm the polarity hypothesis: the presence or absence of a processing cost for SIs is modulated by the polarity of scalar words. This finding suggests that the process of deriving SIs is without a cognitive cost, but that a cost may emerge during verification if the SI expresses a negative proposition.

1. a. The arrow usually lands on green.
   b. The arrow does not always land on green.
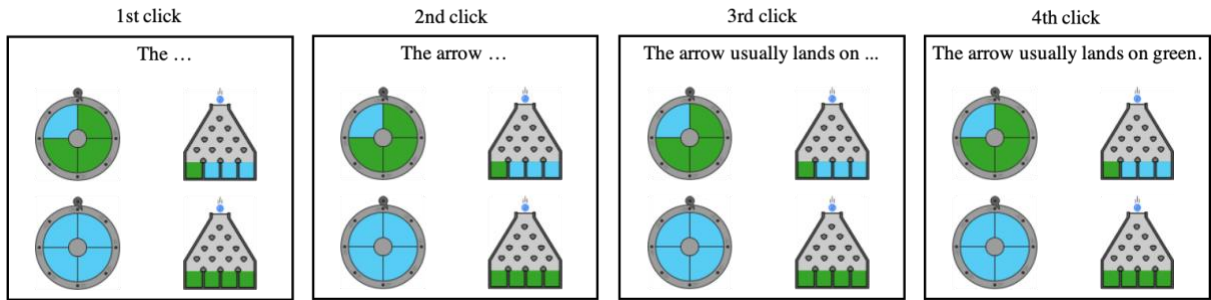   c. The arrow always lands on green.

Fig.1 Example display in each time window in the self-paced decision task.
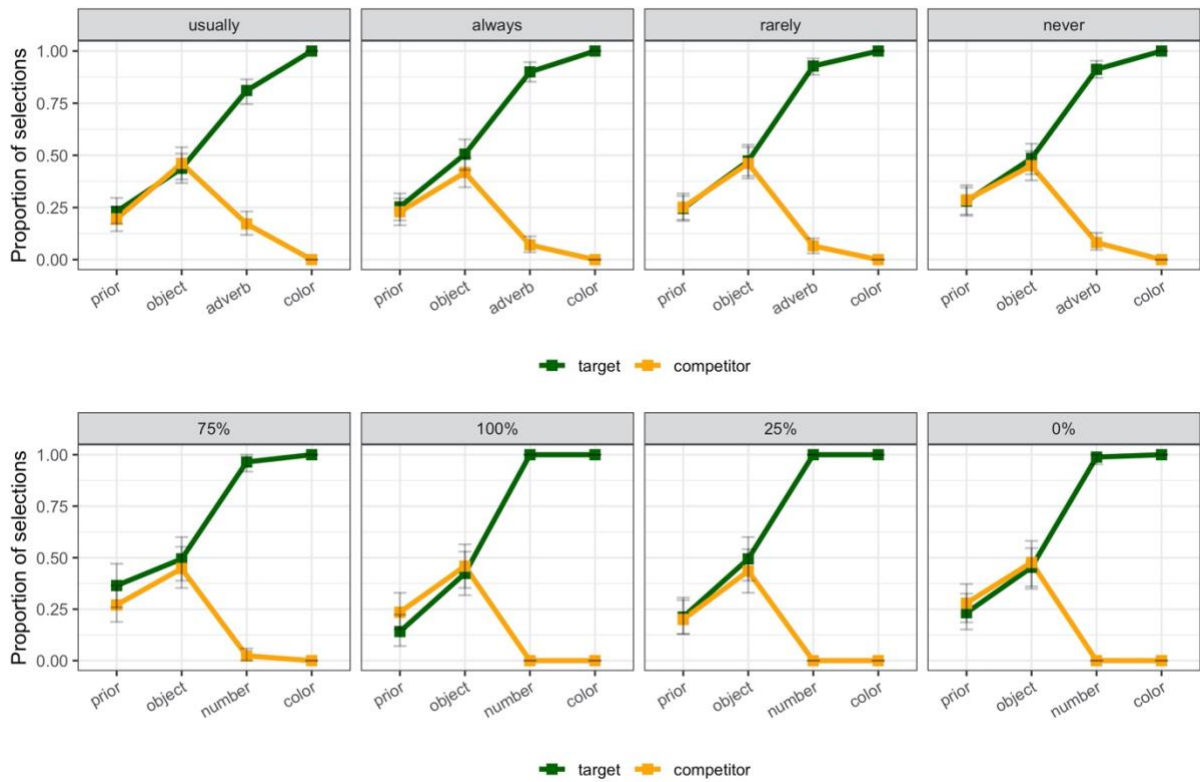

Fig.2 Proportions of clicks on targets and competitors for each condition and time window. Error bars represent 95% CIs.

References: [1] Huang, Y. T., & Snedeker, J. (2009). *Cognitive Psychology, 58*(3), 376–415. [2] Degen, J., & Tanenhaus, M. K. (2015). *Cognitive science, 39*(4), 667-710. [3] Bott, L., & Noveck, I. A. (2004). *Journal of Memory and Language, 51*, 437–457. [4] De Neys, W., & Schaeken, W. (2007). *Experimental Psychology, 54*, 128–133. [5] van Tiel, B., Pankratz, E., & Sun, C. (2019). *Journal of Memory and Language, 105*, 93-107. [6] Qing, C., Lassiter, D., & Degen, J. (2018). In *40th annual conference of the Cognitive Science Society*.