Every quantifier isn't the same: Informativity matters for ambiguity resolution in quantifier-negation sentences

Noa Attali
UC Irvine
nattali@uci.edu

Gregory Scontras
UC Irvine
g.scontras@uci.edu

Lisa S. Pearl
UC Irvine
lpearl@uci.edu

1 Introduction

Empirical work on quantifier-*not* sentences has focused primarily on universal quantifiers, exploring the ambiguity that arises when logical operators interact (e.g., *Everyone didn't go* could mean *No one went* or *Not all went*). In their Rational Speech Act model of this ambiguity resolution, Savinelli et al. (2017) demonstrate that pragmatic factors (such as model priors over the likely world states) stand to explain divergent behavior between children and adults. We extend this work to a broader empirical base, exploring the model's predictions for sentences with a wider range of quantifiers (*some*, *no*); we then test those predictions against behavioral data collected in a series of experiments. We find that a straightforward extension of the Savinelli et al. model captures the range of quantifier-*not* behavior we gather, thereby providing strong support for this cognitive model of probabilistic ambiguity resolution. In particular, the model explains interpretation preferences on the basis of informativity and prior beliefs over world states, such that interpretations that are more informative are preferred.

2 Model

Following Savinelli et al. (2017), we use a Rational Speech Act (**RSA**) model (Frank and Goodman, 2012) to formally articulate the cognitive process that yields observed interpretations of scopally-ambiguous utterances. We adapt Savinelli et al.'s model, which describes scope ambiguity resolution for *every*, to additionally account for *some* and *no* (see Figure 1). Hearing a quantifier-*not* utterance (i.e., *Every/Some/No marble isn't red*), a pragmatic listener L_1 reasons jointly about the true world state (i.e., the number of red marbles in a 3-marble world), the Question Under Discussion (**QUD**, i.e., the implicit topic: *Are all/no marbles red?, How many marbles are red?*), and the scope (surface: quantifier > *not*; inverse: *not* > quantifier).

$$P_{L_1}(w,i,q|u) \propto P(w)\cdot P(i)\cdot P(q)\cdot P_{S_1}(u|w,i,q)$$
$$P_{S_1}(u|w,i,q) \propto exp(\alpha \cdot \log(L_0(x|u,i,q)))$$
$$P_{L_0}(x|u,i,q) \propto \sum_w \delta_{x=[\![q]\!](w)} \cdot P_{L_0}(w|u,i)$$
$$L_0(w|u,i) \propto \delta_{[\![u]\!]^i(w)} \cdot P(w)$$

Figure 1: RSA model. Hearing a quantifier-*not* utterance, a pragmatic listener (L_1) reasons about the true world state w, the likely QUD q, and the scope interpretation i that the speaker (S_1) meant by utterance u. S_1 selects u by reasoning about a naive listener (L_0), who infers the state s on the basis of u's semantics, $[\![u]\!]^i$, a mapping from worlds to truth values parameterized by the scope interpretation i: e.g., $[\![\textit{every-not}]\!]^{\textit{surface}}$ maps only $w = 0$ to true and $[\![\textit{every-not}]\!]^{\textit{inverse}}$ maps only $w = 3$ to false. $\delta_{[\![u]\!]^i(w)} = 1$ when $[\![u]\!]^i(w)$ is true and 0 otherwise. A QUD q maps worlds to partitioned sets of worlds x: e.g., q= *are all red?* maps $w = 3$ to $x = \{3\}$ and $w = 0$ to $x = \{0, 1, 2\}$. $\delta_{x=[\![q]\!](w)} = 1$ when $x =[\![q]\!](w)$ and 0 otherwise.

Parameter setting. Savinelli et al. found that adult-like behavior (e.g., endorsing an *every-not* utterance like *Every marble isn't red* in a scenario where not all marbles are red) required a prior over world states that strongly favored the *all* state; we term this a "high positive expectation". Motivated by these analytic results, we set $P(w)$ in our model to favor all marbles being red ($P(w = 3) = 0.9$). We kept the priors over QUD and scope interpretation uniform, and set $\alpha = 1$ (i.e., no scaling).

Predictions. Under these parameter settings, the model predicts that the proportion of inverse interpretations depends on the quantifier: *every-not*>*no-not*>*some-not* (see Figure 3). The model makes these predictions because the listener L_1 reasons that the speaker S_1 maximizes the probability that the listener L_0 will arrive at the true world state. More specifically, given the prior expectation that *all marbles are red*, both *every-not* (surface scope: *none*; inverse scope: *not all*) and *some-not* (surface scope: *not all*; inverse scope: *none*) are maximally informative ways of conveying that the prior high positive expectation is false. However, there are

Proceedings of the Society for Computation in Linguistics (SCiL) 2021, pages 394-395.
Held on-line February 14-19, 2021

more ways for *not all* to be true than for *none* to be true (and so *not all* maximizes the probability of arriving at the true world state). So, the listener reasons that the speaker is most likely to intend the *not all* meaning for both (i.e., inverse for *every-not* and surface for *some-not*). For *no-not*, there is no strong pressure toward either interpretation; surface scope (*all*) is slightly preferred to inverse scope (*some*) because of a weak specificity implicature that strengthens the meaning of *no-not* to the stronger surface interpretation (*all*) compatible with the high positive expectation.

3 Experiments

To test the modeled listener's predictions, we measured interpretation behavior in a paraphrase-endorsement task (Scontras and Goodman, 2017) wherein participants encounter a potentially-ambiguous utterance (e.g., *Every marble isn't red*) and rate unambiguous paraphrases (e.g., *None/Not all of the marbles are red*). To identify the appropriate paraphrases for surface vs. inverse interpretations of our three potentially-ambiguous utterances, we had participants (N=94) complete a reference task: given a paraphrase, select the scenario that the paraphrase likely described. Participants chose at ceiling the image consistent with the intended scope interpretation for each of the paraphrases (e.g., the image of 0-of-3 red marbles for *None of the marbles are red* and 2-of-3 red marbles for *Not all of the marbles are red*). Having validated our paraphrases, we used them in the paraphrase-endorsement task where participants (N=47) saw three potentially-ambiguous utterances in random order and rated unambiguous paraphrases (Fig. 2).

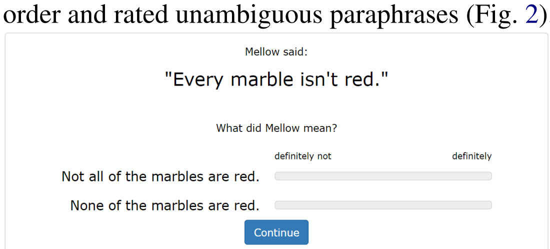

Figure 2: Sample paraphrase endorsement trial.

4 Results and discussion

Our model, which assumes a setting of the world state prior from Savinelli et al., not only captures the qualitative patterns in our data for this range of quantifiers, but also largely captures the quantitative patterns as well. With a high positive expectation for the state of the world, the model predictions

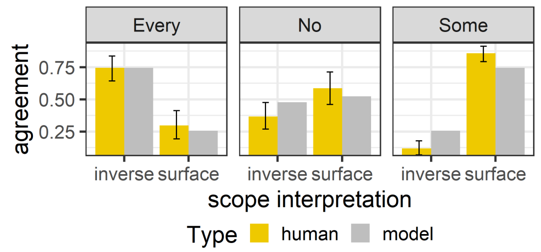

Figure 3: Results comparing model predictions and human data. Grey: Model predictions for L_1 marginal distribution over interpretation i. $P(w = 3) = 0.9$. Yellow: Degree of agreement with each paraphrase in task. Error bars are bootstrapped 95% CIs.

for the pragmatic listener's marginal distribution over scope interpretations align with interpretation patterns in the paraphrase endorsement task (Fig. 3). The results capture the intuition that adults understand a quantifier-*not* sentence as a negation of the high positive expectation.

In the case of *some*, human responses were more categorical than the model predictions, which may be due to *some*'s status as a positive polarity item (PPI) that does not scope under negation (Szabolcsi, 2004). However, our modeling results offer an explanation for why *some* would be a PPI in the first place: interpreting *some* under negation results in an utterance that is either uninformative or has an unlikely meaning, and therefore is not efficient.

More broadly, our adapted RSA model implements the hypothesis that interpretation proceeds through recursive reasoning, hinges on a cooperative, efficient speaker, and builds on certain prior expectations about the world. We empirically support this hypothesis for interpretations of scopally-ambiguous sentences with a universal, existential, and negative quantifier. In the process, we find direct support for the idea that pragmatic factors underpin adult interpretation behavior.

References

M. C. Frank and N. D. Goodman. Predicting pragmatic reasoning in language games. *Science*, 336(6084): 998–998, 2012.

K. Savinelli, G. Scontras, and L. Pearl. Modeling scope ambiguity resolution as pragmatic inference: Formalizing differences in child and adult behavior. In *CogSci*, 2017.

G. Scontras and N. D. Goodman. Resolving uncertainty in plural predication. *Cognition*, 168:294–311, 2017.

A. Szabolcsi. Positive polarity–negative polarity. *Natural Language & Linguistic Theory*, 22(2):409–452, 2004.