

## Learner's preference matches typological pattern of morphological marking: an animate-marking rule is easier to acquire than an inanimate-marking rule

Yiran Chen, Kathryn Schuler  
University of Pennsylvania

**Background.** An animacy hierarchy has been observed in many cases of morphological marking on nouns, including the marking of plural and case (Smith-Stark 1974; Dixon 1994). Across the world's languages, animate nouns are more likely to bear morphological markings than their inanimate counterparts, while the opposite pattern rarely occurs. What drives languages to settle on this universal preference for animate marking? One hypothesis is that some patterns are easier to learn than others due to language-specific or domain-general learning biases. Previous experimental work has shown that learners' preference matches typological frequencies of system of personal pronoun (Maldonado & Culbertson) and harmonic word order (Culbertson & Newport). Here we use an artificial language learning experiment to ask whether an animate-marking rule is easier to learn than an inanimate-marking rule.

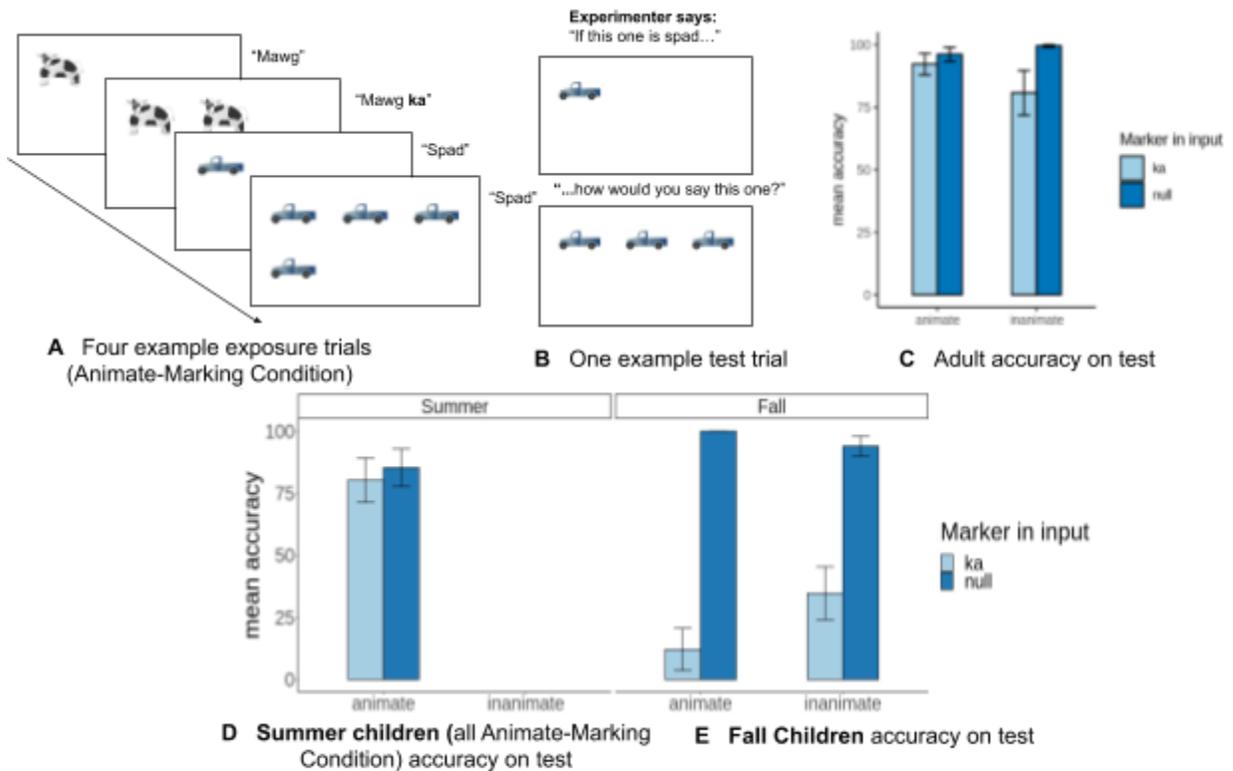
**Method.** We exposed 46 children (age 5.00-8.99 years) and 32 adults (age 18-26 years) to one of two languages in which plural marking is not applied across the board. As shown in Figure 1A, in one language, only animate nouns receive the plural marker "ka"; inanimate nouns' plural form is the same as their singular form (or the plural marker is "null"). We call this first condition Animate-Marking Condition. Reversedly, in Inanimate-Marking Condition, only inanimate nouns receive the plural marker "ka". The six nouns in the artificial language are randomly matched to the six objects in each experiment. At test, we ask children to produce the plural form of animate and inanimate nouns to determine whether they have learned the marking pattern, which is conditioned on animacy (Figure 1B).

**Results.** The predicted pattern is found in adults. As shown in Figure 1C, adults perform equally well on marked nouns compared to unmarked nouns (coefficient = -1.022,  $p > .05$ ). However, there is a significant interaction between condition and whether the noun is marked in input: adults have more difficulty learning the marked form in the Inanimate-Marking Condition (coefficient = -5.824,  $p < .001$ ).

In an initial proof-of-concept (Figure 1D), children run in the Animate-Marking Condition successfully learned the animate marking rule, correctly using the marker on animate nouns and not on inanimate nouns. However, the results of our full follow-up experiment (Figure 1E) are less clear: Children in both Animate-Marking Condition and Inanimate-Marking Condition accurately produce the nouns that are supposed to be unmarked in their language at test (inanimate nouns in Animate-Marking Condition and animate nouns in the Inanimate-Marking Condition). However, their accuracy is significantly worse on test trials in which the noun is marked with "ka" in the input (coefficient = -2.97543,  $p = .000$ ). For them, it is harder to learn all marked forms than unmarked forms. More interestingly, there is a significant interaction between condition and markedness in input: children in Inanimate-Marking Condition are even less accurate at producing the marked nouns in their language than children in the Animate-Marking Condition (coefficient = -1.553,  $p = .000$ ).

We are running follow-up analyses to investigate the possible effects of different experimenters and different test locations on children's dropped accuracy in Animate-Marking Condition.

**Conclusions.** Our preliminary results suggest that adult's learning of conditioned plural marking aligns with the patterns we see in language typology: learning a language in which morphological markings only fall on animates comes easily; learning a language in which only inanimates are marked is harder. Whether this tendency is also observed in children remains to be seen..



**Figure 1.** (A) Examples of four exposure trials in the Animate-Marking Condition, and (B) a single production test trial. Children were exposed to six singular and 12 plural sentences for each of six nouns (three animate and three inanimate) and tested on each of the nouns four times. Mean accuracy on the production test is shown in (C) for adults in the Animate-Marking Condition and Inanimate-Marking Condition, (D) for children run in the Animate-Marking Condition during summer and (E) for children run in fall semester in both Animate-Marking Condition and Inanimate-Marking Condition.

**Selected references:**

1. Smith-Stark, T. C. (1974, April). The plurality split. In Chicago Linguistic Society (Vol. 10, No. 1, pp. 657-672).
2. Maldonado, M., & Culbertson, J. Something about us: Learning first person pronoun systems.
3. Culbertson, J., & Newport, E. L. (2015). Harmonic biases in child learners: In support of language universals. *Cognition*, 139, 71-82.

# Scope Assignment and Scalar Implicatures in Child English: The Role of Working Memory

Shuyan Wang

University of Connecticut

**Introduction:** Acquisitional delays have been observed for scope assignment and scalar implicatures (e.g., Musolino 1998; Chierchia et al. 2001). Both delays have been proposed to follow from children's immature processing capacities. Yet, no one has examined the two phenomena within the same children. This paper provides within-subject data from English-speaking children on processing capacity, scope assignment for structures like (1), and scalar implicatures (<some, all>).

(1). There is a horse chasing every cow. ( $\exists > \forall / \forall > \exists$ )

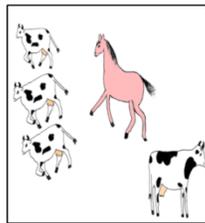
Reinhart (1999, 2004) proposes that *licensing inverse-scope readings and checking scalar implicatures* require *Reference Set Computation* (RSC (2)), which is too taxing for preschoolers. If I assume that young children skip RSC and allow *all* the readings under consideration, the prediction is that English-speaking children with limited processing capacities will allow *some* where the stronger *all* is true, and allow more inverse scope readings where adults hesitate.

(2). *Reference Set Computation:* The reference set consists of pairs <d, i> of derivation and interpretation, and it is motivated by interface needs: A given <d, i> pair is blocked if the same interface effect could be obtained more economically (i.e., there is a more economical <d, i> competitor in the reference set). Reference set computation is triggered *only* by the application of uneconomical procedures. (Reinhart 2005)

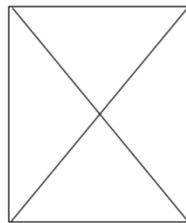
**Method:** A 'covered-box' task was used to test scope assignment and scalar implicatures. Participants need choose a picture based on a sentence. They were trained to choose the hidden picture whenever they thought both visible pictures were incorrect (3).

(3). i. *Scope-assignment:*

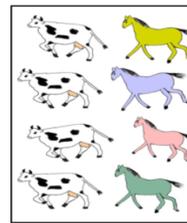
Test sentence: There is a horse chasing every cow.



Distractor



Hidden picture



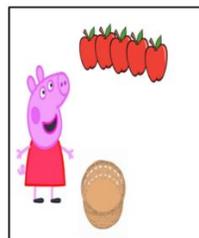
Target (Inverse scope)

ii. *Scalar implicature:*

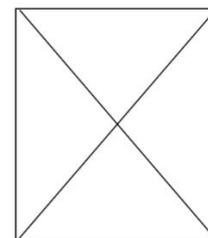
Test sentence: Peppa Pig took some apples.



Target (LU)



Distractor



Hidden picture

i. *Scope assignment:* Two readings were created: surface-scope and inverse-scope. There were 9 items for each reading and 9 fillers in pseudorandom order. Inclusion criteria: a child need be correct on at least 7 of the 9 fillers.

ii. *Scalar implicatures*: The target scale was <some, all>. Two conditions were created: logically-true & pragmatically-informative (LI) (control) and logically-true & pragmatically-underinformative (LU). There were 6 items for each condition and 6 fillers in pseudorandom order. Inclusion criteria: a child need be correct on at least 5 of the 6 control/filler items.

A digit-span task (working memory): Participants need recall a sequence of digits (forwards and backwards) with progressively longer sequences.

Participants: 19 English-speaking children and 14 English-speaking adults participated in all three tasks. Another 6 children only took scalar-implicature study and digit-span test. 16 children (4;00-8;07, mean 4;11) passed scope-assignment screening. 23 children (3;11-9;11, mean 5;08) passed scalar-implicature screening. 14 children (4;00-8;07, mean 4;11) and all adults passed both.

**Results & Discussion**: Children accepted more inverse-scope readings than adults (86.1% and 59.5% respectively, see (4)). Adults showed a significant *inverse* correlation between backward digit-span and acceptance of inverse scope ( $r=-.561, p=.037$ ). Since *almost all* children allowed inverse-scope readings (little variation), the correlation between digit span and acceptance of inverse scope was not found for children. Adults seldom allowed *some* in LU conditions, but children did so frequently (83.3%, see (5)). Children showed a significant *inverse* correlation between digit-span and acceptance of *some* in LU ('forwards':  $r=-.614, p=0.002$ ; 'backwards':  $r=-.755, p<.0001$ ).

(4). *Results for scope assignment*

Participants:	Children (N=16)	Adults (N=14)
Acceptance of surface-scope readings	88.9%	87.6%
Acceptance of inverse-scope readings	<b>86.1%</b>	<b>59.5%</b>

(5). *Results for scalar implicatures*

Participants:	Children (N=23)	Adults (N=14)
Acceptance of 'some' in LI conditions	97.1%	100%
Acceptance of 'some' in LU conditions	<b>83.3%</b>	<b>29.8%</b>

The results for the 14 children passing both screenings are shown in (6). Children showed a significant *inverse* correlation of digit span with acceptance of *some* in LU ('forwards':  $r=-.759, p=.002$ ; 'backwards':  $r=-.764, p=.001$ ). No other significant correlation was found.

In sum, English-speaking children with limited processing capacities allowed more inverse-scope readings than adults and frequently accepted *some* when *all* is true. These findings provide new support for processing accounts for the observed delays in scope assignment and scalar implicatures.

(6). *Results for the 14 children who passed all inclusion criteria*

Acceptance of surface-scope readings	Acceptance of inverse-scope readings	Acceptance of 'some' in LI conditions	Acceptance of 'some' in LU conditions
93.7%	<b>86.5%</b>	95.2%	<b>91.7%</b>

**Selected References:**

Chierchia, G. et al. (2001) The acquisition of disjunction: Evidence for a grammatical view of scalar implicatures. *BUCLD 25*. Lidz, J. (2016) Quantification in child language. *Oxford Handbook of Developmental Linguistics*. Reinhart, T. (1999) The processing cost of reference set computation: Guess patterns in acquisition. *OTS Working Papers in Linguistics*.

**SYNTAX-DRIVEN ASYMMETRIES IN MORPHOLOGICAL DECOMPOSITION  
AS REVEALED BY MASKED VISUAL PRIMING**

Roberto Petrosino & Jon Sprouse – *University of Connecticut*

**INTRODUCTION.** A large body of masked priming studies has shown that, before being visually recognized, words are decomposed in smaller units that seem to correspond to morphemes [1]. This procedure, called *morphological decomposition*, seems to be fairly sophisticated as it occurs in words that are made of more than one morpheme (morphologically transparent words; e.g., *driver*→{*drive*}-{*er*}); but it does not occur in words that contain a root plus a non-affixal ending (e.g., *brothel*↯{*broth*}-{*el*}, where *el* is not an English suffix). The currently most-accepted model of decomposition – here referred to as *decomposition-first model* (1DM) – argues that decomposition occurs *before* accessing the information associated to morphemes and stored in the mental lexicon [2]. In this model, decomposition is therefore *automatic* and is predicted not to be affected by any kind of lexical information. Notice that the term “lexical information” in the last sentence is purposefully vague and may refer to a wide spectrum of linguistic properties – from (a) the phono-orthographic realizations of allomorphs (e.g., *teach* ~ *taught*) to (b) syntactic information (e.g., syntactic categories), to (c) meaning. Previous studies have shown that decomposition is sensitive to (a) and insensitive to (c) [3, 4]. In this paper, we explore whether decomposition may be sensitive to (b), and in particular to affixal syntactic restrictions. By this term, we mean the syntactic categories affixes may selectively attach to. For example, the suffix *-able* only attaches to verbs (e.g., *do-able*, but not *\*bliss-able*), while the suffix *-ness* only attaches to adjectives (e.g., *weak-ness*, but not *\*bliss-ness*). To this end, we elicited priming of syntactically licit words (i.e., *detectable*; experiment 1) and syntactically illicit pseudo-words (*blissable*; experiment 2).

**EXPERIMENT 1** asked whether decomposition equally occurs in bimorphemic words suffixed with different suffixes: *-ful*, *-able*, *-ity*, *-ment*, and *-ness*. A total of six conditions were tested (24 word pairs each): the *identity condition* (e.g., *thumb-THUMB*); the *ful-condition* (e.g., *successful-SUCCESS*); the *able-condition* (e.g., *detectable-DETECT*); the *ity-condition* (e.g., *purity-PURE*); the *ment-condition* (e.g., *settlement-SETTLE*); and the *ness-condition* (e.g., *weakness-WEAK*). One hundred and forty native speakers of American English were recruited through Amazon Mechanical Turk and took the experiment online through PsychoJS [5]. Subjects were asked to perform a lexical decision task on the target. The primes were presented for 34 ms and were preceded by a 500ms-long forward mask ('#####') and followed by the target. Fig. 1 reports mean priming effects and Cohen’s *d* (ES, effect size) for each condition over the bars. Response times for each condition were fitted into a linear mixed-effect regression (LMER) model and a Bayes Factor (BF) model. Each model had *log RT* as the dependent variable, RELATEDNESS (i.e., related/unrelated prime) as fixed factor, and SUBJECT and ITEM as random factors (intercepts only). Fig. 1 reports *p*- and *BF*<sub>1,0</sub>-values for each condition under the bars. All conditions showed significant priming effects in the LMER analysis (*ps*<.05). For the interpretation of the BFs, we refer to Jeffreys (1961)’s interpretive table. In the BF analysis, only a subset of them substantially supported the alternative hypothesis (*BF*<sub>1,0</sub>*s*>3): identity, *ful*-, *able*-, and *ness*-conditions; the *ity*- and the *ment*-conditions anecdotally supported the null hypothesis (*1*>*BF*<sub>1,0</sub>*s*>0.33). None of the conditions were significantly different from the other (Dunn-corrected *ps*=1).

**EXPERIMENT 2** asked whether decomposition is dependent on the syntactic restrictions associated to each of the affixes tested above (Table 1). More concretely, it asked whether syntactically illicit forms like *\*blissable*, *\*blissity*, *\*blissment*, *\*blissness* decompose as much as *blissful* does. One hundred and two bare noun roots were paired with the corresponding syntactically illicit affixed primes, and were Latin-Square arranged across five conditions (17 pairs each): *identity condition* (e.g., *skit-SKIT*); *able-condition* (e.g., *skitable-SKIT*); *ity-condition* (e.g., *skitivity-SKIT*); *ment-condition* (e.g., *skitment-SKIT*); *ness-condition* (e.g., *skitness-SKIT*); *unrelated condition* (e.g., *trainee-SKIT*). A *ful-condition* was added

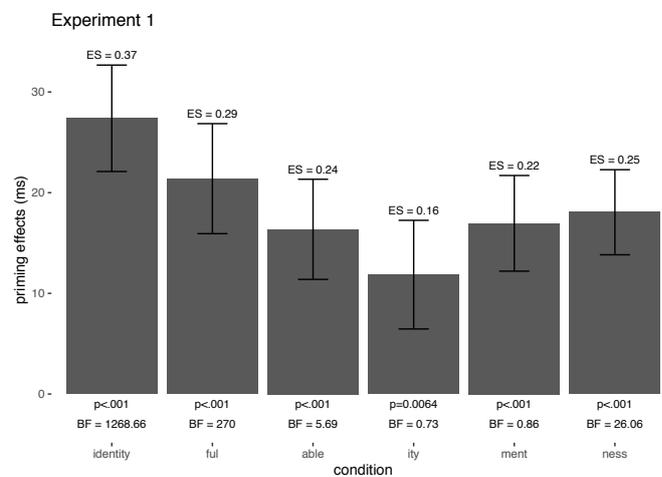


Figure 1

and were Latin-Square arranged across five conditions (17 pairs each): *identity condition* (e.g., *skit-SKIT*); *able-condition* (e.g., *skitable-SKIT*); *ity-condition* (e.g., *skitivity-SKIT*); *ment-condition* (e.g., *skitment-SKIT*); *ness-condition* (e.g., *skitness-SKIT*); *unrelated condition* (e.g., *trainee-SKIT*). A *ful-condition* was added

AFFIX	ATTACHES TO...	YIELDS...
<i>able</i>	V	A
<i>ity</i>	A	N
<i>ment</i>	V	N
<i>ness</i>	A	N

Table 1

as the licit condition (separate from the Latin-Square design above) to compare the illicit affixed conditions to. In this condition, 17 pairs had their targets preceded by a licit *ful*-affixed derived prime (e.g., *graceful-GRACE*) and 17 pairs had their targets preceded by an unrelated, bimorphemic prime (e.g., *thrower-GRACE*). One hundred and forty native speakers of American English participated in the experiment. Response times were collected and analyzed in the same way as described for experiment 1. Fig. 2 reports mean priming effects and Cohen’s *d* (ES, effect size) over the bars; the *p*- and *BF*-values are under the bars. In the LMER analysis, all conditions showed significant priming effects ( $ps < .001$ ). In the BF analysis, only a subset of them substantially supported the alternative hypothesis ( $BF_{1,0} > 3$ ): namely, the identity and the *ful*-condition (licit), and the *ity*- and *ness*-conditions; the *able*- and *ment*-conditions were instead found to substantially support the null hypothesis ( $BF_{1,0} > 0.33$ ). Dunn-corrected post-hoc analysis revealed that the comparisons involving the identity or the *ful*-conditions, and the *able*- or the *ment*-conditions were significant ( $ps < .05$ ,  $BF_{1,0} > 3$ ).

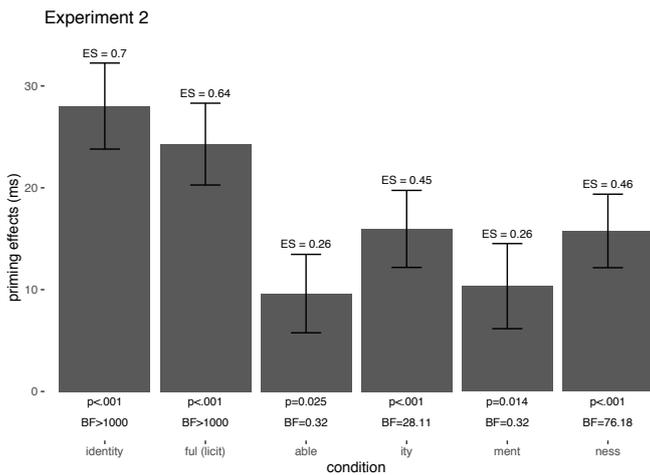


Figure 2

**GENERAL DISCUSSION.** The ultimate goal of the two experiments above was to test the sensitivity of decomposition to syntactic affixal restrictions. The decomposition-first model expects them not to affect decomposition at all, as they are stored in the lexicon. The results of the two experiments described above are summarized in Table 2 below. Experiment 1 asked whether decomposition depends on the suffix used in each of the five conditions tested: *-ful*, *-able*, *-ity*, *-ment*, and *-ness*. In our statistical analyses, the BF calculations revealed that the *ful*-, *able*-, and *ness*-conditions strongly supported the alternative hypothesis, whereas the *ity*- and *ment*-condition did not. This asymmetry might be explained as an effect of affix productivity: the suffixes *-ity* and *-ment* are indeed considered less productive than the suffixes *-ful*, *-able*, and *-ness* [6]. Experiment 2 asked whether decomposition is affected by violations to the affixal syntactic restrictions

defined in Table 1. The BF analysis suggested that the *ity*- and *ness*-conditions strongly supported the alternative hypothesis, while the *able*- and *ment*-conditions strongly supported the null hypothesis. Comparison of the effects across the two experiments seems to suggest that decomposition of the suffix *-able* was affected by presentation of illicit suffixed forms, whereas decomposition of the suffixes *-ness* was not. The lack of priming in the *ity*- and *ment*-conditions of experiment 1 prevents us from properly interpreting the effects elicited for the same conditions in experiment 2.

CONDITION	EXPERIMENT 1		EXPERIMENT 2	
	<i>example</i>	<i>priming?</i>	<i>example</i>	<i>priming?</i>
able	<i>detectable-DETECT</i>	✓	<i>skittable-SKIT</i>	✗
ity	<i>purity-PURE</i>	✗	<i>skittity-SKIT</i>	✓
ment	<i>settlement-SETTLE</i>	✗	<i>skitment-SKIT</i>	✗
ness	<i>weakness-WEAK</i>	✓	<i>skitness-SKIT</i>	✓

Table 2

**CONCLUSIONS.** Though the results are hard to be fully interpreted in terms of the questions being asked, two theoretically interesting conclusions can be drawn. First, decomposition seems to occur for some affixes only. This does not necessarily challenge decomposition-first models; rather, it may suggest that decomposition only occurs for a subset of affixes that are available for decomposition without any reference to the lexicon. We hypothesized that only productive affixes are available for decomposition. In this sense, *-ful*, *-able*, and *-ness* are available for decomposition because highly productive, whereas *-ity* and *-ment* are not available for decomposition because they are not as productive. Second, decomposition seems to be affected by syntactic restrictions of a subset of the affixes available for decomposition: in particular, decomposition of *-able* seems to be affected if its syntactic restriction is violated, but decomposition of *-ness* does not. Further testing is needed to validate these tentative conclusions and clarify the potential contribution of syntactic information onto early decomposition.

## ***De re* interpretation in belief reports--An experimental investigation**

Yuhan Zhang & Kathryn Davidson

Department of Linguistics, Harvard University

### **Introduction**

The subject of the sentence *In 1978, the president built a huge skyscraper in NYC* can, under a *de re* reading refer to the current president, or under a *de dicto* reading can refer to the holder of the presidential role within the temporal operator (*in 1978*). Despite the variety of theoretical approaches to formalize this distinction in natural language semantics (Cresswell & von Stechow, 1982; Keshet, 2008; Pearson, Forthcoming; Percus, 2000; von Fintel & Heim, 2011; see a summary in Keshet & Schwarz, 2019), it is generally accepted that both interpretations should be generated by the grammar, although in some contexts one reading may be argued to be strongly inaccessible, as, e.g. argued by Nelson (2019) for the lack of *de re* interpretation in (1).

(1) Sally *believes* that *her brother* is happy. (Nelson, 2019: supplement (3))

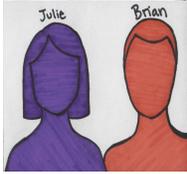
[supposed *de re* context: Sally observes someone laughing, not knowing he is her brother.]

The presumed availability of both *de re* and *de dicto* readings under intensional operators sits at the center of intensional semantics, yet as far as we know, there has yet to be any quantitative investigation of these readings and the factors affecting their accessibility. Following the literature, we expect the grammar to generate both readings with additional heuristics resolving this ambiguity in conversation; a key question then becomes whether contexts can be devised in which one reading is, in fact, *unavailable* or *inaccessible*. Here, we test the availability and accessibility of *de re* readings under two different intensional operators (both attitude embedding verbs: *believe*, *say*) and two different QUDs (*real world facts*, *subject's mental state*) to lay the groundwork for understanding factors that go into the ambiguity, and occasional lack thereof.

### **Experiment**

Participants (128 native speakers of English,  $\bar{X}_{age} = 38.2$ ,  $N(\text{male}) = 76$ ) were recruited through Amazon Mechanical Turk to participate in an agreement rating task with three mistaken identity stories (inspired by Marti 2006)(Table 1). Our contexts vary in **whether the holder of the belief report can identify the entity in the real world** under a particular description (e.g. “the poem that Nina wrote” in Table 1): if so, our sentences make both *de re* and *de dicto* readings true; however, if the belief holder cannot identify the entity based on the description, then the sentence is only true with a *de re* interpretation. This basic context manipulation was between subjects (with all three stories the same value for a given participant). Second, since *de re* interpretation is evaluated in the real world and *de dicto* in a belief world, we created an explicit question to which the target sentence is an answer such that **the question asks about either real world facts or the status of the belief world**. This manipulation was also between subjects and we hypothesize that an interaction effect should be present -- in our condition where the sentence only permits *de re* reading but the belief world is the emphasis of QUD, the agreement rating would be lower. Third, we compared between subjects **the use of *believe* and *say* verbs** in terms of this *de re/de dicto* accessibility, predicting that since *say* prefers literal report and the description of the target entity is changed in the report, *say* would have overall lower agreement rating. In the end we have a 2x2x2 entirely between-subjects design, with a dependent measure in which participants were asked for agreement with a statement (i.e. ‘Do you agree with Brian’s last statement according to this story?’), then required to shift a slider scale starting from numerical value 50 to somewhere between ‘highly disagree’ (0) and ‘highly agree’ (100). After incorrect responses to the comprehension question were deleted, 377 trials remained for analysis.

Table 1: Schema of test conditions within the *poem* example

	<p><b>[CONTEXT]</b>          Julie is one of several judges of a poetry competition and Brian is the editor of the poetry collection. Julie encounters a fabulous poem (without knowing who wrote it) and tells Brian: “This poem is awesome. I believe it is going to win the prize.” Brian looks into the submission pile and discovers that Nina was the author.</p>	
<p><b>[CONDITION: <i>de re</i> ONLY]</b>          However, Brian doesn't share his discovery with Julie about the authorship of this poem. So Julie doesn't know who wrote this promising poem.</p>	<p><b>[CONDITION: <i>de re</i> &amp; <i>de dicto</i>]</b>          Brian shares his discovery with Julie about the authorship of this poem. So Julie knows that it was Nina who wrote this promising poem.</p>	
<p><b>[CONTEXT DIALOGUE]</b>          Jim: “Do you know which poem will win the prize?”          Brian: “I think they are still finalizing the result, but Julie had some idea.”</p>		
<p><b>[CONDITION: QUD-real world emphasis]</b>          Jim: “So which one will win?”</p>	<p><b>[CONDITION: QUD-belief world emphasis]</b>          Jim: “So what does Julie think?”</p>	
<p><b>[CRITICAL SENTENCE expected to have different readings under different conditions]</b>          Brian: “Well, she believes/said that the poem that Nina wrote will win the prize.”</p>		

### Data Analysis & Result

We analyzed two features from the data: (1) the direction to which the slider moved indicates the basic category of agreement or disagreement; (2) the distance from the center indicates the degree to which participants agreed or disagreed. The categorical direction data indicates that all conditions elicit more agreement than disagreement, supporting the general availability of *de re* interpretation (Fig.1). In the *believe* verb condition with only *de re* interpretation and a belief-world emphasis, people were significantly less likely to agree than when both *de re* and *de dicto* are set to be true (Fisher’s exact test: odds ratio = 0.122,  $p^{***} < .001$ ; odds ratio = 0.319,  $p^* < .05$ ), presenting a *de re* reading with decreased inaccessibility. Furthermore, *believe* elicited more agreement than *say* in the *de dico* + *de re* condition with real-world emphasis (odds ratio = 4.583,  $p^* < .05$ ). This cross-verb difference is also reflected through the distribution of the distance from the center (Fig.2), as the agreement degree of *say* is more scattered across the scale where *believe* tends to aggregate on the ‘highly agree’ end, suggesting that accessing *de re* is harder for propositions with *say* than those with *believe*.

### Conclusion

We analyzed participants’ judgments of sentences under contexts which supported both *de re* and *de dicto* readings versus contexts which only supported *de re* readings. We found *de re* to be more inaccessible when belief states were highlighted by varying the question under discussion, suggesting at least some contextual factors could be experimentally controlled for specific interpretation. Since *believe* and *say* also behaves differently, our immediate next steps include testing the accessibility of *de re* and *de dicto* interpretations mediated by other intensional operators such as temporal phrases, for a better understanding of principles that affect these readings in the intensional domain.

Figure 1: Percentage of agreement under 2x2x2 design (with Fisher's exact test score)

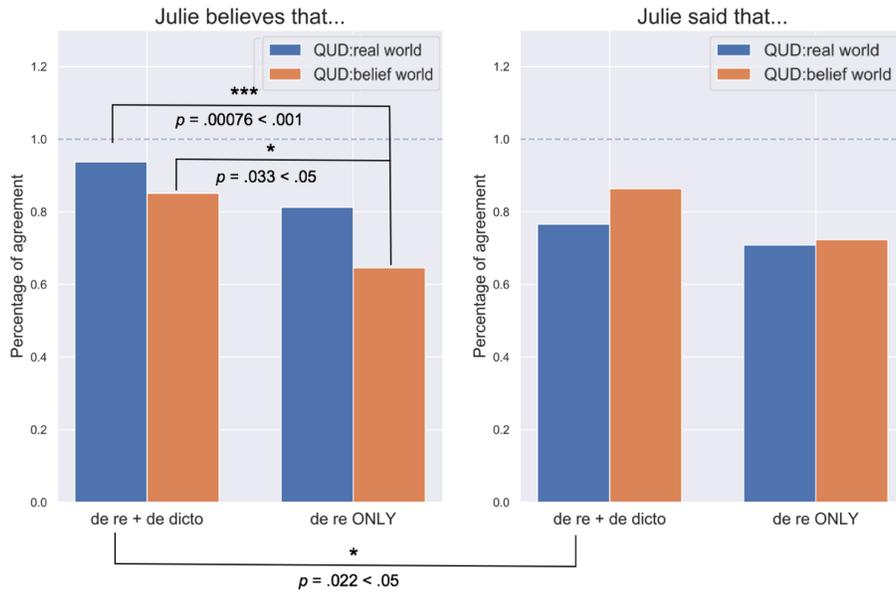
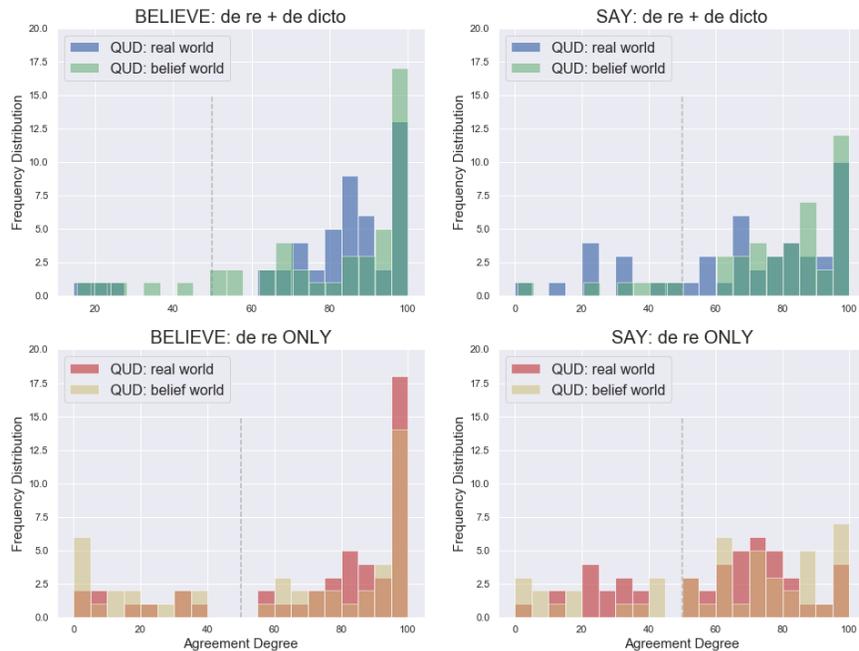


Figure 2: Distribution of distance from center under 2x2x2 design



References [1]. Cresswell & von Stechow. (1982). De re belief generalized. [2]. Keshet. (2008). *Good intensions: Paving two roads to a theory of the de re / de dicto distinction*. [3]. Keshet & Schwarz. (2019). De Re/De Dicto. [4] Marti, L. (2006). Restoring Indefinites to Normalcy: An Experimental Study on the Scope of Spanish algunos. [5]. Nelson, M. (2019). Propositional Attitude Reports. [6]. Pearson, H. (Forthcoming). Attitude verbs. [7]. Percus, O. (2000). Constraints on Some Other Variables in Syntax. [8]. von Fintel, K., & Heim, I. (2011). *Intensional Semantics*.