



The representation and processing of distributivity and collectivity: ambiguity vs. underspecification

JAKUB DOTLAČIL 

ADRIAN BRASOVEANU 

**Author affiliations can be found in the back matter of this article*

RESEARCH

][ubiquity press

ABSTRACT

Sentences with plural expressions can receive at least two interpretations. For example, the sentence *The boys hold a balloon* could mean that the boys as a group jointly hold one balloon (the collective reading) or that each boy holds one balloon, which would imply that as many balloons were held as there are boys (the distributive reading). Building on Frazier et al. (1999), we show that the human processor favors collective readings. Crucially, the preference for collective readings is only observed when the distributive reading has to be established through the means of phrasal distributivity (e.g., triggered by distributive quantifiers), and the preference disappears in case of lexical distributivity (e.g., the distributive interpretation of *win*). The findings provide evidence for different mental representations of the two types of distributivity and shed light on why the processor exhibits a default preference for collective interpretations.

CORRESPONDING AUTHOR:

Jakub Dotlačil

Utrecht University, Trans 10,
3512JK Utrecht, Netherlands

j.dotlaci@gmail.com

KEYWORDS:

semantic ambiguity in processing; pluralities; distributivity; collective interpretation; null operators

TO CITE THIS ARTICLE:

Dotlačil, Jakub and Adrian Brasoveanu. 2021. The representation and processing of distributivity and collectivity: ambiguity vs. underspecification. *Glossa: a journal of general linguistics* 6(1): 6. 1–22. DOI: <https://doi.org/10.5334/gjgl.1131>

Plural expressions pose a challenge to communication. Consider, for instance, a plural noun phrase like *my classmates* and its role in the following two sentences:

- (1) a. My classmates had a coffee.
b. My classmates surrounded the podium.

At first glance, it might seem that the interpretation of the noun phrase and the interpretation of the sentences in which the noun phrase appears are straightforward. The plural expression *my classmates* simply refers to those individuals that are my classmates, and the sentences are statements about those individuals. However, on closer inspection, (1a) and (1b) turn out to communicate very different statements. The predicate *had a coffee* in (1a) assigns a property to individual classmates: the sentence states that each of them had one coffee. In contrast to (1a), the predicate *surrounded the podium* does not assign a property to individual classmates, but to the whole group: (1b) makes a statement about my classmates as a group to the effect that together, they surrounded the podium.

Examples (1a) and (1b) include predicates that clearly distinguish between the two roles that the plurality *my classmates* plays in interpretation, but this is often not the case in regular conversation. For example, the sentence in (2) below makes two mutually exclusive statements. It might express that individual items cost \$50, in which case all the items as a group cost more than \$50, or it might express that all the items as a group cost \$50, in which case the individual items cost less than that. Following the formal semantics literature, we will henceforth refer to the first interpretation as *distributive* and the second interpretation as *collective* (see Landman 1995; Nouwen 2012; Dotlačil to appear; Champollion to appear and references therein for more details on the two interpretations).

- (2) The items cost exactly \$50.

The fact that sentences with plural noun phrases can be understood in two very different ways raises a non-trivial issue about alignment between production and comprehension. How do speakers and hearers converge on one interpretation when talking about multiple entities? And how are the two interpretations represented in grammar?

In semantic research, the dominant account of the distributive/collective ambiguity assumes that the two interpretations are structurally different. It is assumed that language has a silent operator, *D*, which in most instances can be paraphrased as ‘each’. When the operator is present, the resulting sentence receives the distributive reading, and when it is not, the sentence receives the collective reading (Roberts 1990, Lasersohn 1995; see Schwarzschild 1996 and Landman 2000 for a more nuanced picture). This position has been supported in psycholinguistic research, in particular, by findings in the processing of distributive and collective interpretations presented in Frazier et al. (1999).¹

In the study of Frazier et al. (1999), participants read sentences with a conjoined noun phrase subject and whose verb+object predicates were compatible with collective and distributive interpretations. The predicates were disambiguated towards their collective interpretation (by the adverb *together*) or their distributive interpretation (by the quantifier *each*). When the disambiguator followed the predicate, as in (3), encountering the distributive disambiguator (*each*) yielded more processing difficulties than encountering the collective disambiguator (*together*). Participants read the underlined region more slowly in (3a) than in (3b), and they regressed more often in (3a) than in (3b) from the same region.

- (3) a. Sam and Maria carried one suitcase each at the airport.
b. Sam and Maria carried one suitcase together at the airport.

Importantly, when the disambiguators preceded the predicate, (4), no contrast between *each* and *together* was observed. That is, (4a) did not show processing difficulties compared to (4b).

¹ See also Brooks and Braine (1996); Pagliarini et al. (2012); Dobrovie-Sorin et al. (2016) for offline studies of distributive and collective interpretations.

This excluded the possibility that processing difficulties observed in (3) could simply be due to the fact that distributive readings are more complex in their interpretation.

- (4) a. Sam and Maria each carried one suitcase at the airport.
 b. Sam and Maria together carried one suitcase at the airport.

To explain the data, Frazier et al. (1999) propose that distributive/collective readings are a case of ambiguity, similar to the kind of lexical ambiguity exemplified by homonyms. Assuming that the processor has to commit to one interpretation when it encounters an ambiguity, and that the processor prefers collective readings of ambiguous predicates, readers will run into problems in the post-predicate disambiguation towards distributivity, (3a). No such problems are expected for (3b). Furthermore, since the sentences in (4) are disambiguated before the predicate is read and the processor has to commit to a reading, (4) should not cause difficulties either. This explanation treats the contrasts in (3) and (4) as a kind of semantic garden-path effect, and fits well with the dominant position in semantic research: the effect arises because readers by default assume the absence of the *D* operator when reading (3), which is incompatible with the later disambiguation towards the distributive reading in (3a).

One issue that remains open and that we address in this paper is why readers prefer the collective reading of predicates like *carried one suitcase* (see also Boylan et al. 2011 for additional support that this preference exists). Frazier et al. (1999) mention two possible explanations, but note that their experiment cannot decide between them. On one hand, it is possible that the processor selects the structure/representation that is the simplest, while still compatible with the processed words.² Since the distributive reading requires the presence of an extra silent operator *D*, the processor picks the collective reading if possible. Under this view, the preference for collective readings is driven by *structural considerations*. Alternatively, it is possible that readers select the reading that is simpler in the sense of introducing fewer events and entities.³ Consider the sentence *Sam and Martha carried one suitcase*. The distributive reading requires the existence of two events of carrying one suitcase, and it is compatible with there being two suitcases in total. In contrast, only one suitcase and one event of carrying one suitcase is assumed in the collective reading. Under this view, the preference for collective readings is driven by *interpretational complexity considerations*.

We add a third possible explanation for the pattern in (3) and (4). Frazier et al. (1999) treat the two instances of *together* in (3b) and (4b) as identical lexical items that are syntactically integrated in more or less the same way, and this also holds for the two instances of *each* in (3a) and (4a). The adverb *together* is, indeed, commonly treated as a manner adverb irrespective of whether it appears in a pre-VP or a post-VP position.⁴ The situation for *each* is less straightforward. The linguistic literature takes the postverbal *each* in (3a) and the preverbal *each* in (4a) to be related but nonetheless distinct items, each associated with a specific syntax and semantics. Preverbal *each* is usually taken to be a VP modifier (much like *together*), but postverbal *each*, a.k.a. binominal *each*, is syntactically part of the NP it modifies (Burzio 1986, Safir and Stowell 1988). The two quantifiers also differ in their semantics. Postverbal, but not preverbal, *each* places a restriction on the meaning of the direct object, requiring it to express cardinality. This restriction is the reason for the unacceptability of (5a) below: the direct object NP *every movie* is a universal quantifier and does not express cardinality, unlike *one movie* in (5b) (Szabolcsi 2011, a.o.). The sentence in (5c) shows that preverbal *each* places no such restrictions on the object. Differences like the one shown in (5) led semanticists to analyze postverbal *each* differently from preverbal *each* (see Zimmermann 2002; Blaheta 2003; Dotlačil 2012; Champollion 2016b).

- (5) a. *The men saw every movie each.
 b. The men saw one movie each.
 c. The men each saw every movie.

² This parsing strategy, dubbed the Principle of Minimal Attachment, is one of the guiding principles of the Garden-Path Model (Frazier 1978).

³ Frazier et al. (1999) argue that this explanation could follow from the Referential Theory of Crain and Steedman (1985).

⁴ However, even *together* differs slightly in meaning depending on whether it appears pre-verbally or post-verbally, as discussed in Schwarzschild (1993).

With this in mind, let us again consider the argument in Frazier et al. (1999). The observed slowdown for postverbal *each* compared to *together*, as opposed to no difference between preverbal *each* and *together*, was taken to reveal the preference of the parser for the collective interpretation of the ambiguous predicate, followed by reinterpretation in (3a). However, since the two types of *each* have different syntactic and semantic characteristics, the slowdown observed in (3a) could be attributed to the syntactic and semantic differences between preverbal and postverbal *each*. For example, it could be attributed to the fact that readers have to consider the extra semantic requirements that postverbal *each* places on the object. That is, the observed slowdown could be due to the *complexity of postverbal each*.

We present two novel self-paced reading experiments whose goal is to replicate the findings of Frazier et al. (1999) with distributive disambiguators other than *each*, and to explain why readers exhibit a preference for collective readings.

In the first experiment, we follow the design of Frazier et al. (1999), but instead of using *each* as the disambiguator of distributive readings, we employ *individually*, which is a manner adverb and should not show a distinct behavior in its preverbal and postverbal positions compared to other manner adverbs like *together*. This is supported by (6), which shows that, in contrast to *each*, the position of *individually* is not sensitive to the type of object. Furthermore, (7) shows that *individually*, unlike *each*, can be conjoined with *together*, corroborating the claim that *individually* and *together* are both manner adverbs.

- (6) a. The women have (each) built the rafts (*each).
b. The women have (individually) built the rafts (individually).
- (7) a. *The women built rafts each and together.
b. The women built rafts individually and together.

Given the close affinity of *individually* to *together*, the effect of *individually* in the two positions can be directly compared to the effect of *together*. In other words, using only adverbs as disambiguators minimizes the risk of confounds in the current study, and should eliminate the possibility that the findings of Frazier et al. (1999) are due to a greater complexity of postverbal *each* compared to preverbal *each*.

In the second experiment, we leverage the observation made in the semantic literature (Link 1987; Roberts 1990; Lasersohn 1989; 1995; Moltmann 1997; Winter 2000; 2001; Kratzer 2008; Dotlačil 2010; Champollion 2010; 2016a; 2017; de Vries 2017, among others) that the *D* operator is only needed for non-lexical (syntactic) predicates. As an example, take (8) below. The predicate *smile* expresses a property that holds of single individuals, that is, (8) states that every boy smiled. However, there is no reason to assume that this reading is generated by means of a *D* operator. It is simply part of our lexical knowledge that smiling is a property that holds of single individuals. When we hear (8), we learn that a group of boys has the property of smiling, but because of this inherent meaning of *smile*, which is roughly ‘the individual’s lips move upwards’, it follows that whenever smiling is truthfully predicated of a group, it must be truthfully predicated of all (or at least a majority of) the individuals in that group (see also Krifka 1989; Yoon 1996; Landman 2000; Kratzer 2008, a.o., for more details).

- (8) The boys smiled.

The situation is different for syntactically complex predicates like *carried one suitcase* in (9) below. Even if we assume that carrying is a property of atomic individuals, we would only arrive at the meaning that each boy did the carrying action, and the boys all carried one (and the same) suitcase. The issue is the quantifier *one suitcase*: unless we assume the presence of a *D* operator, we do not have a way to derive that there were multiple suitcases, one per boy. That is, without a *D* operator, (9) relates a group of boys and one suitcase by means of the predicate *carry*. The *carry* relation might distribute down to individual boys in the group, but it still relates each of them to the one suitcase under discussion.

- (9) The boys carried one suitcase.

However, we can get a distributive reading for (9) if we postulate the presence of a covert *D* operator. Recall that *D* can be roughly interpreted as *each*, and note that the paraphrase *the boys each carried one suitcase* does express that multiple suitcases were carried, one per boy.⁵

The second experiment will study whether distributive/collective reinterpretation is processed differently when the predicate is lexical, as in (8), and when it is constructed in syntax, as in (9), whose distributive reading requires a quantificational element, the *D*-operator. As we will see, there is indeed a contrast between cases in which distributivity can be purely lexical and cases in which distributivity has to be phrasal and triggered by a quantificational element. We take that as evidence that the preference for collective readings is not due to the interpretational simplicity of collective readings, and argue that our findings are compatible with the position that the processor chooses simpler structures, i.e., structures without the *D* operator.

2 EXPERIMENT 1

The design of this experiment followed Frazier et al. (1999). We investigated the processing of sentences that had a plural entity in subject position, using plural definites, e.g., *the students*. Each sentence included either the adverb *individually*, which forces a distributive interpretation,⁶ as in (10a) and (10c) below, or the adverb *together*, which forces a collective reading,⁷ as in (10b) and (10d). The second manipulation was the position of these adverbs: they appeared either preverbally/early, as in (10a) and (10b), or postverbally/late, as in (10c) and (10d).

- (10) a. The girls individually wrote a sonnet after they had read Shakespeare.
 (EARLY, INDIVIDUALLY)
- b. The girls together wrote a sonnet after they had read Shakespeare.
 (EARLY, TOGETHER)
- c. The girls wrote a sonnet individually after they had read Shakespeare.
 (LATE, INDIVIDUALLY)
- d. The girls wrote a sonnet together after they had read Shakespeare.
 (LATE, TOGETHER)

The logic behind this experiment was identical to that of Frazier et al. (1999). If readers non-randomly select one reading when encountering a plural NP + predication, and if a subsequent reinterpretation is costly, we expect to observe processing difficulties either in (10c) or (10d), once the disambiguating adverb following the predicate is read and integrated.

Frazier et al. (1999) observed that distributive reinterpretation was costly, and accounted for this processing difficulty by assuming that the processor initially selected the collective reading. By the same reasoning, we would expect to see processing difficulties in (10c) as compared to (10d). If, however, (10d) turns out to be more difficult than (10c), we can conclude that the processor selects the distributive interpretation by default, and the reinterpretation towards a collective reading is what incurs an additional processing cost.

The early cases, (10a) and (10b), control that the processing cost is due to reinterpretation, and it is not caused by the inherent costs associated with distributive or collective readings *simpliciter*.

⁵ For a good formal explanation why we assume this contrast between lexical predicates as in (8), which do not require the *D* operator, and phrases/syntactically-complex predicates, which do, see Champollion (2019) and references therein. To be sure, it does not follow that all multi-word expressions require the *D* operator for their distributive reading, since it is possible that deriving the meaning of some multi-word expressions is not necessarily syntactically mediated. We will come back to this point in the General Discussion section (§4).

⁶ See Moltmann (2005) for the semantics of *individually*. As Moltmann (2005) shows in detail, when *individually* modifies predicates specifying spatio-temporal locations, it also has an irrelevant, spatial-separation reading. We used predicates that exclude this interpretation of *individually*.

⁷ See especially Schwarzschild (1993), Laserson (1995) and Kratzer (2008) for the semantics of *together*. Strictly speaking, *together* does not force a collective reading, since it is compatible with both a collective and a cumulative reading. But this finer distinction between non-distributive readings is not relevant for our purposes. The only important point is that the distributive reading is excluded. The fact that *together* excludes the distributive reading has been experimentally confirmed (for adults) in Syrett and Musolino (2013).

2.1 METHOD

2.1.1 Participants

The participants were 87 undergraduate students at University of California Santa Cruz, all self-identified native speakers of English. They received course (extra-)credit for their participation.

2.1.2 Procedure and items

Participants received a link to participate in the self-paced reading study; we used a non-cumulative moving window self-paced reading paradigm (Just et al. 1982).⁸ The experiment was run on a UCSC hosted installation of the IBEX platform. Participants took the experiment online.

After clicking on the link, each participant first saw general instructions about the self-paced reading procedure and identified whether he/she is a native speaker of English (non-native speakers could participate and receive (extra-)credit so that there would be no incentive to lie). After general instructions, participants could familiarize themselves with the methodology on three practice items. The practice session was followed by the experiment.

The experiment consisted of 40 fillers and 28 target items. An example of an experimental item in all four conditions is given in (10); see the Appendix for the list of all items. Each item consisted of four conditions. The items were rotated through the four conditions in four lists, using the standard counterbalanced Latin square design. Each participant was assigned to one of the lists; the order of the stimuli was randomized for every participant. Every item and every filler was followed by a yes-no comprehension question. Every question checked whether participants paid attention to the previous dashed sentence. For every item, the questions were identical across all four conditions. The questions were unambiguous and had a clear correct answer. For example, the question following (10) was:

(11) Did the girls read Dickens?

When participants answered incorrectly, they were notified of their mistake. Yes and no responses were distributed roughly evenly (33 out of 68 questions had a yes-response).

2.2 RESULTS AND DISCUSSION

As expected, responses to questions did not pose significant problems to participants. The median correct response to fillers and target items was 95% correct, with 85 participants having at least 88% of responses correct. Two participants, however, were clear outliers, answering only 79% and 82% of the questions correctly. These participants were removed from all subsequent analyses. The final number of participants whose data we analyzed: 85.

Prior to the analysis of reading times (RTs), we removed extremely fast (<50 ms) and slow (>3000 ms) responses, as is common in analyses of self-paced reading studies (see, for example, Futrell et al. 2018). Less than 1 percent of all the data was eliminated in this way. Secondly, we log-transformed RTs to mitigate their characteristic right-skewness. We probed for (very fast or very slow) outliers among readers by checking mean logRTs per participant. No participant's mean logRT was more than 2.5 standard deviations away from the grand mean, so we retained the RT data from all 85 participants.

Following Trueswell et al. (1994), among others, we factored out the influence of word length and word position by running a linear mixed-effects regression that had intercept-only random effects for subjects and two fixed effects – word length in characters and word position in the sentence. The resulting residualized log reading times were used for all subsequent analyses.

The main regions of interest (ROIs) were the adverb, the predicate and the three words following the disambiguated predicate (the spillover region). The ROIs are underlined in (12).

- (12) a. The girls individually wrote a sonnet after they had read Shakespeare.
b. The girls together wrote a sonnet after they had read Shakespeare.
c. The girls wrote a sonnet individually after they had read Shakespeare.
d. The girls wrote a sonnet together after they had read Shakespeare.

⁸ Participants read the sentence word by word with all the words initially hidden (dashes of the appropriate length are displayed where the words should be). Each press of the SPACE bar reveals the next word and hides the preceding one.

The graphical summary of logRTs starting with adverbs in the early position is presented in **Figure 1**. LogRTs are also summarized in a table in Appendix.

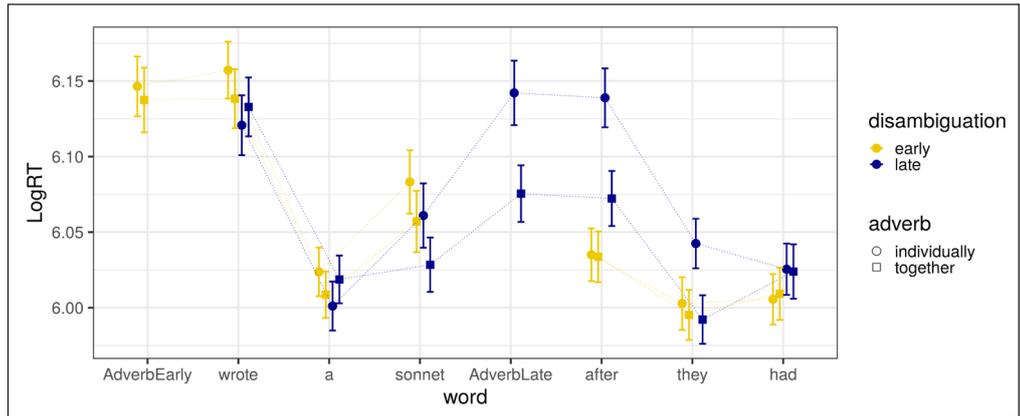


Figure 1 Exp. 1, LogRTs per word in the predicate, adverb and spillover regions. The position AdverbEarly shows LogRTs for adverbs *individually* and *together* in the early position. The position AdverbLate shows LogRTs for adverbs *individually* and *together* in the late position.

For each region, we constructed a mixed-effects model with three treatment-coded fixed effects: ADVERB (TOGETHER OR INDIVIDUALLY; the former was the reference level), POSITION (EARLY OR LATE, the former was the reference level); the interaction of ADVERB and POSITION. Following one of the recommendations in Barr et al. 2013, all our models included the maximal random-effect structure for subject and items. Since the `lme4` R package, which is commonly used for estimating mixed-effects models, did not converge with the maximal random-effect structure, we used Bayesian models estimated with Stan (Carpenter et al. 2017) and `brms` (Bürkner 2017). The structure of prior distributions and the sampling details of the models are provided in the Appendix.

In **Figure 2**, we show the posterior distributions of the fixed effects on two words, the adverb and the first word in the spillover region. We focus on these two words because these were the only cases in which the 89% credible intervals of at least one parameter exclude 0. On the adverb, we see that in the early position *individually* decreases residualized logRTs (posterior distribution of INDIVIDUALLY: 89% credible interval: $\langle -0.08, -0.014 \rangle$, median = -0.05 , $p(\beta < 0) = 0.98$).^{9,10} We also observe an interaction of LATE POSITION:INDIVIDUALLY, which almost fully removes the speedup due to *individually* (posterior distribution of INDIVIDUALLY:LATE POSITION: 89% credible interval: $\langle 0.005, 0.12 \rangle$, median = 0.05 , $p(\beta > 0) = 0.96$). On the first spillover word, the word *after*, we see that LATE POSITION leads to a slowdown in reading (posterior distribution of LATE POSITION: 89% credible interval: $\langle 0.01, 0.08 \rangle$, median = 0.04 , $p(\beta > 0) = 0.98$). Second, there is a positive

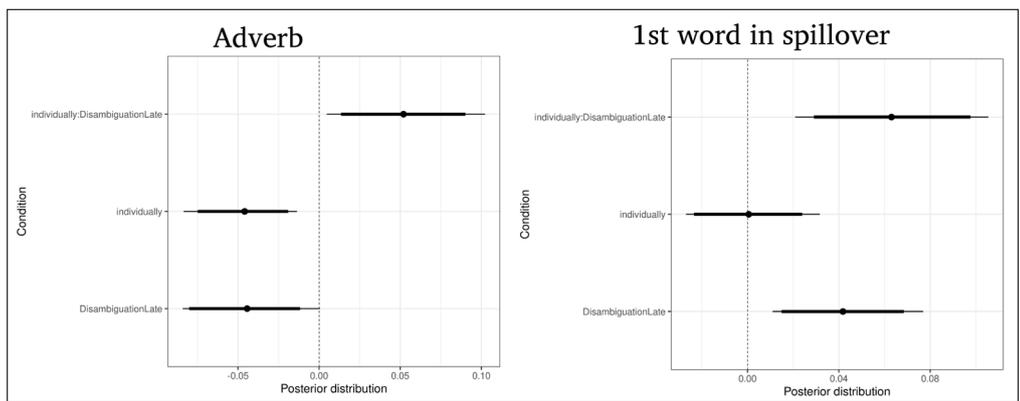


Figure 2 Posterior distributions of parameters of interest on Adverb and the first word in the spillover (the word *after* in (12)). A thick horizontal line represents 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents 89% HDPI of a parameter. The dashed vertical line is at the position of 0. Filled circles represent median values.

⁹ The 89% credible interval specifies the interval in which the parameter falls with the probability of 89%. $p(\beta < 0) = x$ specifies the probability that the parameter is negative is x (positive if we consider $p(\beta > 0) = x$). We can inspect all distributions but are particularly interested in posterior distributions that are predominantly positive or negative (i.e., $p(\beta > 0)$ or $p(\beta < 0)$ is 0.95 or higher, the credible intervals do not span zero). We present the 89% credible interval, rather than 95% credible intervals, following the suggestions and conventions proposed in Kruschke (2014) and McElreath (2020).

¹⁰ Based on the graph in Figure 1, this finding might look surprising since in the early position, *individually* and *together* show almost the same logRTs. However, the first adverb is longer by 4 letters so differing lengths mask the effect. The effect is observable when we consider residualized log-reading times. When we examine the model that uses raw (non-residualized) logRTs as the dependent variable, we do not observe any speed-up due to adverb type.

interaction of LATE POSITION and INDIVIDUALLY (posterior distribution of LATE POSITION:INDIVIDUALLY: 89% credible interval: $\langle 0.02, 0.11 \rangle$, median = 0.063, $p(\beta > 0) = 0.99$). No other region of interest shows a posterior distribution of a parameter that is overwhelmingly ($p(\beta) > 0.95$) positive or negative.

Focusing on the effect observed on the spillover, we see that postverbal adverbs lead to a slowdown in reading, and the slowdown is even more pronounced when the postverbal adverb is the distributive disambiguator *individually*. The latter effect can also be observed on the adverb, even though the interaction INDIVIDUALLY:LATE POSITION is less dominantly positive on that word. Both effects are replications of Frazier et al. (1999).

The slowdown caused by postverbal adverbs on the first spillover word is orthogonal to our main research question. It possibly reflects the cost associated with adverb-predicate integration. No such cost would be observable for preverbal adverbs, since these were read before the predicate and could have already been integrated while reading the predicate itself. The effect might alternatively be due to the fact that adverbs are usually preferred and more frequent in preverbal than postverbal positions (see the frequency data in Ko 2016). While the preference is well-known for subject-oriented and speaker-oriented adverbs (e.g., Haider 2004), the results of the present experiment suggest that the preference might also hold for manner adverbs such as *together* and *individually*.

More importantly for us, the postverbal-position slowdown is further modulated by the positive interaction LATE:INDIVIDUALLY. This shows that the postverbal adverb disambiguating towards the distributive reading incurs a processing cost above and beyond any processing cost caused by postverbal adverbs in general. The observed interaction is compatible with the position that readers select the collective interpretation of ambiguous predicates, and that the distributive reinterpretation is costly. The findings generalize the results of Frazier et al. (1999) beyond the disambiguator *each*, and are incompatible with the hypothesis that the processing difficulty associated with postverbal *each* can be fully explained by the semantic complexity of binominal *each*.

There are two alternative explanations that we can discard. First, one could think that the observed slowdown caused by the postverbal adverb *individually* is just due to some general cost associated with this adverb (e.g., because the adverb *individually* is less frequent and morphologically more complex than *together*). If this was so, however, we would also expect *individually* to be read more slowly than *together* in preverbal position, which does not actually happen. In fact, we observed the exact opposite: *individually* is read faster than *together* preverbally when we control for word length. Furthermore, the effect of preverbal adverb type does not spill over to any of the words on the predicate *wrote a sonnet*, showing again that *individually* is not intrinsically harder to process than *together*.

The second alternative explanation we can discard is that distributive readings themselves cause processing difficulties, rather than the reanalysis towards distributivity. It would follow from this second alternative explanation that the early disambiguator towards distributivity should also be a cause of processing difficulties, again contrary to the experimental results.

We thus conclude that the cost of reinterpretation from collectivity to distributivity is genuine. What remains unclear is why the processor should prefer collective readings. The second experiment addresses this question.

3 EXPERIMENT 2

After excluding the option that the results of Frazier et al. (1999) are simply due to the complexity of postverbal *each*, we are left with two possible explanations for why the processor should prefer collective over distributive interpretations. One explanation is that the processor is sensitive to structural economy considerations and prefers structures without the *D* operator. Alternatively, the processor is sensitive to interpretational economy/complexity considerations, and prefers collective readings because they are simpler in the sense that they introduce fewer entities, or other simpler interpretational objects compared to distributive readings. For example, in the sentence *Two boys carried one suitcase*, the collective reading postulates just one event and one suitcase, whilst the distributive reading postulates two events and is compatible with there being two suitcases, such that each event+suitcase is associated with one of the two boys.

Since the collective reading postulates fewer events, and possibly also fewer objects, than the distributive reading, the preference for interpretational economy selects the collective reading.

In semantic research, it has been argued at least since Scha (1981) that the *D* operator is not needed across the board to obtain distributive readings for any predicate. Consider (13).

- (13) a. The girls met.
b. The girls slept.

The predicate *meet* does not express a property of single individuals but a property of groups. Under most semantic theories, it is assumed that the predicate applies to some kind of plurality – a plurality of girls in (13a). The lexical meaning of the predicate will ensure that even though the predicate applies to the group, we learn something about individual girls when hearing (13a) – for example, that every girl, or at least a majority of the girls, was at the same location.

A similar analysis can be given to (13b). Even though we deal with a predicate that expresses a property of individuals (*sleep*), it is possible to assume that the predicate *sleep* applies to the plurality of girls. Due to the lexical meaning of the predicate, we learn something about individual girls from (13b) – that is, that all (or at least most of) the girls are asleep. This is sometimes dubbed a *vagueness* account of distributivity (cf. Winter 2000), or an *underspecification* account (cf. Champollion to appear). We will use the latter label.

It is hard to see how the underspecification account could generalize to examples like (14).

- (14) The girls slept on a narrow bed.

Even if we assume that *sleep* can apply to a group and express that every individual, or most individuals, sleep, as we did above, we would only derive the meaning that there was one narrow bed and every girl slept on it. However, the distributive reading of (14) is more appropriately paraphrased as ‘the girls each slept on a narrow bed’ and requires, in the pragmatically most plausible interpretation, that there were as many beds as there were girls.

More generally, the underspecification account derives distributivity as a by-product of the lexical meaning in (13b). The same strategy, however, would only work for (14) if we assume that the whole predicate *slept on a narrow bed* is stored in the lexicon. If it is not, something else has to be responsible for deriving the distributive reading of (14).

That something else, so the semantic literature argues, is phrasal distributivity, triggered by the distributivity operator, or the *D* operator for short. In most accounts (see Champollion to appear for a recent summary), the *D* operator appears at the VP level and modifies the whole predicate. In (14), this would be the predicate *slept on a narrow bed*. For our purposes, it suffices to say that the operator is interpreted in the same way as *each* (Heim et al. 1991; see Schwarzschild 1996 for a more nuanced view), i.e., the *D*-enriched sentence (14) can be paraphrased as ‘the girls each slept on a narrow bed’. This paraphrase correctly highlights that the distributive reading expresses that each girl slept on a different narrow bed, and there were as many beds as girls.

Before proceeding, let us reiterate one crucial issue about *D* operators in relation to (14). We need the *D* operator not because the verb *sleep* is understood distributively, but because the whole predicate *sleep on a narrow bed* is so interpreted – in particular, narrow beds have to covary with the girls. The distributive interpretation of this whole complex VP poses a challenge to the underspecification account, since the underspecification account should operate only on lexical units. In contrast to that, other VPs that are complex, but that receive a distributive interpretation exclusively driven by their verbs, would not be problematic for the underspecification account. Consider the following example:

- (15) The boys heard a noise.

The most likely interpretation is that there was one particular noise, but each boy heard the noise individually, on his own. The paraphrase signals that hearing is distributive. Yet, since the noise does not covary with the boys, this paraphrase can be captured by assuming that *hear* is specified as signaling the relation between individuals (each individual who hears) and some object (what is heard), and the distributive reading arises through the lexical means. That is, since the distributivity is fully located on the verb only, no *D* operator is needed. See

Let us now see how the phrasal and lexical distributivity might help us explain why we observe the preference for collective readings in processing. If the parser preferred simpler structures, i.e., structures lacking a null operator such as the *D* operator, we should see processing preferences for collective readings of syntactically-built predicates as in (14), but we should not see such preferences for one-word predicates, (13). On the other hand, if the parser generally preferred collective readings, we should see across-the-board processing preferences for collectivity, whether the predicates are syntactically-assembled (14) or not (13b). For example, if the parser preferred interpretations with fewer semantic objects, then the parser should preferentially select the collective reading of both (14) and (13b), since in both cases the collective interpretation postulates only one collective event, unlike the distributive reading.

To test these two hypotheses, we designed a self-paced reading experiment with 8 conditions, exemplified in (16). The predicate could be one word only and the reading could be established based on the lexical knowledge, or the predicate consisted of several (3 or 4) words. In the first case, the predicate was a verb, e.g., *won*. In the second case, the predicate was the same verb followed by an object, e.g., *won an award*. For all the items, the verb was chosen in such a way that it could be followed by a direct object, but an intransitive version in which the object is absent would also be grammatical.

Both collective and distributive readings were possible with the chosen predicates, in both the transitive and the intransitive version. As in Experiment 1, the predicate was disambiguated as collective or distributive by two adverbs, namely *collectively* and *individually*. We use *collectively* instead of *together* since *collectively* matches the other disambiguation adverb, *individually*, more closely in length and morphological structure, while still forcing the non-distributive reading. The disambiguation could happen either before or after the predicate. All conditions are presented in (16).

- (16)
- a. The girls individually won an award during the science fair.
EARLY, INDIVIDUALLY, OBJECT PRESENT
 - b. The girls collectively won an award during the science fair.
EARLY, COLLECTIVELY, OBJECT PRESENT
 - c. The girls won an award individually during the science fair.
LATE, INDIVIDUALLY, OBJECT PRESENT
 - d. The girls won an award collectively during the science fair.
LATE, COLLECTIVELY, OBJECT PRESENT
 - e. The girls individually won during the science fair.
EARLY, INDIVIDUALLY, NO OBJECT
 - f. The girls collectively won during the science fair.
EARLY, COLLECTIVELY, NO OBJECT
 - g. The girls won individually during the science fair.
LATE, INDIVIDUALLY, NO OBJECT
 - h. The girls won collectively during the science fair.
LATE, COLLECTIVELY, NO OBJECT

As the items were relatively complex, we also included an acceptability subtask into the experiment. This was used to check that predicates with and without objects were equally acceptable in the distributive and the collective disambiguation, i.e., both interpretations were equally possible with either predicate.

Let us summarize our predictions for the self-paced reading part of the experiment. If the preference for collectivity is driven by structural considerations, we expect the processor to choose collective readings only for syntactically complex predicates, i.e., when the direct object is present. Consequently, the reinterpretation towards distributivity should incur processing costs only for object-present cases. When the object is absent, the processor has no reason to prefer collective readings over distributive ones. That is, we expect a three-way interaction between adverb position, adverb type and the presence of a direct object. However, if the preference for collective readings is driven only by considerations of interpretational complexity measured in terms of the number of entities postulated in the semantic model, we expect late

distributive disambiguation to cause processing difficulties regardless of the presence of the object. This is because the collective reading of any predicate postulates fewer entities, i.e., fewer events, than the distributive reading of the same predicate. Thus, a two-way interaction between the position of the adverb and the type of adverb is predicted in this case.

3.1 METHOD

3.1.1 Participants

The participants were 55 undergraduate students at University of California Santa Cruz, all self-identified native speakers of English. They received course (extra-)credit for their participation.

3.1.2 Procedure and items

As in Experiment 1, the study was run online. Participants received a link that directed them to the experiment, created on the IBEX platform and hosted on UCSC servers. As in Experiment 1, we used the non-cumulative moving window self-paced reading paradigm.

The participants first received general instructions, which were followed by three practice items and the experiment. The study consisted of 32 target items and 96 fillers. Every item had 8 conditions. An example of an experimental item in all eight conditions is given in (16); see the Appendix for the list of all items. The items were distributed in eight lists using the standard counterbalanced Latin square design. Each participant was assigned to one of the lists and was presented the stimuli (items and fillers) in random order. Every item and every filler was followed by a yes-no comprehension question and, unlike in Experiment 1, by an acceptability judgment of the self-paced reading sentence. Yes-no comprehension questions checked whether participants paid attention to the previous self-paced reading sentence. For every item, the questions were identical across all conditions. The questions were unambiguous and had a clear correct answer. For example, the question following (16) was:

(17) Were the girls at the science fair?

In the acceptability judgment subtask, participants had to rate on a five-point Likert scale (-2, -1, 0, 1, 2) the acceptability of the sentence they had just read. The sentence did not reappear for the acceptability judgment task; that is, participants had to rely on their memory to make the judgment. No answers were considered as correct or incorrect for the acceptability judgment task, and this was also stated in the introductory instructions. However, the comprehension questions had correct answers, and whenever a participant made a mistake in answering a comprehension question, they were directly notified of that.

3.2 RESULTS AND DISCUSSION

The comprehension questions were not difficult for participants: the average accuracy was 90% of answers correct. However, two outlier participants answered only 73% and 66% of questions correctly, so they were excluded from subsequent analyses. All the other participants answered at least 79% of the questions correctly.

The descriptive summary of the acceptability data is provided in [Table 1](#). The acceptability data was analyzed using mixed-effects ordinal probit models, with the following treatment-coded

ADVERB	DISAMBIGUATION	OBJECT	RESPONSE (MEDIAN)	RESPONSE (MEAN)	RESPONSE (SD)
collectively	early	null	1	1.036	0.950
collectively	early	object	1	1.104	0.886
collectively	late	null	1	1.096	0.967
collectively	late	object	1	0.915	0.953
individually	early	null	1	0.967	0.973
individually	early	object	1	1.029	0.972
individually	late	null	1	1.175	0.921
individually	late	object	1	0.802	1.066

Table 1 Descriptive summary of the acceptability subtask. Participants responded by selecting a value on the five-point Likert scale (-2, -1, 0, 1, 2).

fixed effects: OBJECT (ABSENT or PRESENT, the former was the reference level), ADVERB (COLLECTIVELY or INDIVIDUALLY, the former was the reference level), POSITION (EARLY or LATE, the former was the reference level); the interaction of ADVERB and POSITION; the interaction of OBJECT and POSITION; the interaction of ADVERB and OBJECT; and the three-way interaction ADVERB, POSITION, OBJECT. The model included the full random-effect structure for subjects and items. The Bayesian model was estimated using Stan and *brms* (cf. Bürkner and Vuorre 2019). The structure of prior distributions and the sampling details of the model are provided in the Appendix.

Posterior distributions are graphically summarized in *Figure 3*. Negative values indicate decreased acceptability, while positive values indicate increased acceptability. Only one parameter has a posterior distribution with a 89% credible interval that excludes 0: the negative interaction of LATE POSITION:OBJECT (89% credible interval: $\{-0.69, -0.15\}$, median = -0.41 , $p(\beta > 0) = 0.992$). That is, sentences with adverbs following multi-word predicates were judged as worse, irrespective of adverb type. We take this to once again indicate that English speakers disprefer adverbs in postverbal position. This finding from the acceptability subtask matches one of the findings in Experiment 1, where increased RTs were observed for postverbal adverbs (recall that in Experiment 1, all predicates consisted of a verb + an object). 89% credible intervals of all the other parameters include zero. In particular, no factor with INDIVIDUALLY shows a fully negative or a fully positive credible interval, suggesting that readers did not find distributive readings of the predicates in the experiment as less or more acceptable than collective readings. This suggests, in turn, that items in this experiment were just as compatible with distributive construals as they were with collective construals. Thus, if we are to observe any difficulties with distributive readings when we turn to incremental processing in the self-paced reading part of the study, we can be reasonably confident that these are not due to the decreased acceptability of the distributive or collective interpretation of the predicates.¹¹

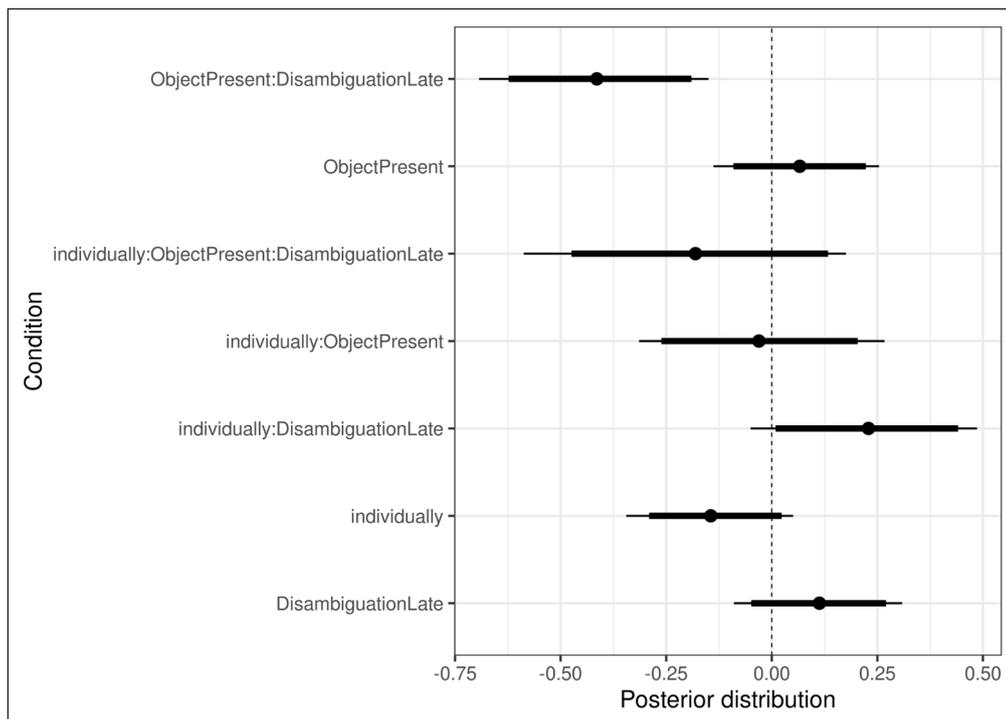


Figure 3 Posterior distributions of parameters of interest in the acceptability study. A thick horizontal line represents the 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents the 89% HDPI of a parameter. The dashed vertical line marks the position of 0. Filled circles represent median values. Negative values indicate that the parameter decreases acceptability, while positive values indicate that the parameter increases acceptability.

For the analysis of RTs, we proceeded in the same way as in Experiment 1. We removed RTs that were faster than 50 ms and slower than 3000 ms, and we log-transformed RTs. As in

¹¹ This point should be kept in mind when considering all the items in the experiment. While both the distributive and the collective reading are easily available for *win/win an award*, our example in (16), other predicates used in the experiment might intuitively be harder to interpret under the collective reading, e.g., *ran/ran a mile* (see the Appendix for the full list of items). However, the collective reading seems possible in this last case as well, as witnessed by the fact that the predicate can combine with *collectively*, irrespectively of the presence of the object. The collective reading signals that the activity is shared in close time-space proximity, and the people participating in it likely share the same goal, in contrast to the distributive version of *ran/ran a mile* (see also the discussion in Lasersohn 1995 of the role of proximity in the interpretation of the anti-distributive adverb *together*). In any case, what is crucial for us is that the version without an object does not differ from the version with an object with respect to the plausibility of the reading forced by *collectively*. The fact that predicates do not differ from each other in this respect is confirmed by the acceptability study.

Experiment 1, we probed for (very fast or very slow) outliers among readers by checking mean logRTs per subject. Two subjects had a mean logRT more than 2.5 standard deviations away from the grand mean and were removed from the final analysis of RTs. The final number of participants whose data we analyzed: 51. The influence of word length and word position was also factored out in the same way as in Experiment 1.

The main regions of interest (ROIs) were the adverb, the predicate and the three words following the disambiguated predicate, the spillover region. The ROIs are underlined in (18).

- (18) a. The girls individually won an award during the science fair.
 b. The girls collectively won an award during the science fair.
 c. The girls won an award individually during the science fair.
 d. The girls won an award collectively during the science fair.
 e. The girls individually won during the science fair.
 f. The girls collectively won during the science fair.
 g. The girls won individually during the science fair.
 h. The girls won collectively during the science fair.

The graphical summary of logRTs is presented in [Figure 4](#). LogRTs are also summarized in a table in the Appendix.

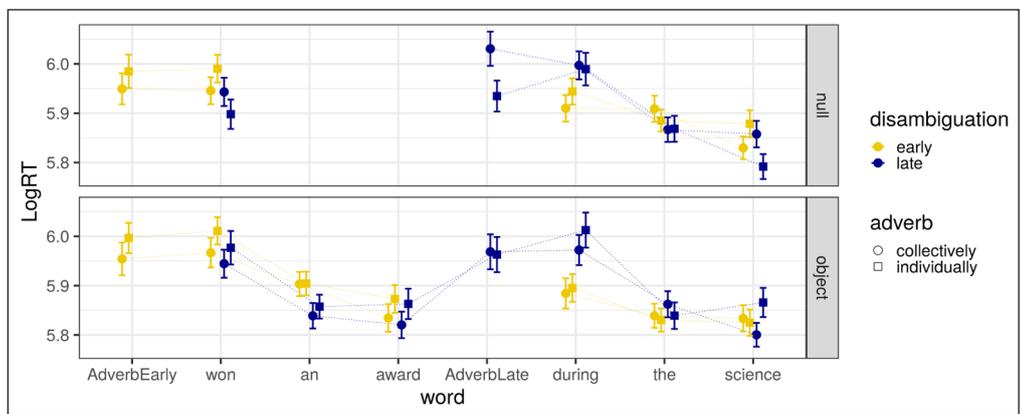


Figure 4 Exp. 2, logRTs per word in the predicate, adverb and spillover regions. The position AdverbEarly shows logRTs for the adverbs *individually* and *collectively* in the early position. The position AdverbLate shows logRTs for the adverbs *individually* and *collectively* in the late position. The top graph shows RTs for conditions without the object, the bottom graph shows RTs for conditions in which the object was present.

For each region, we estimated a Bayesian mixed-effects model with the same fixed effects as in the acceptability study (ADVERB, POSITION, OBJECT and their interactions). The models had a maximal random-effect structure and were estimated using Stan and *brms*. The structure of prior distributions and the sampling details of the models are provided in the Appendix.

Posterior distributions of the parameters on the adverb and spillover regions are summarized graphically in [Figure 5](#), [Figure 6](#), [Figure 7](#), and [Figure 8](#).

[Figure 5](#), corresponding to the adverbs *individually/collectively*, shows that the LATE POSITION stands out as a positive effect (89% credible interval: (0.02, 0.14), median = 0.08, $p(\beta > 0) = 0.987$). The slowdown due to LATE POSITION is most likely related to the slowdown caused by late adverbs in Experiment 1 and, possibly, to the interaction of LATE ADVERB:OBJECT in the acceptability study. These results repeatedly show that readers incur a processing cost when integrating adverbs that appear in a post-predicate position. The effect, while interesting on its own, is not of primary concern to us. Another effect stands out in the same region: the interaction LATE POSITION:INDIVIDUALLY (89% credible interval: (-0.21, -0.04), median = -0.13, $p(\beta < 0) = 0.991$). This signals that late disambiguation towards distributivity with lexical predicates (i.e., predicates that consist just of one word, the intransitive verb) speeds up reading. The speed-up is diminished in the case of phrasal predicates, since most of the posterior probability mass of the three-way interaction OBJECT PRESENT:LATE POSITION:INDIVIDUALLY is positive (median = 0.067, $p(\beta > 0) = 0.82$).

[Figure 6](#), corresponding to the first word in the spillover (the word *during* in (18)), shows that the posterior distribution of LATE POSITION is once again positive (89% credible interval: (0.028, 0.144), median = 0.09, $p(\beta > 0) = 0.991$). This slowdown has likely the same explanation as the same slowdown observed on the adverb.

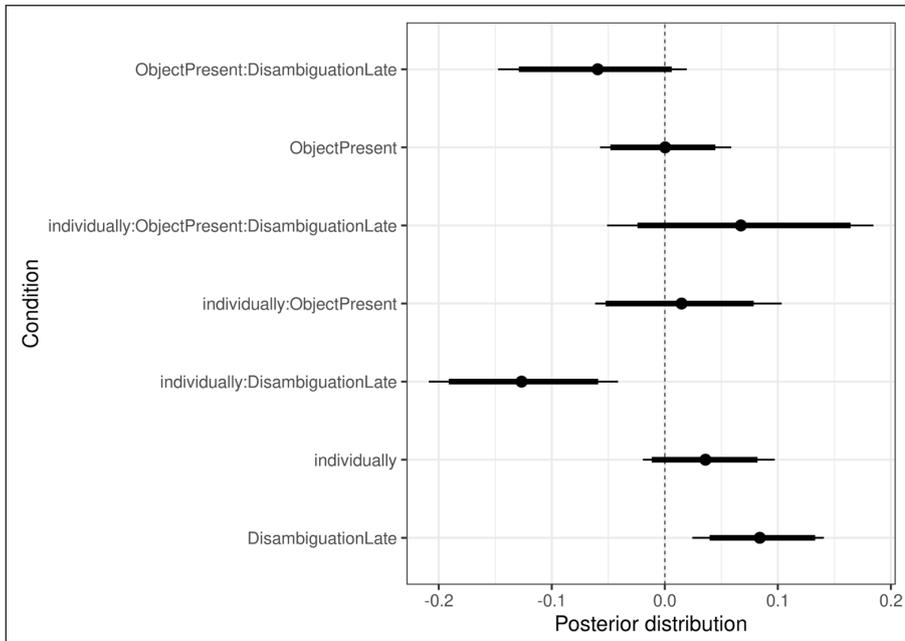


Figure 5 Posterior distributions of parameters of interest on the adverb *individually/collectively*. A thick horizontal line represents the 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents the 89% HDPI of a parameter. The dashed vertical line marks the position of 0. Filled circles represent median values.

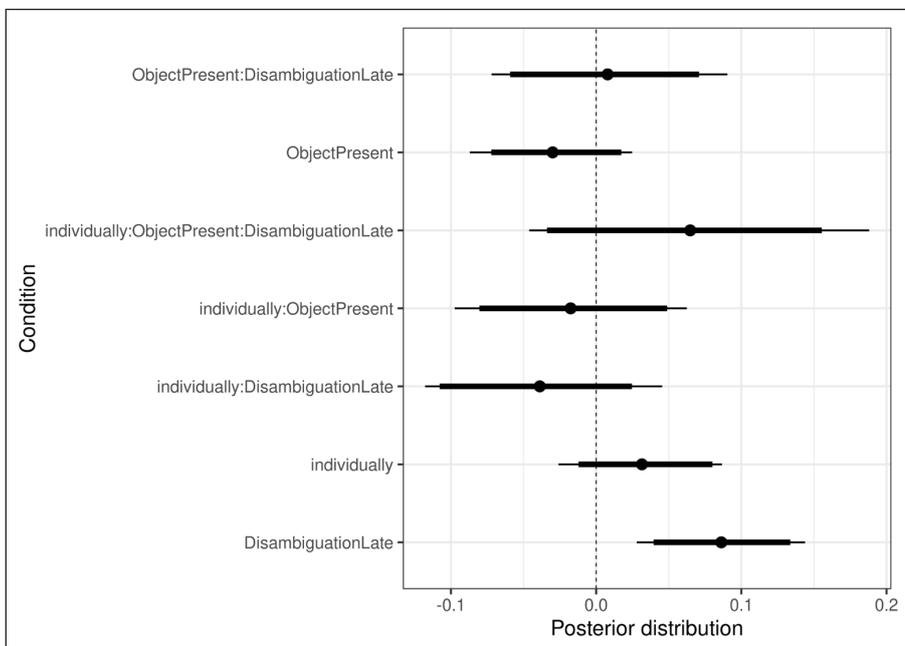


Figure 6 Posterior distributions of parameters of interest on the 1st word in the spillover. A thick horizontal line represents the 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents the 89% HDPI of a parameter. The dashed vertical line marks the position of 0. Filled circles represent median values.

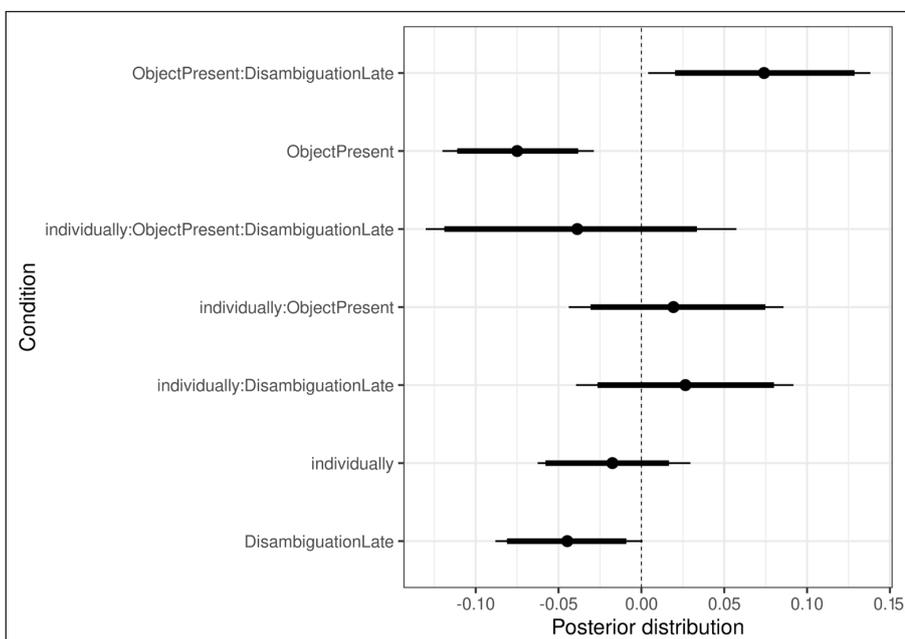


Figure 7 Posterior distributions of parameters of interest on the 2nd word in the spillover. A thick horizontal line represents the 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents the 89% HDPI of a parameter. The dashed vertical line marks the position of 0. Filled circles represent median values.

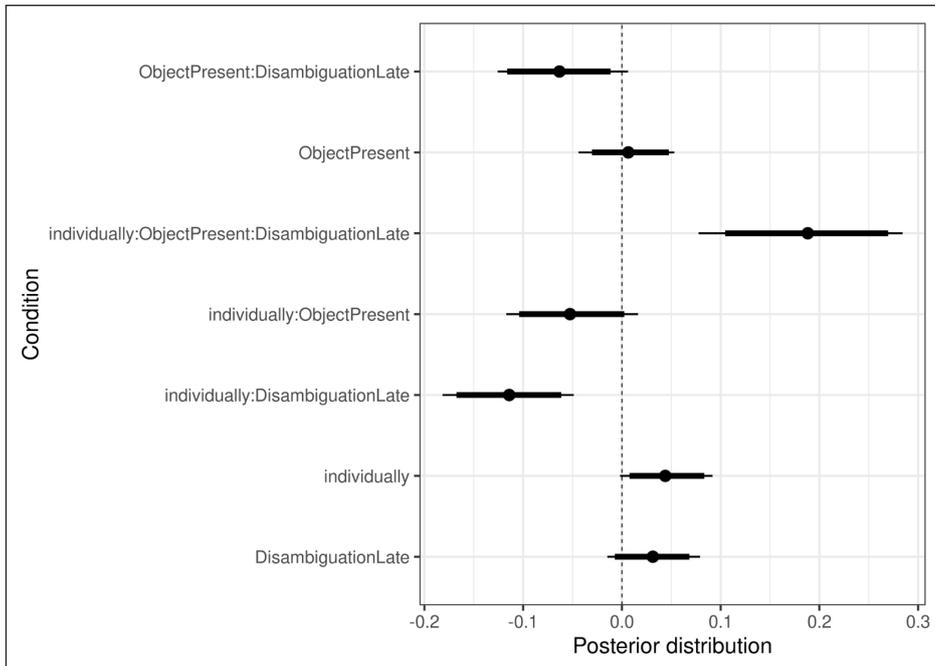


Figure 8 Posterior distributions of parameters of interest on the 3rd word in the spillover. A thick horizontal line represents the 80% highest posterior density interval (HDPI) of a parameter, a thin horizontal line represents the 89% HDPI of a parameter. The dashed vertical line marks the position of 0. Filled circles represent median values.

Figure 7, corresponding to the second word of the spillover (the word *the* in (18)), shows that OBJECT PRESENT clearly affects reading times by causing a speed-up (89% credible interval: $\langle -0.12, -0.03 \rangle$, median = -0.075 , $p(\beta < 0) = 0.995$). This effect is probably caused by the fact that objects add extra information and increase the predictability of the follow-up regions. This facilitating role of direct object is orthogonal to our main investigation. We also see that OBJECT PRESENT:LATE POSITION is associated with a slowdown (89% credible interval: $\langle 0.004, 0.14 \rangle$, median = 0.074 , $p(\beta < 0) = 0.961$), which matches the decreased acceptability of the same condition.

Finally, **Figure 8**, corresponding to the third word of the spillover (the word *science* in (18)), shows two conditions that clearly stand out. First, there is a negative interaction of LATE POSITION and ADVERB INDIVIDUALLY (89% credible interval: $\langle -0.18, -0.05 \rangle$, median = -0.114 , $p(\beta < 0) = 0.997$). Second, there is a dominant positive effect associated with the three-way interaction OBJECT PRESENT, LATE POSITION and ADVERB INDIVIDUALLY (89% credible interval: $\langle 0.07, 0.28 \rangle$, median = 0.19 , $p(\beta > 0) = 0.998$). Thus, while the late disambiguation towards distributivity with lexical predicates (i.e., predicates that consist just of one word, the intransitive verb) speeds up reading, the effect is clearly reversed when the predicate consists of the verb + an object. In that case, disambiguation towards distributivity slows down reading. The effect can also be clearly observed in **Figure 9** on the word *science* (the last word in the figure).

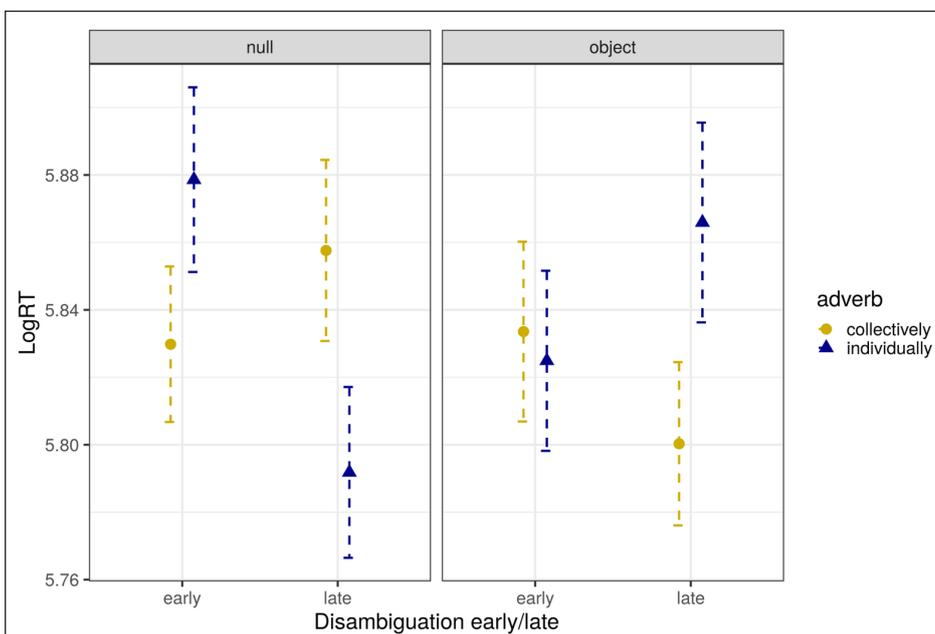


Figure 9 Exp. 2, LogRTs on the third word of the spillover, i.e., the third word after the predicate + the adverb (the word *science* in (18)). The graph shows a three-way positive interaction between ADVERB, POSITION and OBJECT: late distributive disambiguation speeds up reading when the object is null (left panel), but slows down reading when the object is present (right panel).

Based on the two experiments, we can conclude that readers have processing difficulties when they read a predicate that is in principle compatible with both collective and distributive interpretations, and they learn afterwards that the predicate should be interpreted distributively.

The findings extend the results of Frazier et al. (1999) in two significant ways. First, we found that the slowdown is not caused just by the postverbal disambiguator *each*, but appears with another distributive disambiguator, the adverb *individually*. This is important to establish since *each* in postverbal position affects interpretations in other ways than just enforcing distributivity. Observing the effect with the disambiguator *individually* gives us stronger evidence that it is really the reinterpretation towards distributivity that incurs processing cost.

Second, we saw that the distributive reinterpretation of predicates consisting of a verb + an object is costly. But the distributive reinterpretation of one-word predicates (intransitive verbs) is not harder than the collective reinterpretation. This second finding fits well with formal semantics research arguing that, in the case of one-word predicates, the distributive/collective reading can just follow from our lexical knowledge, but in the case of complex predicates, that is, multi-word expressions whose interpretation is syntactically mediated, something more has to be said – specifically, the presence of a *D*-operator is needed.¹² Thus, the distributive reinterpretation is costly only when distributive interpretations are triggered by a covert quantificational element that can take scope, the *D*-operator. Our findings match well with research on structural priming of distributive and non-distributive readings of pluralities, which provides evidence that the two readings involve different mental and structural representations (Maldonado et al. 2017). We will now discuss in detail what our results mean for the theory of distributivity and collectivity, and the processing of these readings.

4.1 AMBIGUITY

Frazier et al. (1999) use their results to argue that the distinction between collectivity and distributivity is a case of ambiguity, not underspecification. This theoretical conclusion is based on the results of Frazier and Rayner (1990), which show that the human processor encounters difficulties (observable as increased reading times) when an ambiguous word (a homonym) is introduced first and disambiguated later towards its dispreferred reading. Crucially, no such difficulty is found when a word with multiple senses is introduced first and one of its senses is instantiated later, regardless of the sense that is instantiated. The question whether the distributive/collective distinction is a matter of ambiguity or underspecification has not been fully resolved in theoretical linguistics (see Lasersohn 1995 and Schwarzschild 1996 for opposing views, and Nouwen 2012; Champollion to appear for a summary of the issues) and thus, processing data might be very important in this debate.

Our experiments add an important qualification to Frazier et al. (1999). Given the results of Experiment 2, we have evidence that phrasal predicates, i.e., predicates constructed in syntax, bias readers towards collective readings, while lexical predicates (consisting of only one word) exhibit a preference for distributive readings. One possibility is that we deal with ambiguity for both predicate types, but each type establishes different preferences (see below as to why the human processor might prefer one reading over the other).

Alternatively, it could be the case that only phrasal predicates are truly ambiguous. Lexical predicates (like the predicate *won*) are underspecified, and the observed preference for distributivity has a different source. For example, world knowledge makes the predicates used in the experiment more easily compatible with their distributive interpretation than the collective one (see Scontras and Goodman 2017 for a formally explicit account that could model this world-knowledge driven preference). This explanation would go against Frazier and Rayner (1990) and suggest that there are, after all, cases of underspecification in which reinterpretation towards one of the senses can affect reading times. The explanation is compatible with the

¹² Other operators beside the *D*-operator have been postulated to explain the existence of distributive readings (see the discussion in Landman 2000). Another possibility is that the ambiguity might be located elsewhere, for example, in the subject noun phrase (but see Lasersohn 1995 for discussion). The *D*-operator based explanation, however, is the dominant view, which is why we focus on it.

work in theoretical linguistic research arguing that only phrasal distributivity/collectivity is a case of (structural) ambiguity.

More generally, our data provide processing evidence that phrasal and lexical distributivity differ from each other. Theoretical work that assumes two different mechanisms for the two types of distributivity are compatible with our findings. This is true for the theories that postulate a *D* operator for phrasal distributivity, while they let the distributivity of lexical predicates arise through lexical means, e.g., the vagueness/underspecification of the lexical element (Link 1987; Roberts 1990; Lasersohn 1989; 1995; Moltmann 1997; Winter 2000; 2001; Kratzer 2008; Champollion 2010; Dotlačil 2010; Champollion 2016a; 2017; de Vries 2017; see Glass 2018 and Champollion 2019 for discussion). The theories that postulate only one mechanism for both types of distributivity (Landman 1989; Schwarzschild 1996; Beck and Sauerland 2000; Glass 2018) would need a different explanation for why phrasal and lexical cases exhibit a different processing profile.

4.2 WHAT TRIGGERS THE PREFERENCE FOR COLLECTIVE READINGS

We mentioned two possible explanations for why readers prefer the collective interpretation of a predicate that could be interpreted as distributive or collective: (i) the processor might prefer the simplest structure, and selects the collective reading because the structure that generates this reading lacks the *D* operator; (ii) alternatively, the processor selects the collective reading because its interpretation is less complex in that it postulates fewer events and objects.

Our results are compatible with the first explanation. We saw that readers choose the collective reading of a predicate when the distributive reading would require the *D* operator, preferring a simpler syntactic structure without the operator.

But the second explanation is not immediately compatible with our results. In particular, we would have to explain why readers prefer the collective reading of a multi-word predicate – a verb + an object such as *won an award* – but they do not prefer the collective reading of a one-word predicate – the verb *won* by itself. Specifically, in the case of multi-word predicates, the distributive reading leads one to assume that there are multiple objects (multiple awards that were won) and multiple winning events, one per person + award. The latter assumption follows from a standard requirement in event semantics that every event should have uniquely specified thematic roles (Parsons 1990). Under the interpretational-simplicity hypothesis, this difference between collective and distributive readings leads readers to prefer the collective interpretation. But the very same hypothesis makes incorrect predictions for one-word predicates like the verb *won*. The distributive reading of this predicate also establishes multiple winning events (one per agent), while the collective reading requires only one event (since there is only one agent, the group that won). Given that, we would expect that lexical predicates like *won* should also exhibit a preference for collective interpretations, which has not been observed in our data.

The proponents of this second explanation/hypothesis could argue that the contrast between multi-word predicates and one-word predicates is due to the different ‘ambiguity status’ of the two predicate types. If the processor has to choose one of the readings only for ambiguous predicates (as argued in Frazier and Rayner 1990), and if only multi-word predicates are ambiguous, we could maintain that the processor is driven by simplicity in interpretation. The results in the second experiment would follow because only in the case of multi-word predicates does the processor have to explicitly represent the complexities of distributive/collective readings. That is, the processor has to make a choice only in the case of phrasal predicates (which are truly ambiguous).

We see this as a potential explanation for part of our data. It is still unclear under this account why one-word predicates like the intransitive verb *won* should show any preference for either reading, in particular, a preference for *distributive* interpretations. After all, if such predicates are vague with respect to collectivity/distributivity, readers should not be forced to choose either interpretation, according to Frazier and Rayner (1990). Of course, one could give up this assumption and conjecture that readers make a choice even in the case of underspecified predicates like *won*. But then it is unclear why readers do not choose the collective reading here,

just as they do in case of multi-word predicates – particularly since the collective interpretation of *won* requires fewer events than its distributive interpretation.

If we assume that the processing preference is driven by the preference to have structures without the *D* operator, then we straightforwardly explain why the processor prefers the collective reading for multi-word predicates like *won an award*. This position is, furthermore, compatible with the finding that one-word predicates are preferably interpreted as distributive. The reason is that collective and distributive readings of one-word predicates cannot be distinguished structurally, and therefore, other factors like pragmatic considerations, world knowledge etc., can and will greatly affect a reader's choice.¹³

Stepping back, we see that only predicates constructed in syntax incur a processing cost when they have to be reinterpreted as distributive. There are two possible ways to explain this fact. The first option is to interpret the data as showing that the choice of the preferred interpretation is driven by simplicity in syntax. The second option is to interpret the data as showing that the processor only has to make a choice when it encounters a predicate that was constructed in syntax. Under either interpretation, it becomes clear that the parser has to be sensitive to the distinction between predicates constructed in syntax, i.e., predicates that have to carry the quantificational *D* operator for their distributive reading, and predicates that can carry their distributive/collective reading in their lexical meaning.

4.3 MULTI-WORD EXPRESSIONS AND LEXICAL DISTRIBUTIVITY

We took the results of Experiment 2 to indicate that syntactically-complex predicates exhibit different biases during incremental interpretation than predicates that are not syntactically complex. Strictly speaking, though, we just see that multi-word predicates (*won an award*) show different preferences than one-word predicates (*won*). This brings up one potential confound in our study: could it be that we observe the three-way interaction between *ADVERB*, *POSITION* and the presence of *OBJECT* just because the multi-word predicates are longer (they consist of at least three words) than one-word predicates? It is known that the cost of reanalysis towards the marked interpretation is more pronounced when readers are allowed a longer time to build and commit to the preferred, default interpretation (Frazier and Rayner 1982). Under this explanation, the observed three-way interaction could simply be a consequence of the longer span during which the default collective reading has been built up (i.e., the reading has become entrenched).

But this hypothesis cannot account for the fact that, in the case of one-word predicates, we see the exact opposite effect: readers prefer late disambiguation towards distributive, rather than collective, readings. This is unexpected if certain predicates are preferably understood as collective, and only the length of these predicates modulates the cost of reanalysis towards the marked interpretation.

A related question is whether there would be any way to distinguish between predicates that consist of several words and predicates that are constructed in syntax. Are there multi-word predicates that should be seen as single units from the perspective of semantics, and that do not need quantificational elements (like the *D*-operator) to resolve distributive/collective ambiguities? Addressing this question opens up an interesting research area where the study of plurality, processing and the lexicon/syntax interface interact. Clearly, though, this research program goes beyond the topic of this article. We only note that, in the semantic literature, it has been proposed that some multi-word predicates might be treated as atomic units

¹³ Alternatively, one-word predicates might in general prefer distributive readings. That is, once we remove structural considerations like the requirement for the *D* operator, the distributive reading is preferred by the processor.

Notice that the acceptability study argues against the possibility that we inadvertently introduced a contrast between one-word and multi-word predicates in Experiment 2, such that predicates in the former group have a plausibility bias towards distributivity, and predicates in the latter group have a bias towards collectivity. If there was a collectivity bias for multi-word predicates, we would expect that *individually* should cause degraded acceptability when the object is present, as *individually* would clash with this hypothetical bias of the multi-word predicate. However, no such decrease in judgments is observed in the acceptability study. The adverb *individually* seems to be just as compatible with multi-word predicates in Experiment 2 as *collectively* is. In other words, both interpretations are equally acceptable, suggesting no such plausibility bias is present. We would like to thank an anonymous reviewer for discussion of this point.

and, consequently, that such predicates might lack a preference for collectivity. Idiomatic expressions like *take a nap* and predicates with bleached verbs like *have a coffee* belong here (see Vries 2015).¹⁴

4.4 PHONOLOGICALLY NULL SEMANTIC OPERATORS

One issue that has sparked interest in semantic and psycholinguistic literature is whether real-time studies of incremental interpretation can provide processing evidence for null semantic operators, and for logical-form displacements that need to be posited for interpretational reasons (Pylkkänen and McElree 2006). This research line has mostly focused on the processing of quantifier scope (see Brasoveanu and Dotlačil to appear for a recent summary) and argument and aspect coercion (see Piñango et al. 2006; Pickering and Ferreira 2008).

In this paper, we turned our attention to the incremental processing of pluralities. The results of the two reported self-paced reading studies show that there is a difference between the processing of phrasal vs. lexical distributivity, and that this contrast aligns well with the theoretical research that postulates a null semantic operator in the former, but not in the latter, case. Insofar as this interpretation is correct, it can be seen as providing further support for the research line pursued in Pylkkänen and McElree (2006). Null semantic operators impact processing, and real-time/online methodologies like self-paced reading, which provide a window into the nature of incremental processing, can be used as an independent source of evidence for the presence of null semantic operators in an LF structure.

5 CONCLUSION

We studied the interpretation of plural expressions like *the students*, focusing on the question of how the processor establishes the distributive/collective interpretation of predicates accompanying such plural expressions.

We reported the results of two self-paced reading experiments. In the first experiment, we built on Frazier et al. (1999) and replicated their finding that the human processor favors collective readings over distributive ones. Based on the results of the second experiment, we argued that the preference for collective readings is only observed when the distributive reading arises via phrasal distributivity, and the preference disappears in the case of lexical distributivity.

The findings provide evidence that the distinction between collective and distributive readings is a case of ambiguity, at least for phrasal predicates. They also indicate that the processor might give more weight to some sources of evidence during incremental interpretation. Specifically, syntactic information and structural simplicity might contribute stronger processing constraints than mental model simplicity, at least when this simplicity/complexity is measured in terms of the number of events or entities that the processor needs to postulate during incremental interpretation.

ADDITIONAL FILE

The additional file for this article can be found as follows:

- **Appendix.** The Appendix lists the items of Experiment 1 and Experiment 2. It provides descriptive summaries for both experiments and includes details about Bayesian models. DOI: <https://doi.org/10.5334/gjgl.1131.s1>

ACKNOWLEDGEMENTS

We would like to thank the *Glossa* anonymous reviewers for their detailed comments on previous versions of this paper. We would also like to thank Nathan Arnett, Valentina Bianchi, Jenny Doetjes, Donka Farkas, Lyn Frazier, Joseph King, Jess Law, Megan Moodie, Rick Nouwen, Elena Pagliarini, Milica Petranovic, Jennifer Spender, Hanna de Vries, and the audiences and

¹⁴ de Vries (2015) argues that, in fact, any verb + simple indefinite object with the article *a* can become one unit, fully lexicalized, and does not need the *D* operator. This lexicalization, however, is optional, which makes her position compatible with our experimental data – if we assume that this lexicalization is not a processing default.

participants of the CUNY 2015 Conference on Human Sentence Processing, the Groningen Linguistics Department Colloquium Series and the joint *Language, Logic and Cognition* and *S-lab* lab meetings at UC Santa Cruz for their comments and feedback at various stages of this research project. The usual disclaimers apply.

FUNDING INFORMATION

Jakub Dotlačil was supported by a Rubicon grant and VENI grant from the Netherlands Organization for Scientific Research. Adrian Brasoveanu was supported by an SRG grant from the UC Santa Cruz Committee on Research.

COMPETING INTERESTS

The authors have no competing interests to declare.

AUTHOR AFFILIATIONS

Jakub Dotlačil  orcid.org/0000-0002-5337-8432
Utrecht University, Trans 10, 3512JK Utrecht, Netherlands

Adrian Brasoveanu  orcid.org/0000-0001-8488-1113
University of California Santa Cruz, Stevenson Faculty Services, 1156 High Street, Santa Cruz, CA 95064, US

REFERENCES

- Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68. 255–278. DOI: <https://doi.org/10.1016/j.jml.2012.11.001>
- Beck, Sigrid, and Uli Sauerland. 2000. Cumulation is needed: A reply to Winter (2000). *Natural Language Semantics* 8. 349–371. DOI: <https://doi.org/10.1023/A:1011240827230>
- Blaheta, Don. 2003. Binominal *each*: evidence for a modified type system. Master's thesis, Brown University, Providence, Rhode Island.
- Boylan, Christine, Dimka Atanassov, Florian Schwarz, and John Trueswell. 2011. A matter of ambiguity: Using eye movements to examine collective vs. distributive interpretations of plural sets. *Presentation at XPRAG 2011*, Barcelona, Spain.
- Brasoveanu, Adrian, and Jakub Dotlačil. to appear. Processing quantification. In Chris Cummins and Napoleon Katzos (eds.), *Handbook of experimental semantics and pragmatics*, 1–19. Oxford, UK: Oxford University Press.
- Brooks, Patricia J., and Martin D.S. Braine. 1996. What do children know about the universal quantifiers *all* and *each*? *Cognition* 60. 235–268. DOI: [https://doi.org/10.1016/0010-0277\(96\)00712-3](https://doi.org/10.1016/0010-0277(96)00712-3)
- Bürkner, Paul-Christian. 2017. brms: An R package for bayesian multilevel models using stan. *Journal of Statistical Software* 80. 1–28. URL <https://www.jstatsoft.org/v080/i01>. DOI: <https://doi.org/10.18637/jss.v080.i01>
- Bürkner, Paul-Christian, and Matti Vuorre. 2019. Ordinal regression models in psychology: A tutorial. *Advances in Methods and Practices in Psychological Science* 2. 77–101. DOI: <https://doi.org/10.1177/2515245918823199>
- Burzio, Luigi. 1986. *Italian syntax: A government-binding approach*. Dordrecht: Kluwer. DOI: <https://doi.org/10.1007/978-94-009-4522-7>
- Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *Journal of statistical software* 76. 1–32. DOI: <https://doi.org/10.18637/jss.v076.i01>
- Champollion, Lucas. 2010. Cumulative readings of every do not provide evidence for events and thematic roles. In Maria Aloni and Katrin Schulz (eds.), *Proceedings of the 17th Amsterdam Colloquium 2009*, 213–222. Springer. DOI: https://doi.org/10.1007/978-3-642-14287-1_22
- Champollion, Lucas. 2016a. Covert distributivity in algebraic event semantics. *Semantics and Pragmatics* 9. 1–65. DOI: <https://doi.org/10.3765/sp.9.15>
- Champollion, Lucas. 2016b. Overt distributivity in algebraic event semantics. *Semantics and Pragmatics* 9. 1–65. DOI: <https://doi.org/10.3765/sp.9.16>
- Champollion, Lucas. 2017. *Parts of a whole: Distributivity as a bridge between aspect and measurement*, volume 66. Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780198755128.003.0008>
- Champollion, Lucas. 2019. Distributivity in formal semantics. *Annual Review of Linguistics* 5. 289–308. DOI: <https://doi.org/10.1146/annurev-linguistics-011718-012528>

- Champollion, Lucas. to appear. Distributivity, collectivity and cumulativity. In Lisa Matthewson, Cecile Meier, Hotze Rullmann, and Thomas Ede Zimmermann (eds.), *Wiley's companion to semantics*. Wiley.
- Crain, Stephen, and Mark Steedman. 1985. On not being led up the garden path: The use of context by the psychological syntax processor. In David Dowty, Lauri Karttunen, and Arnold Zwicky (eds.), *Natural language parsing: Psychological, computational and theoretical perspectives*, 320–358. Cambridge, UK: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511597855.011>
- Dobrovie-Sorin, Carmen, Emilia Ellsiepen, and Barbara Hemforth. 2016. Why are distributive readings dispreferred. In E. Carrilho, A. Fiéis, M. Lobo, and S. Pereira (eds.), *Selected papers from Going Romance 28*, Lisbon, 83–102. Amsterdam: John Benjamins. DOI: <https://doi.org/10.1075/rllt.10.05dob>
- Dotlačil, Jakub. 2010. Anaphora and distributivity. A study of *same*, *different*, reciprocals and *others*. Doctoral Dissertation, Utrecht University.
- Dotlačil, Jakub. 2012. Binominal *each* as an anaphoric determiner: Compositional analysis. In Ana Aguilar Guevara, Anna Chernilovskaya, and Rick Nouwen (eds.), *Proceedings of Sinn und Bedeutung 16*, 211–225. MIT Working Papers in Linguistics.
- Dotlačil, Jakub. to appear. Semantic approaches to number. In Jenny Doetjes and Patricia Cabredo Hofherr (eds.), *Oxford handbook on number*. Oxford: Oxford University Press.
- Frazier, Lyn. 1978. On comprehending sentences: Syntactic parsing strategies. Doctoral Dissertation, University of Connecticut.
- Frazier, Lyn, Jeremy M. Pacht, and Keith Rayner. 1999. Taking on semantic commitments, II: collective versus distributive readings. *Cognition* 70. 87–104. DOI: [https://doi.org/10.1016/S0010-0277\(99\)00002-5](https://doi.org/10.1016/S0010-0277(99)00002-5)
- Frazier, Lyn, and Keith Rayner. 1982. Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive psychology* 14. 178–210. DOI: [https://doi.org/10.1016/0010-0285\(82\)90008-1](https://doi.org/10.1016/0010-0285(82)90008-1)
- Frazier, Lyn, and Keith Rayner. 1990. Taking on semantic commitments: processing multiple meanings versus multiple senses. *Journal of Memory and Language* 29. 181–200. DOI: [https://doi.org/10.1016/0749-596X\(90\)90071-7](https://doi.org/10.1016/0749-596X(90)90071-7)
- Futrell, Richard, Edward Gibson, Harry J. Tily, Idan Blank, Anastasia Vishnevetsky, Steven T. Piantadosi, and Evelina Fedorenko. 2018. The natural stories corpus. In *Proceedings of LREC 2018, Eleventh International Conference on Language Resources and Evaluation*, 76–82. Miyazaki, Japan.
- Glass, Lelia. 2018. Distributivity, lexical semantics, and world knowledge. Doctoral Dissertation, Stanford University.
- Haider, Hubert. 2004. Pre- and postverbal adverbials in *ov* and *vo*. *Lingua* 114. 779–807. DOI: [https://doi.org/10.1016/S0024-3841\(03\)00051-2](https://doi.org/10.1016/S0024-3841(03)00051-2)
- Heim, Irene, Howard Lasnik, and Robert May. 1991. Reciprocity and plurality. *Linguistic Inquiry* 22. 63–101.
- Just, Marcel A., Patricia A. Carpenter, and Jacqueline D. Woolley. 1982. Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General* 111. 228–238. DOI: <https://doi.org/10.1037/0096-3445.111.2.228>
- Ko, Edwin. 2016. A corpus-based study of variation and change in adverb placement across world englishes. Doctoral Dissertation, Georgetown University.
- Kratzer, Angelika. 2008. On the plurality of verbs. In Johannes Dölling, Tatjana Heyde-Zybatow, and Martin Schäfer (eds.), *Event structures in linguistic form and interpretation*, 269–300. Berlin/New York: Mouton de Gruyter.
- Krifka, Manfred. 1989. Nominal reference, temporal constitution, and quantification in event semantics. In Renate Bartsch, Johan van Bentham, and Peter van Emde Boas (eds.), *Semantics and contextual expressions*, 75–116. Dordrecht: Foris. DOI: <https://doi.org/10.1515/9783110877335-005>
- Kruschke, John. 2014. *Doing bayesian data analysis: A tutorial with R, JAGS, and Stan*. Academic Press. DOI: <https://doi.org/10.1016/B978-0-12-405888-0.00008-8>
- Landman, Fred. 1989. Groups I. *Linguistics and Philosophy* 12. 559–605. DOI: <https://doi.org/10.1007/BF00627774>
- Landman, Fred. 1995. Plurality. In Shalom Lappin (ed.), *The Handbook of Contemporary Semantic Theory*. London: Blackwell.
- Landman, Fred. 2000. *Events and Plurality: The Jerusalem Lectures*. Dordrecht: Kluwer. DOI: <https://doi.org/10.1007/978-94-011-4359-2>
- Laserson, Peter. 1989. On the readings of plural noun phrases. *Linguistic inquiry* 20. 130–134.
- Laserson, Peter. 1995. *Plurality, Conjunction and Events*. Dordrecht: Kluwer. DOI: <https://doi.org/10.1007/978-94-015-8581-1>
- Link, Godehard. 1987. Generalized quantifiers and plurals. In Peter Gärdenfors (eds.), *Generalized quantifiers*, 151–180. Springer. DOI: https://doi.org/10.1007/978-94-009-3381-1_6
- Maldonado, Mora, Emmanuel Chemla, and Benjamin Spector. 2017. Priming plural ambiguities. *Journal of Memory and Language* 95. 89–101. DOI: <https://doi.org/10.1016/j.jml.2017.02.002>
- McElreath, Richard. 2020. *Statistical rethinking: A bayesian course with examples in r and stan*. CRC press. DOI: <https://doi.org/10.1201/9780429029608>

- Moltmann, Friederike. 1997. *Parts and wholes in semantics*. Oxford: Oxford University Press.
- Moltmann, Friederike. 2005. Part structures in situations: The semantics of *Individual and Whole*. *Linguistics and Philosophy* 28. 599–641. DOI: <https://doi.org/10.1007/s10988-005-2142-2>
- Nouwen, Rick. 2012. Plurality. In Maria Aloni and Paul Dekker (eds.), *Cambridge handbook of semantics*. Oxford: Oxford University Press.
- Pagliarini, Elena, Gaetano Fiorin, and Jakub Dotlačil. 2012. The acquisition of distributivity in pluralities. 387–399. Cascadilla Press.
- Parsons, Terence. 1990. *Events in the semantics of english: A study in subatomic semantics*. Cambridge, MA: MIT Press.
- Pickering, Martin J., and Victor S. Ferreira. 2008. Structural priming: A critical review. *Psychological Bulletin* 134. 427–459. DOI: <https://doi.org/10.1037/0033-2909.134.3.427>
- Piñango, Maria Mercedes, Aaron Winnick, Rashad Ullah, and Edgar Zurif. 2006. Timecourse of semantic composition: The case of aspectual coercion. *Journal of psycholinguistic research* 35. 233–244. DOI: <https://doi.org/10.1007/s10936-006-9013-z>
- Pylkkänen, Liina, and Brian McElree. 2006. The syntax-semantic interface: On-line composition of sentence meaning. In Matthew Traxler and Morton Ann Gernsbacher (eds.), *Handbook of psycholinguistics*, 537–577. New York: Elsevier. DOI: <https://doi.org/10.1016/B978-012369374-7/50015-8>
- Roberts, Craige. 1990. *Modal subordination, anaphora, and distributivity*. New York and London: Garland Publishing.
- Safir, Ken, and Tim Stowell. 1988. Binominal *each*. In *Proceedings of the North East Linguistic Society (NELS)* 18. 426–450.
- Scha, Remko. 1981. Distributive, collective, and cumulative quantification. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof (eds.), *Formal methods in the study of language. part 2*, 483–512. Amsterdam: Mathematical Centre Tracts.
- Schwarzschild, Roger. 1993. Plurals, presuppositions and the sources of distributivity. *Natural Language Semantics* 2. 201–248. DOI: <https://doi.org/10.1007/BF01256743>
- Schwarzschild, Roger. 1996. *Pluralities*. Dordrecht: Kluwer. DOI: <https://doi.org/10.1007/978-94-017-2704-4>
- Scontras, Gregory, and Noah D Goodman. 2017. Resolving uncertainty in plural predication. *Cognition* 168. 294–311. DOI: <https://doi.org/10.1016/j.cognition.2017.07.002>
- Syrett, Kristen, and Julien Musolino. 2013. Collectivity, distributivity, and the interpretation of plural numerical expressions in child and adult language. *Language Acquisition* 20. 259–291. DOI: <https://doi.org/10.1080/10489223.2013.828060>
- Szabolcsi, Anna. 2011. *Quantification*. New York: Cambridge University Press.
- Trueswell, John, Michael Tanenhaus, and Susan Garnsey. 1994. Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language* 33. 285–318. DOI: <https://doi.org/10.1006/jmla.1994.1014>
- de Vries, Hanna. 2015. Shifting sets, hidden atoms: the semantics of distributivity, plurality and animacy. Doctoral Dissertation, Utrecht University.
- de Vries, Hanna. 2017. Two kinds of distributivity. *Natural Language Semantics* 25. 173–197. DOI: <https://doi.org/10.1007/s11050-017-9133-z>
- Winter, Yoad. 2000. Distributivity and dependency. *Natural Language Semantics* 8. 27–69. DOI: <https://doi.org/10.1023/A:1008313715103>
- Winter, Yoad. 2001. *Flexibility Principles in Boolean Semantics*. Cambridge, MA: MIT Press. DOI: <https://doi.org/10.7551/mitpress/3034.001.0001>
- Yoon, Youngeun. 1996. Total and partial predicates and the weak and strong interpretations. *Natural language semantics* 4. 217–236. DOI: <https://doi.org/10.1007/BF00372820>
- Zimmermann, Malte. 2002. Boys buying two sausages each: On the syntax and semantics of distance-distributivity. Doctoral Dissertation, University of Amsterdam.

TO CITE THIS ARTICLE:

Dotlačil, Jakub and Adrian Brasoveanu. 2021. The representation and processing of distributivity and collectivity: ambiguity vs. underspecification. *Glossa: a journal of general linguistics* 6(1): 6. 1–22. DOI: <https://doi.org/10.5334/gjgl.1131>

Submitted: 09 October 2019
 Accepted: 19 September 2020
 Published: 19 January 2021

COPYRIGHT:

© 2021 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

Glossa: a journal of general linguistics is a peer-reviewed open access journal published by Ubiquity Press.