Groups versus covers revisited: Structured pluralities and symmetric readings*

Brian Buccola
Department of Linguistics & Languages
Michigan State University

Jeremy Kuhn
Institut Jean Nicod
Département d'Études Cognitives
ENS, EHESS, CNRS, PSL University

David Nicolas
Institut Jean Nicod
Département d'Études Cognitives
ENS, EHESS, CNRS, PSL University

February 24, 2021

Abstract

A number of natural language constructions seem to provide access to structured pluralities — that is, pluralities of pluralities. A body of semantic work has debated both how to model this additional structure and the extent to which it depends on pragmatics. In this article, after controlling for the distinction between ambiguous and underspecified interpretations, we present new data showing that structured pluralities are sometimes but not always available, depending on the form of the plural noun phrase used. We show that these results challenge two longstanding theories of plurality. Finally, we sketch two different ways to account for these data and describe some of the diverging predictions they make.

Keywords: ambiguity, covers, experimental linguistics, groups, inverse linking, plural dynamic semantics, plurals, reciprocity, underspecification

* Authors are listed in alphabetical order. This work benefited from discussion at the Workshop on Cross-Linguistic Semantics of Reciprocals in Utrecht. We would like to thank in particular Denis Paperno, Filipe Hisao Kobayashi, and two anonymous reviewers for *Natural Language Semantics*. The research leading to these results received support from ERC FP7 grant 313610 (SemExp), ERC H2020 grant 788077 (Orisem), and ANR-17-EURE-0017 (FrontCog).

1 Introduction

1.1 Structured pluralities

Under Link's (1983) classic analysis of plurals, a plurality of individuals is represented by their mereological sum, which is another individual. All these objects are entities of type *e*, and the relation between them is represented by algebraic structure in the ontology. This kind of analysis has the advantage that plural terms can be assigned a semantics that accounts simply for natural language entailments (including parallels between plurals and mass terms), without inflating the type theory (see e.g. Gillon 1992; Link 1998; Champollion 2017; Florio and Nicolas 2020; Landman 2020). This mereological (as opposed to set-theoretic) analysis has an important characteristic: while sets can be nested, giving rise to sets of sets, mereological sums cannot. Mereological sums are thus flat, unstructured.

However, as underscored by Landman (1989), a number of natural language constructions seem to provide access to structured pluralities — that is, pluralities of pluralities — visible in the purported readings of sentences with these expressions. For example, the expression *the French students and the Italian students* may *prima facie* seem to refer to a plurality of two pluralities. Subsequent linguistic work has thus debated how to model this additional structure, as well as the extent to which it depends on pragmatics (see Landman 1989; Gillon 1987; Lasersohn 1989; Schwarzschild 1996; Wohlmuth 2018; Grimau 2020, among others). Philosophical work has similarly been interested in the representation of structured pluralities in language and logic (see Oliver and Smiley 2004, pp. 647–665; McKay 2006, pp. 46–53; Linnebo and Nicolas 2008; Grimau 2019, among others). The importance of this issue has been somewhat obscured, though, by subsequent disagreements about the precise empirical landscape.

In this article, we make contributions to this debate on several fronts. Methodologically, we observe that any discussion about what *readings* a sentence has must make a conceptual distinction between ambiguity and underspecification; choosing test sentences accordingly is essential to avoid empirical confounds. Empirically, our primary finding is that structured pluralities are sometimes but not always available, depending on the form of the plural noun phrase used. Strikingly, we show that no existing proposal is able to capture the full range of data. Theoretically, we provide two explicit analytical directions that can be adopted to analyze the facts described. On one analysis, these readings arise via structural mechanisms associated with certain plural noun phrases, while on the other, language plays a distinguished role in the management of context. Depending on which analytical direction is chosen, these findings bear on theories of DP-internal scope-taking, or on theories of dynamic semantics.

1.2 Symmetric readings

A sentence involving a structured plurality is provided in (1), which contains a conjunction of two plural terms. The sentence can be true in a context in which the French students hit the Italian students and the Italian students hit the French students. On standard, mereological theories of plurality (e.g., Link 1983), conjunction of two pluralities generates a flat plurality, so how does the compositional semantics gain access to the two sub-pluralities that hit each other?

(1) The French students and the Italian students hit each other.

Landman (1989) introduces a group-forming operator (↑) that may apply to plural noun phrases. Thus, ↑[*the French students*] *and* ↑[*the Italian students*] denotes a plurality containing two atomic groups. On this logical form, (1) is true if members of each group hit members of the other one.[1]

Schwarzschild (1996) proposes an alternative analysis on which the interpretation of sentences with plural arguments depends on the choice of covers of their denotations (i.e., ways of dividing the pluralities into subpluralities). Any cover that is recoverable from context can in principle serve this role. The reading above is obtained with a cover of the students (here, explicitly mentioned) that divides them into two pluralities: the French students and the Italian students.

Schwarzschild's semantics generates many more readings as well, since every possible cover yields a possible reading, and possible covers are restricted only by pragmatic factors. This is in sharp contrast with Landman's analysis, which is restricted by the syntactic structure of plural noun phrases, according to which each plural noun phrase may denote a group. In favor of his own analysis, Schwarzschild argues that sentences like (2) have an equivalent reading, but that there is no relevant node in the logical form at which a group-forming operator can attach. For sentence (3), too, he suggests that such a reading exists, arguing that (3) can be true in a situation like the one described above.

(2) The students from the two countries hit each other.

(3) The students hit each other.

Let us call the putative reading described above the "symmetric" reading (cf. Winter and Scha 2015; see also the "plurality-collective uses" of Grimau 2020). The empirical question we address in this article is which, if any, of sentences (1–3) genuinely have a symmetric reading. We argue that genuine ambiguity between two or more readings must be dissociated from (mere) underspecification, and

1 In Landman's full system, a further ↑ operator would apply outside this structure, allowing the noun phrase to be the subject of the collective predicate.

we present experimental results to that effect. In Section 2, we show that, contra Landman (1989), sentence (2) *does* have a reading equivalent to the one in (1). But we also show that, contra Schwarzschild (1996), sentence (3) does *not* have this reading, even in a context in which the relevant cover is highly salient. Our results thus pose a challenge to two longstanding views on plurals.

In Section 3, we address the theoretical question of how to account for the symmetric reading of (2) while excluding an analogous reading for (3). We present two potential directions for analysis: first, by enriching Landman (1989) with scope-taking; second, by restricting Schwarzschild (1996) using dynamic semantics. We discuss the diverging predictions these analyses would make.

Finally, we note that symmetric readings are reminiscent of so-called "intermediate readings", which figured prominently in debates about the semantics of plural noun phrases (e.g. Gillon 1987; Lasersohn 1989; Gillon 1990a). Importantly, however, symmetric readings are *not* intermediate readings. Rather than rehash the debate with respect to intermediate readings, we focus here on symmetric readings, where our novel empirical conclusions are robust. In Section 4, we compare our methodologies and results with work on intermediate readings and discuss possible extensions.

2 The empirical and theoretical landscape

2.1 Ambiguity versus underspecification

The theories described above make claims about what *readings* are available for a given sentence. Importantly, the fact that a sentence can be judged true in a particular context does not by itself mean that this context corresponds to a distinct reading of the sentence. In cases of ambiguity, the situations in which a sentence is true under one reading differ from those in which it is true under another reading. Consider the sentence *Sue has a bat*. Under one reading, it is true if Sue has a flying mammal; under another reading, it is true if Sue has a baseball bat. Under a given reading, the sentence may be true in one of these situations but false in the other situation. In cases of underspecification, by contrast, a sentence has a single, underspecified reading, which is true in different situations. For instance, the sentence *Sue has a sibling* is underspecified regarding whether Sue has a brother or a sister. Under this reading, the sentence is true in either of the two situations. In some cases, one reading may be strictly stronger than another. For *Everyone loves someone*, the inverse scope reading entails the surface scope reading; thus, looking solely at the contexts in which the sentence may be judged true is thus not sufficient to determine whether the inverse scope reading exists.

In the case at hand, reciprocal sentences have been claimed to allow a very weak reading based on an "inclusive alternative ordering" (Dalrymple et al. 1998). Under this reading, sentence (3) could be paraphrased as "Each student either hit or was hit by some other student." The sentence would then be true in the "symmetric" situation described above, but this would not indicate that the sentence has a distinct symmetric reading.

To detect genuine ambiguity, one should consider not only when the sentence is true, but also when it is false. Specifically, if a sentence is ambiguous between several readings, then there may be situations in which it is judged true under one reading, and false under another (Gillon 1990a, 2004). In practice, speakers' intuitions about truth and falsity turn out to be not so clear when focusing on simple sentences like (1). To alleviate this problem, one can consider sentences with ellipsis and negation, such as (4), uttered in a context in which two separate covers are relevant.[2]

(4) *Context: This class has only French and Italian students. On Monday, a fight broke out: the French students hit the Italian students, and the Italian students hit the French students. On Tuesday, another fight broke out, but this time within the two groups: the French students hit one another, and the Italian students hit one another.*
On Monday, the French students and the Italian students hit each other, but not on Tuesday.

If sentence (1) has the reading characterized at the outset, then the sentence on this reading should be judged true with respect to Monday, but false with respect to Tuesday. The full sentence in (4) should thus be able to be judged true. By contrast, if sentence (1) has only a single, underspecified reading, then it should be judged true on both Monday and Tuesday, so sentence (4) should be false. Structures of this form thus provide a way to test the existence of symmetric readings for sentences (1–3).

2.2 Experimental results

These judgments were tested in an online experiment with 30 native speakers of English. Eight short context paragraphs were created, paired with three sentences

2 Note that this test holds equally well for structural ambiguity and pronominal ambiguity (as in Schwarzschild's analysis). In sentence (i), for example, the pronoun may have either Bill or John as an antecedent, but whichever it is, it must be the same individual on both Monday and Tuesday.

(i) On Monday, Bill told John that he would win, but not on Tuesday.

containing the target NP structures: "the As and the Bs", "the Xes from the two Ys", and "the Xes". In order to judge the availability of the target reading, subjects were asked to read each context paragraph, then evaluate (on a continuous scale) whether each of the three sentences could be used as a true description of that context. A full paradigm is presented in (5).

(5) *Context: This zoo has two types of tigers — Indian tigers and Chinese tigers — who typically live together. In April, there were two special exhibits, one on Indian animals and one on Chinese animals, so the zookeeper separated the tigers into two groups according to their country of origin. In May, it was mating season, so to have careful control over breeding, the zookeeper again separated the tigers into two groups, but this time by sex.*

   a. In April, the zookeeper separated the Indian tigers and the Chinese tigers, but not in May.

   b. In April, the zookeeper separated the tigers of the two countries, but not in May.

   c. In April, the zookeeper separated the tigers, but not in May.

In the present experiment, we focused only on sentences that can support symmetric interpretations. As a result, all sentences had a reciprocal flavor, even if only one test sentence contained an overt reciprocal. Nevertheless, the examples tested involved various linguistic phenomena: *same*, *each other*, *together*, the subjects of verbs (*interact*, *exchange*, *shake hands*), and the objects of verbs (*separate*, *connect*). Full details on methodology and analysis are provided in the Appendix.

Judgments on the eight trios of sentences showed that both (a) sentences and (b) sentences can be judged as true in the relevant context, but that the (c) sentences cannot, as shown in Figure 1. Statistically, a significant difference was found between (a) and (c) sentences and between (b) and (c) sentences, but not between (a) and (b) sentences. This suggests that the (a) and (b) sentences share a reading — the symmetric reading discussed for (1) — that the (c) sentences don't have. This contradicts the predictions of both Landman's analysis and Schwarzschild's analysis, as summarized in Table 1. Note that it is unlikely that Schwarzschild could explain the data via pragmatics, since the (c) sentences were always presented following a rich context that made the relevant cover salient.

Finally, we observe that many of the tested examples involve reciprocal-like meanings, and that the interpretation of reciprocals is itself subject to a number of semantic and pragmatic effects (see Dalrymple et al. 1998; Dotlačil 2013; Sabato and Winter 2012; Poortman et al. 2018; Haug and Dalrymple 2020, among others). The present pattern of results is nevertheless not specific to reciprocals. First, we note that one trio of sentences involved the predicate *work together*, which resists

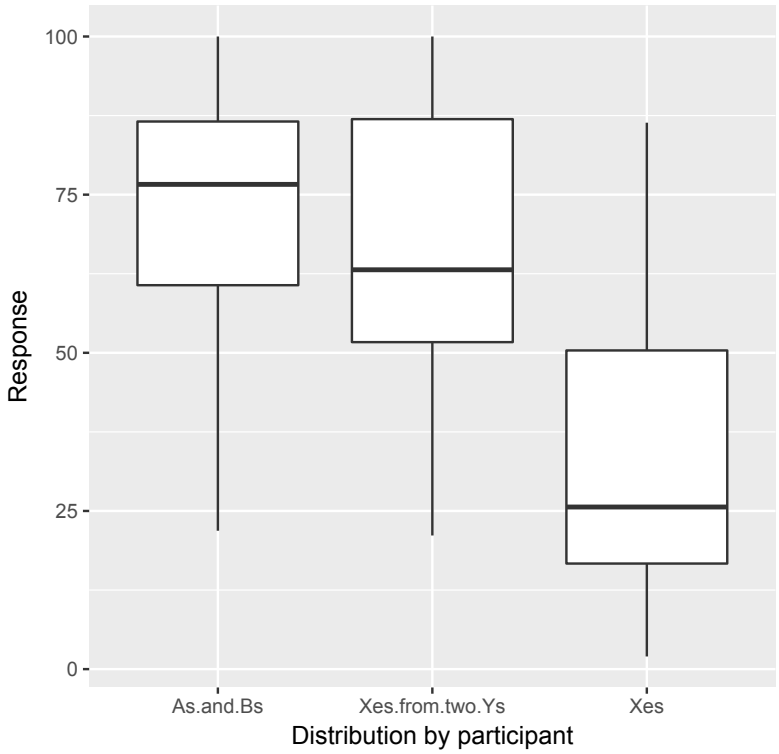

Figure 1: Distribution of responses by participant for the three NP structures on a scale from 0 to 100.

the addition of an overt *each other*. More importantly, the present results arise from manipulating the noun phrase, not the verb phrase. Thus, although the ultimate analysis should of course interface with a semantics of reciprocals, a theory of reciprocals is not sufficient to explain the pattern of results found here. For example, many theories of reciprocals derive the precise semantics of a reciprocal expression from world knowledge about the predicate (e.g., Dalrymple et al. 1998; Poortman et al. 2018). However, any of these theories must then apply this reciprocal meaning to its noun phrase argument; if two potential arguments have the same denotation, then the resulting sentential semantics is the same. Our results show that the noun phrases above do *not* yield identical sentential meanings. The pattern of results thus bears on the semantics of plural noun phrases, not just on the semantics of reciprocals.

		Landman	Schwarzschild	Results
a.	the As and the Bs	✓	✓	✓
b.	the Xes from the two Ys	*	✓	✓
c.	the Xes	*	✓	*

Table 1: Summary of experimental results in relation to existing theories. Check marks in the first two columns indicate where each theory predicts the sentence to allow a symmetric reading. Check marks in the final column indicates where a symmetric reading was in fact found to be available.

3 Directions for analysis

These results paint a picture that is challenging to all current theories. In particular, we find that the symmetric reading of (1) corresponds to a distinct logical form that is not derivable by simple pragmatic means. On the other hand, this reading is not likely to be due to group-forming operators, nor to "generalized conjunction" (Partee and Rooth 1983), since the reading is also available for (2), whose subject is a single noun phrase, without conjunction. How are we to analyze (2)? Here, we sketch two analytical directions — one an enrichment of Landman (1989); the other a revision of Schwarzschild (1996) — then describe the diverging predictions they make.

3.1 Group-forming operators plus scope

On Landman's analysis, we would ideally like to assign the noun phrase in (6a) an interpretation equivalent to the one in (6b).

(6) a. the students from the two countries

b. $\uparrow$[the students from France] $\oplus$ $\uparrow$[the students from Italy]

Such an interpretation can be obtained by combining Landman's group-forming operator with a mechanism of scope taking and the operation of "collectivity raising" from Winter (2001). We assume that (6a) is assigned the structure in (7), in which the DP *the two countries* has been quantifier-raised outside of a structure containing a group-forming operator. Following Kobele (2010)'s analysis of inverse linking, this structure is assigned the interpretation in (8).

(7)

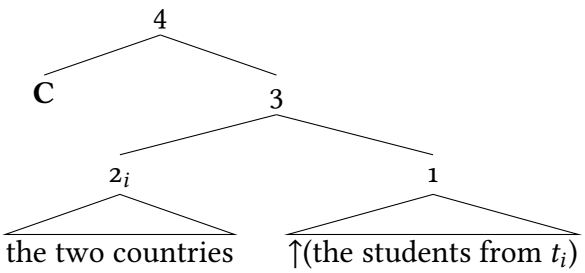

(8) $\mathbf{C}(\lambda P.\text{the two countries}(\lambda t.P(\uparrow \text{the students from}(t))))$

Node 1 denotes the (atomic) group of students from country i. Node 2 is a generalized quantifier that is true of all predicates that contain each of the two contextually salient countries (here, France and Italy). Using the compositional system of Kobele (2010), Node 3 returns another generalized quantifier: the set of predicates that contain both the atomic group of students from France and the atomic group of students from Italy. Finally, we apply Winter (2001)'s **C** operator, defined in (9), which transforms a generalized quantifier into a (lifted) plurality.

(9) a. $\min = \lambda Q.\lambda A.Q(A) \wedge \forall B \in Q[B \subseteq A \rightarrow B = A]$

b. $\mathbf{E} = \lambda A.\lambda P.\exists X[A(X) \wedge P(X)]$

c. $\mathbf{C} = \lambda Q.\mathbf{E}(\min(Q))$

Given a generalized quantifier, the function "min" returns the set of all its minimal predicates. Thus, applied to Node 3, it returns the singleton set containing the set $\{\uparrow$ (students from France), $\uparrow$ (students from Italy)$\}$. This is precisely the plurality desired in (6b). (Winter (2001) models pluralities as sets instead of sums, but the translation can be made easily.) Existential raising in (9b) asserts that there is some such plurality that has the property denoted by the predicate.[3] With the predicate *hit each other*, the resulting sentence is true if the group of students from France hit the group of students from Italy, and vice versa.

This analysis thus explains our data by enriching Landman's framework with more recent hypotheses regarding scope-taking and the relation between generalized quantifiers and plurality.

3.2 Covers plus dynamic semantics

A second strategy retains Schwarzschild (1996)'s analysis in terms of covers, but places restrictions on the pluralities that can be involved. We propose to use the framework of dynamic semantics in order to determine what singular and

3 Existential raising plays a less trivial role in cases with indefinites, where "min" does not return a singleton set, as in *The students from two countries hit each other.*

plural discourse referents are recoverable from a given discourse context. One thread of work on dynamic semantics focuses on how plural discourse referents are introduced and manipulated (van den Berg 1996; Nouwen 2003; Brasoveanu 2008). On these theories, when one plurality is placed in a semantic relation with another plurality, the dynamic system represents not only the two pluralities, but also the thematic relation between them. Thus, the DP in (10a) generates an information state like the one in (10b); here, horizontal rows indicate that the "from" relation holds between the values of x and y.

(10) a. the^{y} students from the^{x} two countries

b. G:

x	y
France	student 1
France	student 2
Italy	student 3
Italy	student 4

Recent work has argued that these semantic associations — and the subpluralities that are created by them — can be accessed by linguistic items elsewhere in the sentence (Dotlačil 2013; Kuhn 2017). One can modify Schwarzschild's analysis to be similarly sensitive to the relations established in the discourse representation. Schwarzschild's analysis involves two variables: a plurality, and a cover over that plurality. For a plural information state G, we let $G|_{x=d}(y)$ be the set of values that y takes on those rows that map x to d. Collecting the sets as d ranges over the values of x provides a cover of y with respect to x. For the information state in (10b), $G(y/x)$ generates the cover $\big\{\{\text{student 1}, \text{student 2}\}, \{\text{student 3}, \text{student 4}\}\big\}$.

(11) $G(j/i) = \big\{S : \exists d[d \in G(i) \wedge G|_{i=d}(j) = S]\big\}$ (Kuhn 2017)

By restricting Schwarzschild's cover variables to only those pluralities that are dynamically accessible, we rule out the cases of overgeneration that arise from a purely pragmatic theory. In particular, if no semantic relation is established between two pluralities, then no dependency is established in the information state. As an example, the discourse in (12a) produces an information state as in (12b), which encodes a trivial relation in which every student is associated with every country. (Note: The fact that each student is associated with each country is just a result of the definition of "random assignment" in van den Berg (1996); it does not entail any relation between the students and the countries.)

(12) a. Two^{x} countries are represented in the class. The^{y} students hit each other.

b. G :

x	y
France	student 1
France	student 2
France	student 3
France	student 4
Italy	student 1
Italy	student 2
Italy	student 3
Italy	student 4

For this information state, $G(y/x) = \big\{\{\text{student 1}, \text{student 2}, \text{student 3}, \text{student 4}\}\big\}$, and $G(y/y) = \big\{\{\text{student 1}\}, \{\text{student 2}\}, \{\text{student 3}\}, \{\text{student 4}\}\big\}$, but no choice of variables will provide the necessary cover for the relevant reading. This predicts that the second sentence in (12a) cannot receive a symmetric reading.

This analysis thus explains our data by restricting Schwarzschild's framework using recent developments on the dynamic semantics of plurals.

3.3 Predictions of the two directions

The two analytical directions make differing predictions on several fronts.

Cumulative readings. First, we observe that the DP in (13) exhibits a cumulative interpretation between the students and the two countries: each of the students comes from one of the two countries, and each country is the origin of at least one of the students. A relatively common way to derive a cumulative interpretation is via pluralization of the predicate (Beck and Sauerland 2000) — in this case, pluralization of the preposition *from*, as in (14a).[4] The double-star operator is defined in (14b).

(13) the students from the two countries

(14) a. the students **from the two countries

b. $^{**}R = \lambda X \lambda Y . \forall x \in X[\exists y \in Y[R(x)(y)]] \wedge \forall y \in Y[\exists x \in X[R(x)(y)]]$

On the other hand, close inspection of the structure in §3.1 reveals that an equivalent interpretation is derived from a rather different logical form on the scope-taking analysis. On the logical form in (15), the effect of cumulativity is generated by an anaphoric dependency. To paraphrase: "for each of the two countries, include the students *from that country*". Of note, the resulting logical

4 For a recent overview discussing other analyses of cumulativity, see Champollion (2020).

form ends up mirroring the analysis that Winter (2000) proposes for cumulative readings generally, which analyzes *the soldiers hit the targets* as equivalent to *the soldiers hit their targets*.[5]

(15) the two countries λx [the students from x]

As a matter of fact, it may be the case that either of these logical forms is available for the DP in (13). Evidence for the availability of both logical forms can be found by adding a numeral to the DP. Empirically, we observe that the two sentences in (16) can both be used to describe the same situation, in which each state is represented by two senators.

(16) *Context: Each state has exactly two senators.*

a. The twelve senators from those six states voted against the bill.

b. The two senators from those six states voted against the bill.

In order to capture this synonymy, the two sentences must be assigned different structures. On a logical form with pluralization of *from*, neither numeral is in the distributive scope of the other, so we can derive an interpretation of (16a) which also has twelve senators total. On a logical form with an anaphoric dependency, we can derive an interpretation of (16b) with twelve senators total, since the numeral *two* appears in the quantificational scope of *those six states*. The fact that both sentences can be used in this context provides evidence in favor of the availability of two distinct logical forms.

Turning to the case at hand, the two analyses in §3.1 and §3.2 make different predictions regarding what logical forms should be available. On the dynamic revision of Schwarzschild (1996), both polyadic quantification and anaphoric dependency will generate a dependency relation, so both will generate an information state of the correct form to provide a non-trivial cover variable. Thus, the symmetric reading should be available on either logical form. On the other hand, the scopal enrichment of Landman (1989) only allows the logical form in (15). Because *the two countries* raises out of the restrictor of the lower NP, the trace that remains below automatically introduces an anaphoric dependency.

The two analyses thus make differing predictions when it comes to (17). The dynamic revision of Schwarzschild (1996) predicts that (17) will allow a symmetric reading in a situation with ten students or with twenty students. The scopal enrichment of Landman (1989) predicts that (17) will only allow a symmetric

5 On the other hand, the analysis in §3.1 does not need to subscribe to other analytical assumptions of Winter (2000). Specifically, it is not committed to the availability of anaphoric dependencies everywhere — only to the fact that anaphoric dependencies may be generated by certain scope-taking operations.

reading in a situation with twenty students.

(17) The ten students from the two countries hit each other.

Scope islands. Another prediction regards the interaction of the symmetric reading with scope islands. Because the scopal enrichment of Landman (1989) relies on the ability of *the two countries* to take wide scope, introducing an island boundary between the two DPs should rule out the necessary logical form. Since relative clauses are generally observed to introduce scope islands, the scopal analysis thus predicts that (18) will not allow the symmetric reading. On the other hand, it is not clear if the predictions are significantly different for the dynamic analysis. Beck and Sauerland (2000) show that island boundaries may also block the cumulative reading that is derived by polyadic quantification, since the double-star operator would need to apply to a constituent that spans an island boundary. Thus, it is possible that the island boundary in (18) simply reduces the acceptability of the cumulative reading across the board.

(18) The students ⟨who come from the two countries⟩ hit each other.

Relatedly, it has been observed that certain prepositions, such as *with*, seem to be resistant to inverse linking interpretations (May and Bale 2006). The scopal analysis would thus predict that (19) should similarly resist a symmetric reading.

(19) *Context: Each student found one of two solutions.*
The students with the two solutions were on opposite teams.

Cross-sentential anaphora. Another prediction on which the two analyses differ regards the availability of cross-sentential anaphora. Since the dynamic analysis is built on a system developed for cross-sentential anaphora, it predicts that a plural pronoun in one sentence should be able to access a plural dependency established in a previous sentence. Under the dynamic analysis, the second sentence in (20) is thus predicted to have a symmetric reading. In contrast, the scopal analysis depends on sentence-internal mechanisms, so does not predict a symmetric reading for (20).

(20) The students come from two countries. They hit each other.

Discussion. Our own judgments on these sentences are not conclusive, and certainly not as clear as our judgments on the original trio of sentences, for which our judgments coincided with the experimental results. We found similar mixed

results for these sentences in an informal survey of eight trained linguists. Thus, if indeed these sentences have intermediate availability of the symmetric reading, it is not fully clear how to interpret this middling acceptability, as both avenues outlined in this section predict categorical classification of the sentences as having the relevant reading or not. Building gradience into the analyses could potentially involve further elaboration of the role of context, as in Grimau (2020), and/or further specification of the processing difficulties of scope-taking, as in Wurmbrand (2018). For example, Grimau (2020) proposes that additional pragmatic mechanisms exist to generate structured pluralities — but only at significant cost.[6]

Furthermore, we note that it is possible to modify each of the theories to generate predictions that converge towards the other. For example, a post-suppositional analysis of numerals (Brasoveanu 2013) may provide a way for the dependency analysis to allow the "ten students" reading of sentences with numerals. Similarly, there is quite a bit of variation in the dynamic literature regarding how cumulative readings are derived (van den Berg 1996; Brasoveanu 2013; Henderson 2014). These analytical choices have the potential to restrict the interpretations available on the dynamic analysis, potentially causing partial convergence with the scope-taking analysis.

4 Comparison with intermediate readings

Intuitively, the symmetric reading of (1) involves two distinct sub-pluralities: the French students, and the Italian students. This is reminiscent of plural sentences that have been argued to have so-called "intermediate readings", in addition to collective and distributive readings. In (21), for example, an intermediate reading could correspond to a situation in which the boys baked one cake and the girls baked a second cake.

(21) The students baked a cake.

a. *Collective reading:* The students baked a cake all together.

b. *Distributive reading:* Each student baked a cake.

c. *Intermediate readings:* Distinct sub-pluralities of students baked a cake.

Whether intermediate readings are genuinely available for sentences like (21) has been the subject of debate (Gillon 1987; Lasersohn 1989; Gillon 1990b; Schwarzschild 1996; Champollion 2016; Wohlmuth 2018). On Schwarzschild (1996)'s

6 Such a mechanism could potentially account for the rich variety of covers, sometimes visually-presented, argued to be available by Schwarzschild (1996) — with the caveat that proper controls should be taken to ensure that these correspond to true *readings* of the sentences in question.

theory, a contextually provided cover is fed to a silent operator that attaches to the verb phrase. The sentence is then true if the original verb phrase holds of each element in the cover. If each element in the cover is a singleton, one gets the distributive reading. If not, one gets an intermediate reading. Crucially, the denotation of the plural noun phrase is flat, e.g. the plurality of students in (21). By contrast, in analyses without a role for context (Landman 1989; Lasersohn 1989), there is no way to derive an intermediate reading.

Notably, though, symmetric readings like the one described for (1) are *not* intermediate readings — the corresponding intermediate reading of (1–3) would be true just in case the French students hit each other and the Italian students hit each other. Still, in order to analyze symmetric readings, Schwarzschild (1996, ch. 6) once again appeals to covers. More precisely, for reciprocal sentences (with or without an overt *each other*), Schwarzschild takes the context to provide two covers: a first cover specifies which pluralities are operative; a second cover specifies which operative pluralities are related to which. In the case of (1), two operative pluralities are taken to be salient: the plurality of French students, and the plurality of Italian students. Just as for intermediate readings, these covers are provided by context, so in a rich enough context, it should be possible to interpret (3) in the same way as (1).

Intermediate readings and symmetric readings are thus generated by the same basic mechanisms: group-forming operators for Landman, contextually provided covers for Schwarzschild. Hence, one would expect general compositional principles to be shared between the two phenomena. If there are differences in the availability of symmetric readings and intermediate readings, these must arise from the specific semantics of reciprocity or distributivity.

For intermediate readings, since the empirical facts have been subject to debate, Wohlmuth (2018) conducted experiments on the interpretations available for sentences like (21), with a plural subject and a predicate containing an indefinite. Wohlmuth reports that such sentences are judged true in an intermediate context just as easily as in a distributive context. She concludes that one should adopt a theory like Schwarzschild's, in which a contextually provided cover is fed to a silent distributivity operator.

Wohlmuth's results differ from those reported here. While Wohlmuth claims that a contextually provided cover is *always* available for intermediate readings, we have argued that this is not the case for symmetric readings. These differences may be explained in one of two ways. First, we note that the methodology adopted in Wohlmuth (2018) does not control for ambiguity versus underspecification as discussed in Section 2.1. It is thus possible that a more consistent set of results would emerge from a parallel paradigm, as in (22).

(22) *Context: On Monday, the French students baked one cake and the Italian students baked a second cake. On Tuesday, the boys baked one cake and the girls baked a second cake.*

a. On Monday, the French students and the Italian students baked a cake, but not on Tuesday.

b. On Monday, the students from the two countries baked a cake, but not on Tuesday.

c. On Monday, the students baked a cake, but not on Tuesday.

If subjects are unable to judge (22c) true in this context (thereby coinciding with our results above), this would suggest that the reading accessed for sentence (21) is the result of underspecification: the sentence would be true just in case there *exists* a cover made of pluralities of students, each plurality of which baked a cake. A new semantic mechanism, independent from what Schwarzschild and Landman propose, would then have to be postulated in order to derive this underspecified reading.[7]

Alternatively, it may be the case that symmetric readings and intermediate readings are not equally available. If this is the case, this difference must emerge from the specific semantics of reciprocity and distributivity.

5 Conclusion

Landman (1989) and Schwarzschild (1996) provide two clear and well-known perspectives on the debate regarding the way in which and the degree to which structured pluralities are represented in natural language. We have provided new data showing that this debate remains open. When we control for the distinct *readings* of plural sentences like (1–3), neither Landman (1989) nor Schwarzschild (1996) is able to capture the full pattern of judgments.

We have seen that either analytical perspective can be modified to capture the observed pattern, but that these modifications entail new theoretical commitments. An analysis based on Landman (1989) can capture the remaining attested reading, but it needs to assume a mechanism of inverse linking plus further type shifters,

7 Bar-Lev (2019) argues that the ability to access such an underspecified reading varies across predicates, depending on the homogeneity of the predicate in question (i.e., whether the plural predication has the inference that all subpluralities behave identically with respect to the predicate). In our own results on symmetric readings, homogeneity plays no role, for two reasons: (i) all the predicates tested are homogeneous, and (ii) the target reading involves a collective predicate acting on a plurality of just two entities (themselves groups or pluralities), so there are no other proper parts of this plurality to which the predicate could apply.

such as Winter (2001)'s **C** operator. An analysis based on Schwarzschild (1996) can rule out the unattested reading, but it needs to assume a rather powerful framework of dynamic semantics. In either case, there remain holes that would need to be filled by future research. For example, the dynamic system would need to be fleshed out with a compositional semantics that makes the necessary discourse referents available for both (1) and (2). Moreover, depending on the empirical facts, either analysis may need to be extended to apply to other linguistic constructions that seem to introduce structured pluralities (Grimau 2020).

We hope that this investigation may serve as the start of a more detailed study of theories of structured pluralities and symmetric readings, of the predictions these theories make, and how best to test these predictions.

6 Appendix: Experimental results

6.1 Methods

Participants. A pilot study was run with 9 participants, of which 3 were removed for failure on attention checks. A power analysis with the remaining participants indicated that 6 participants would be sufficient to find the predicted effect with 0.8 power and a significance level of 0.05. For the full experiment, we requested 30 participants through Amazon's Mechanical Turk. We analyzed the results of 23 participants; the other 7 were removed for failure on attention checks.

Materials and procedure. Eight short context paragraphs were created, paired with three sentences containing the target NP structures: "the As and the Bs", "the Xes from the two Ys", and "the Xes". The eight predicates that were tested are listed in (23), and complete materials are available on OSF: https://osf.io/9wk7c/.

(23) __ hit each other; The zookeeper separated __; __ shook hands; __ worked together; __ read the same books; Sally connected __; __ interacted; __ exchanged letters.

In order to judge the availability of the target reading, subjects were asked to read each context paragraph, then evaluate whether each of the three sentences could be used as a true description of that context. Ratings were made on a continuous scale with endpoints labeled "The sentence **cannot** be true" and "The sentence **can** be true". Instructions explicitly told subjects that ambiguous sentences should be judged as "can be true" if the sentence is true on at least one interpretation. (For example, "Sarah touched the dog with the stick" can be true in a context in which Sarah picked up a stick and used it to touch the dog, even though there is

a second reading of the sentence that is false in this context.) The eight contexts were presented in a random order without fillers, with all three sentences on the same screen as the context paragraph (For motivation for this methodology, see Schütze and Sprouse 2013; Marty, Chemla, and Sprouse 2020).

6.2 Results

Linear mixed-effects models were used to model subjects' responses (Baayen, Davidson, and Bates 2008). Participants and predicates were included as random variables, with the maximal converging model used for each comparison (Barr et al. 2013). The NP structure was included as the predictor variable; for each pair of NP structures, we compared the model with the predictor variable to the model without it. Model fit was assessed using chi-square tests on the log-likelihood values of competing models. These methods were preregistered on OSF: https://osf.io/5edm9/.

Means and standard deviations by participant for the three NP structures, on a scale from 0 to 100, are presented in Table 2, with distribution presented in Figure 1 in the main text. The analysis revealed a significant difference between "the Xes" and each of the other two NP structures ((a) vs. (c): $\chi^2(1) = 15.40$, $p < .001$; (b) vs. (c): $\chi^2(1) = 18.42$, $p < .001$). No significant difference was found between "the As and the Bs" and "the Xes from the two Ys", ($\chi^2(1) = 2.71$, $p = 0.10$).

		Mean	S.D.
a.	the As and the Bs	71.85	20.78
b.	the Xes from the two Ys	66.14	21.10
c.	the Xes	35.26	24.10

Table 2: Means and standard deviations by participant for the three NP structures

References

Baayen, R. H., D. J. Davidson, and D. M. Bates (2008). Mixed-Effects Modeling with Crossed Random Effects for Subjects and Items. In: *Journal of Memory and Language* 59.4, pp. 390–412. DOI: https://doi.org/10.1016/j.jml.2007.12.005.

Bar-Lev, Moshe E. (2019). Specification and Homogeneity in Plural Predication. Unpublished manuscript. URL: https://semanticsarchive.net/Archive/WQ5ODgyY/.

Barr, Dale J., Roger Levy, Christoph Scheepers, and Harry J. Tily (2013). Random Effects Structure for Confirmatory Hypothesis Testing: Keep it Maximal. In: *Journal of Memory and Language* 68.3, pp. 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). In: *Natural Language Semantics* 8.4, pp. 349–371. DOI: https://doi.org/10.1023/A:1011240827230.

Brasoveanu, Adrian (2008). Donkey Pluralities: Plural Information States versus Non-Atomic Individuals. In: *Linguistics and Philosophy* 31.2, pp. 129–209. DOI: https://doi.org/10.1007/s10988-008-9035-0.

Brasoveanu, Adrian (2013). Modified Numerals as Post-Suppositions. In: *Journal of Semantics* 30.2, pp. 155–209. DOI: https://doi.org/10.1093/jos/ffs003.

Champollion, Lucas (2016). Covert Distributivity in Algebraic Semantics. In: *Semantics and Pragmatics* 9.15, pp. 1–65. DOI: https://doi.org/10.3765/sp.9.15.

Champollion, Lucas (2017). *Parts of a Whole: Distributivity as a Bridge Between Aspect and Measurement*. New York, NY: Oxford University Press.

Champollion, Lucas (2020). Distributivity, Collectivity, and Cumulativity. In: *The Wiley Blackwell Companion to Semantics*. Ed. by Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann, and Thomas Ede Zimmermann. Hoboken, NJ: John Wiley & Sons, Ltd. DOI: https://doi.org/10.1002/9781118788516.sem021.

Dalrymple, Mary, Makoto Kanazawa, Yookyung Kim, Sam McHombo, and Stanley Peters (1998). Reciprocal Expressions and the Concept of Reciprocity. In: *Linguistics and Philosophy* 21.2, pp. 159–210. DOI: https://doi.org/10.1023/A:1005330227480.

Dotlačil, Jakub (2013). Reciprocals Distribute over Information States. In: *Journal of Semantics* 30.4, pp. 423–477. DOI: https://doi.org/10.1093/jos/ffs016.

Florio, Salvatore and David Nicolas (2020). Plurals and Mereology. In: *Journal of Philosophical Logic*. DOI: https://doi.org/10.1007/s10992-020-09570-9.

Gillon, Brendan S. (1987). The Readings of Plural Noun Phrases in English. In: *Linguistics and Philosophy* 10.2, pp. 199–219. DOI: https://doi.org/10.1007/BF00584318.

Gillon, Brendan S. (1990a). Ambiguity, Generality, and Indeterminacy: Tests and Definitions. In: *Synthese* 85.3, pp. 391–416. DOI: https://doi.org/10.1007/BF00484835.

Gillon, Brendan S. (1990b). Plural Noun Phrases and their Readings: A Reply to Lasersohn. In: *Linguistics and Philosophy* 13, pp. 477–485. DOI: https://doi.org/10.1007/BF00630751.

Gillon, Brendan S. (1992). Towards a Common Semantics for English Count and Mass Nouns. In: *Linguistics and Philosophy* 15, pp. 597–639. DOI: https://doi.org/10.1007/BF00628112.

Gillon, Brendan S. (2004). Ambiguity, Indeterminacy, Deixis, and Vagueness: Evidence and Theory. In: *Semantics: A Reader*. Ed. by Stephen Davis and Brendan S. Gillon. New York, NY: Oxford University Press, pp. 157–190.

Grimau, Berta (2019). In Defence of Higher-Level Plural Logic: Drawing Conclusions from Natural Language. In: *Synthese*. DOI: https://doi.org/10.1007/s11229-019-02399-z.

Grimau, Berta (2020). Structured Plurality Reconsidered. In: *Journal of Semantics*. DOI: https://doi.org/10.1093/jos/ffaa012.

Haug, Dag Trygve Truslew and Mary Dalrymple (2020). Reciprocity: Anaphora, Scope, and Quantification. In: *Semantics and Pragmatics* 13.10. DOI: https://doi.org/10.3765/sp.13.10.

Henderson, Robert (2014). Dependent Indefinites and their Post-Suppositions. In: *Semantics and Pragmatics* 7.6, pp. 1–58. DOI: https://doi.org/10.3765/sp.7.6.

Kobele, Gregory M. (2010). Inverse Linking via Function Composition. In: *Natural Language Semantics* 18.2, pp. 183–196. DOI: https://doi.org/10.1007/s11050-009-9053-7.

Kuhn, Jeremy (2017). Dependent Indefinites: The View from Sign Language. In: *Journal of Semantics* 34.3, pp. 407–446. DOI: https://doi.org/10.1093/jos/ffx007.

Landman, Fred (1989). Groups, I. In: *Linguistics and Philosophy* 12.5, pp. 559–605. DOI: https://doi.org/10.1007/BF00627774.

Landman, Fred (2020). *Iceberg Semantics for Mass Nouns and Count Nouns: A New Framework for Boolean Semantics*. Cham, Switzerland: Springer Nature Switzerland AG. DOI: https://doi.org/10.1007/978-3-030-42711-5.

Lasersohn, Peter (1989). On the Readings of Plural Noun Phrases. In: *Linguistic Inquiry* 20.1, pp. 130–134. URL: https://www.jstor.org/stable/4178619.

Link, Godehard (1983). The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretical Approach. In: *Meaning, Use, and Interpretation of Language*. Ed. by Reiner Bäuerle, Christoph Schwarze, and Arnim von Stechow. Berlin, Germany: Walter de Gruyter, pp. 303–323. DOI: https://doi.org/10.1515/9783110852820.302.

Link, Godehard (1998). *Algebraic Semantics in Language and Philosophy*. Stanford, CA: CSLI Publications.

Linnebo, Øystein and David Nicolas (2008). Superplurals in English. In: *Analysis* 68.3, pp. 186–197. DOI: https://doi.org/10.1093/analys/68.3.186.

Marty, Paul, Emmanuel Chemla, and Jon Sprouse (2020). The Effect of Three Basic Task Features on the Sensitivity of Acceptability Judgment Tasks. In: *Glossa: A Journal of General Linguistics* 5.1, pp. 1–23. DOI: https://doi.org/10.5334/gjgl.980.

May, Robert and Alan Bale (2006). Inverse Linking. In: *The Wiley Blackwell Companion to Syntax*. Ed. by Martin Everaert and Henk van Riemsdijk. Malden, MA: John Wiley & Sons, Ltd, pp. 639–667. DOI: https://doi.org/10.1002/9780470996591.ch36.

McKay, Thomas J. (2006). *Plural Predication*. New York, NY: Oxford University Press. DOI: https://doi.org/10.1093/acprof:oso/9780199278145.001.0001.

Nouwen, Rick (2003). Plural Pronominal Anaphora in Context: Dynamic Aspects of Quantification. PhD thesis. Utrecht, Netherlands: Utrecht University. URL: https://hdl.handle.net/1874/630.

Oliver, Alex and Timothy Smiley (2004). Multigrade Predicates. In: *Mind* 113.452, pp. 609–681. DOI: https://doi.org/10.1093/mind/113.452.609.

Partee, Barbara and Mats Rooth (1983). Generalized Conjunction and Type Ambiguity. In: *Meaning, Use, and Interpretation of Language*. Ed. by Reiner Bäuerle, Christoph Schwarze, and Arnim von Stechow. Berlin, Germany: Walter de Gruyter, pp. 361–383. DOI: https://doi.org/10.1515/9783110852820.361.

Poortman, Eva B., Marijn E. Struiksma, Nir Kerem, Naama Friedmann, and Yoad Winter (2018). Reciprocal Expressions and the Maximal Typicality Hypothesis. In: *Glossa: A Journal of General Linguistics* 3.1, pp. 1–30. DOI: https://doi.org/10.5334/gjgl.180.

Sabato, Sivan and Yoad Winter (2012). Relational Domains and the Interpretation of Reciprocals. In: *Linguistics and Philosophy* 25.3, pp. 191–241. DOI: https://doi.org/10.1007/s10988-012-9117-x.

Schütze, Carson T. and Jon Sprouse (2013). Judgment Data. In: *Research Methods in Linguistics*. Ed. by Robert J. Podesva and Devyani Sharma. Cambridge, UK: Cambridge University Press, pp. 27–50.

Schwarzschild, Roger (1996). *Pluralities*. Dordrecht, Holland: Kluwer Academic Publishers.

van den Berg, Martin (1996). Some Aspects of the Internal Structure of Discourse: The Dynamics of Nominal Anaphora. PhD thesis. Amsterdam, Netherlands: University of Amsterdam. URL: https://hdl.handle.net/11245/1.111452.

Winter, Yoad (2000). Distributivity and Dependency. In: *Natural Language Semantics* 8.1, pp. 27–69. DOI: https://doi.org/10.1023/A:1008313715103.

Winter, Yoad (2001). *Flexibility Principles in Boolean Semantics: The Interpretation of Coordination, Plurality, and Scope in Natural Language*. Cambridge, MA: MIT Press.

Winter, Yoad and Remko Scha (2015). Plurals. In: *Handbook of Contemporary Semantic Theory*. Ed. by Shalom Lappin and Chris Fox. 2nd ed. Malden, MA: John Wiley & Sons, Ltd, pp. 77–113. DOI: https://doi.org/10.1002/9781118882139.ch3.

Wohlmuth, Kata (2018). Nonatomic Distributive Readings: An Experimental Perspective. In: *Conference of the Student Organization of Linguistics in Europe (ConSOLE)*. Vol. 26, pp. 139–156. URL: https://www.universiteitleiden.nl/binaries/content/assets/geesteswetenschappen/lucl/sole/console-xxvi.pdf.

Wurmbrand, Susanne (2018). The Cost of Raising Quantifiers. In: *Glossa: A Journal of General Linguistics* 3.1, pp. 1–40. DOI: https://doi.org/10.5334/gjgl.329.