The polysemy of container pseudo-partitives

Karen Duek — University of California, Santa Cruz
Adrian Brasoveanu — University of California, Santa Cruz[1]

Abstract. In a series of experimental studies, we investigate the nature of the meaning multiplicity associated with container pseudo-partitives, e.g., *glass of water* can refer to the container (the glass) or the containee (the portion of water), asking if its distinct interpretations are available simultaneously in copredication environments or if they are in competition. The studies compare the acceptability of container pseudo-partitives to lexical polysemes (*book*) and homonyms (*date*), for which copredication has been claimed to be licit and anomalous respectively, and show that pseudo-partitives behave more like polysemes than homonyms, which is unexpected for current accounts. But we also see that copredication with pseudo-partitives is acceptable to a lesser extent than with polysemes, pointing to the possibility that copredication should be rethought in terms of the gradient costs it imposes, and that the multiple meanings of phrasal polysemous expressions like pseudo-partitives might be asymmetrically available in copredication environments.

Keywords: polysemy, pseudo-partitives, copredication / zeugma, experimental semantics

1. Introduction

This paper is in broad terms concerned with the different shades of meaning multiplicity found in natural language, in particular in the nominal domain. Though pervasive, nominal meaning multiplicity is not uniform. A basic distinction can be made between two major classes of meaning multiplicity: homonymy and polysemy. Whereas homonyms are single linguistic expressions associated with multiple meaning representations that are clearly distinct and unrelated (i.e., this is an instance of ambiguity in the technical, formal-semantics sense of the term), polysemes are associated with a single, complex meaning representation in which multiple aspects of a concept relate to each other in systematic and productive ways.

This distinction is reflected in the fact that the senses of homonyms stand in competition, while the senses of a polyseme are typically available simultaneously. The zeugma, a.k.a. copredication, test targets precisely this difference. The example in (1-a) illustrates the anomaly that results from trying to force the single instance of the homonym *date* to satisfy the selectional requirements of *moldy*, which selects for the fruit sense of *date*, and *sarcastic*, which selects for its human sense. In (1-b), *enthralling* requires *book* to be interpreted as an abstract, propositional object, while *fall* demands that we conceive of *book* as a physical object. But unlike the *date* example, the result is not zeugmatic.

[1]We are grateful to Pranav Anand, Amy Rose Deal, Donka Farkas, the CUSP 6 and UCSC S-circle audiences, as well as the anonymous SuB 2014 reviewers and the SuB 2014 audience. The usual disclaimers apply.

(1) a. #The moldy [*inanimate*] **date** was very sarcastic [*animate*].
b. An enthralling [*abstract*] **book** fell [*concrete*] off the table.

The need for a typology of meaning multiplicity for nouns has been long recognized for lexical items, but less is currently known about the possibilities left open for meaning multiplicity at the phrasal level: the possibility of constructing polysemous representations out of syntactically complex expressions has been less systematically explored. This paper will consider the case of container-containee meaning multiplicity exhibited by English pseudo-partitives (*glass of wine*, *box of books*) as a window into this matter. Pseudo-partitives can be interpreted in at least 3 ways:

(2) a. Marina broke the glass of wine that was on the table. *container*
b. Marina drank the glass of wine that was on the table. *containee / concrete portion*
c. This recipe calls for a glass of wine. *measure*

In (2-a), *break* requires the pseudo-partitive to be interpreted as a container; we will refer to this interpretation as the container reading. In (2-b), to meet the selectional requirements of *drink*, the same phrase must be interpreted as a containee, i.e., a concrete portion of wine; we will refer to this as the containee, or the concrete portion, reading. Finally, in (2-c), there is no requirement that there be any actual container involved, or any actual containee (concrete portion), only that there is an abstract amount of wine that would be enough to fill a glass (the measure reading).

In this paper, we will be concerned only with the first two readings, leaving aside the measure interpretation. We report two experimental studies whose main goal was to examine the container/containee meaning multiplicity exhibited by pseudo-partitives. The experiments aim to distinguish between two families of accounts that attribute this multiplicity to different sources:

(3) Two ways to account for the container/containee meaning multiplicity of pseudo-partitives:
a. **Pseudo-partitives are ambiguous**: the container noun is basically ambiguous between a lexical head (the container reading) and a functional head (the containee reading) and depending on its lexical vs. functional nature, it occupies different syntactic positions and has distinct denotations (Selkirk 1977, Corver 1998, Grimshaw 2007, Landman 2004, Rothstein 2009, Partee and Borschev 2012).
b. **Pseudo-partitives are complex type constructors**: a polysemous denotation is built compositionally (this novel account builds on Pustejovsky 1995 and Asher 2011).

Ambiguity approaches take the phrasal meaning multiplicity of pseudo-partitives to be an instance of homonymy: different meanings are associated with distinct, competing representations, so they should not be simultaneously available. In contrast, a complex-type constructor account assimilates the container/containee multiplicity to polysemy, which predicts that they can be simultaneously

available in the same representation. Thus, the zeugma / copredication test should provide adequate evidence for one or the other type of approaches. Theoretically, what is at stake is whether polysemy is restricted to the lexical domain, or if the grammar offers compositional mechanisms to derive the same sort of meaning flexibility at the phrasal level.

We report here two acceptability-judgment studies in which native English speakers were asked to assess zeugma / copredication sentences with container pseudo-partitives. As baselines for the pseudo-partitive judgments, we asked speakers to also judge zeugmatic sentences that involved lexical homonymy and lexical polysemy. The studies show that container pseudo-partitives are unlike homonyms: copredication with pseudo-partitives does not lead to clear anomaly in the same way that it does with homonyms. The results are compatible with the hypothesis that the output of the semantic composition needed to interpret pseudo-partitives is polysemous in nature. The implication is therefore that complex semantic representations exhibiting meaning multiplicity are not exclusively associated with lexical items, which are inherently conceptually complex, but that particular constructions can manipulate simple nominal meanings to build such complex representations compositionally.

But the studies also show that pseudo-partitives are not exactly like lexical polysemes: we can observe a small effect that distinguishes the clearly acceptable cases of lexical polysemy and the somewhat less acceptable cases of phrasal polysemy. Thus, overall, our results point to the possibility that copredication should be rethought in terms of the gradient costs it imposes.

2. Background

Current accounts of the container/containee ambiguity take each reading to correspond to a distinct syntactic and semantic representation, thereby assimilating the meaning multiplicity exhibited by pseudo-partitives to homonymy. Partee and Borschev (2012), for instance, take the container reading to arise as a consequence of the container word being the head of the nominal phrase / projection. The container word has a relational denotation and takes the containee nominal as its complement. When we predicate something of a pseudo-partitive, that predicate is taken to be true of the container (together with the substance that fills it).

In contrast, the concrete portion and measure readings are derived if the containee word is the head of the nominal phrase and the container word combines with a numeral to form a nominal modifier for the containee noun. Partee and Borschev (2012) recognize that their account is incompatible with the possibility of copredication, which nonetheless seems possible and non-zeugmatic at least between the container and containee / concrete portion readings, as in (4). They suggest that if these facts indeed hold, a more appropriate account of this meaning multiplicity should approximate it to the representation of lexical polysemy.

(4) a. The glass of wine that was a cabernet had a chip on the rim.

b. The glass of wine that had a chip on the rim was a cabernet.

The possibility of copredication for polysemes has motivated theories that significantly enrich (*i*) the structure of the lexicon, such as Pustejovsky (1995), or (*ii*) the system of types and the mechanisms of predication employed in formal semantics (Asher 2011, Asher and Pustejovsksy 2013). For concreteness, we consider the account put forward in Asher (2011) more closely.

Recall that the ability to satisfy multiple, sometimes incompatible, semantic selectional requirements is the most crucial diagnostic tool for nominal polysemy. Likewise, this is often the criterion used to determine the meaning of an ambiguous element, i.e., a homonym, in a particular context. Therefore progress in accounting for homonymy and polysemy is dependent on developing a precise formal notion of what selectional restrictions are and what it means to meet them. In Asher's account, selectional requirements are modeled as presuppositions at the level of types. This is made possible by significantly expanding the system of types to mirror conceptual organization, including subtype relations. For example, a predicate like *drink* specifies that its argument must be not only be an individual, but a liquid, physical entity.

The result is a 'two-tier' model of lexical meaning: lexical meanings are still modeled as lambda terms, the functional 'structure' of which encodes the 'coarse-grained' compositionality-related part of their meaning (just as in Montagovian semantics), but these terms also include fine-grained information about the richly structured types of the λ-bound variables, which is needed to model polysemy, selectional restrictions and related phenomena. An important component of the formalization is a λ-bound variable π of the presuppositional type Π that encodes the typing constraints a predicate places on its arguments. The type of propositions is thus $\Pi \Rightarrow t$: a proposition can be evaluated only with respect to a context that satisfies all of its type presuppositions; if they cannot be satisfied via binding or accommodation, the proposition cannot be evaluated.

The crucial part for us is how the type of a polyseme is established in this framework. Since the polyseme can justify multiple, incompatible type presuppositions, its type will have to be more complex. For example, the type of the object 'contributed' by a noun like *book* must be complex enough to satisfy both a physical-object type presupposition and an abstract-object type presupposition. Such complex types enter semantic composition in a specific way: the complex-type variable (or more generally, term) 'introduced' by a noun like *book* is distinct from the variables it contributes as arguments to other predicates in the sentence, even though the former are related to the latter. For example, *book* continues to contribute a variable of a complex, physical-*cum*-abstract type to the semantic representation, i.e., a variable whose type is the dot-type PHYSICAL • INFO, even after it combines with a predicate like *heavy*, which exclusively selects for a physical-object type. The fact that the dot-type variable continues to be available correctly predicts that a predicate selecting for an abstract-object type can successfully be added to this representation – see, for example, the felicitous example of copredication in (1-b) above.

If container pseudo-partitives behave like lexical polysemes with respect to copredication, it would be desirable to account for this in a similar way. The account could either (*i*) give a dot-type meaning to the container noun (building on Partee and Borschev 2012 and related accounts), or (*ii*) give a dot-type meaning to the entire structure, but not to the container.

Two empirical questions thus arise. The first one is whether container pseudo-partitives in English exhibit a similar pattern of copredication acceptability as lexical polysemes. If so, the second question is whether the container noun on its own is polysemous, or only the full construction is (thus making it a genuine example of phrasal polysemy). If the behavior of pseudo-partitives is close to the one exhibited by polysemes but not identical (which is actually what we will see), this can be taken as modest evidence that pseudo-partitives are an example of phrasal polysemy.

3. Experiment 1

The main goal of Experiment 1 was to investigate how speakers treat copredications with container pseudo-partitives. In order to do so, we first seek to establish that the copredication test distinguishes between lexical polysemy and homonymy when using coarse-grained, offline ('non-real-time') measures, in particular, acceptability judgments. We are unaware of previous studies assessing the acceptability of copredication directly. The most closely related studies in the literature are the eye-tracking studies reported in Frazier and Rayner (1989), who found that the distinction between homonymy and polysemy is reflected in the fine-grained, real-time behavioral measures that eye-tracking provides. In particular, Frazier and Rayner (1989) found that participants committed to a particular sense in the case of homonyms (even in neutral contexts that did not bias them one way or another), but no such commitment was made in the case of polysemes. Assuming offline acceptability judgments can distinguish between homonyms and polysemes, they will serve as baselines when we evaluate the acceptability of copredication with container pseudo-partitives, enabling us to to locate pseudo-partitives on the homonymy-to-polysemy spectrum.

3.1. Design and materials

We tested the acceptability of copredication with homonyms, polysemes and container pseudo-partitives. We manipulated three factors: (*i*) the senses selected by the first predicate, e.g., for a polyseme like *book*, whether the predicate selects for its concrete or abstract sense; (*ii*) the senses selected by the second predicate, which were the same as the senses selected by the first predicate; and finally, (*iii*) the type of meaning multiplicity: homonymy, polysemy, pseudo-partitive with a mass containee (*bottle of water*), or with a count containee (*box of books*). That is, Experiment 1 had a $2 \times 2 \times 4$ factorial design. There were 16 items for homonyms, 16 for polysemes, 16 for pseudo-partitives with count containees, and 20 items for pseudo-partitives with mass containees, for a total of 68 items. Every item was passed through the corresponding 4 conditions (2 senses for the first predicate $\times$ 2 senses for the second predicate) and 4 lists were generated, each of which included every item exactly once, with the items rotated through the 4 conditions across the 4 lists; the participants were rotated through these 4 lists (Latin square design).

The stimuli for this experiment (and the subsequent one) always involved a restricted relative clause modifying the subject nominal, and the two predicates were provided by the relative clause on one hand, and the matrix VP on the other.[2] The nouns and biasing predicates in the homonymy and polysemy conditions were based on the ones used in Frazier and Rayner (1989) whenever possible; modifications were made whenever translating their items into our copredication frame resulted in an unnatural sentence. Homonyms always alternated between an animate and a inanimate sense (e.g., *date*, *bat*), and polysemes always alternated between a sense involving a concrete, physical-object and a sense involving an abstract object 'stored' in the concrete one (*book*, *newspaper*). Example items are provided below.

(5) **Homonymy**

a.	The date that **tasted** bitter was **bought** just yesterday.	*inanimate-inanimate*
b.	The date that **walked** in late was very **rude** to Jane.	*animate-animate*
c.	The date that **tasted** bitter was very **rude** to Jane.	*inanimate-animate*
d.	The date that **walked** in late was **bought** just yesterday.	*animate-inanimate*

(6) **Polysemy**

a.	The novel that got some great **reviews** was a terrifying **thriller**.	*abstract-abstract*
b.	The novel that got **soaked** in coffee was **found in the sale bin**.	*concrete-concrete*
c.	The novel that got some great **reviews** was **found in the sale bin**.	*abstract-concrete*
d.	The novel that got **soaked** in coffee got some great **reviews**.	*concrete-abstract*

(7) **Mass containee**

a.	The jug of lemonade John **broke** had lemons **painted** on it.	*container-container*
b.	The jug of lemonade I **drank** was too **sweet**.	*containee-containee*
c.	The jug of lemonade my grandfather **broke** was too **sweet**.	*container-containee*
d.	The jug of lemonade I **drank** had lemons **painted** on it.	*containee-container*

(8) **Count containee**

a.	The tupperware of cookies made of blue **plastic** is **sealed shut**.	*container-container*
b.	The tupperware of cookies I **baked** today is all **chocolate** chip.	*containee-containee*
c.	The tupperware of cookies made of blue **plastic** is all **chocolate** chip.	*container-containee*
d.	The tupperware of cookies I **baked** today is **sealed shut**.	*containee-container*

Homonyms are expected to be judged as less acceptable in the 2 mismatching conditions than in the matching ones, while polysemes should be judged as equally acceptable in all conditions. Under the ambiguity approach, pseudo-partitives should display the same behavior as the homonyms.

[2]Other structures could be used in the copredication test, e.g., coordinations, adjectival modification, appositives etc. Empirically investigating if these copredication structures yield roughly similar results is left for a future occasion.

Under the complex type approach, pseudo-partitives should behave like the polysemes. A third possibility is that pseudo-partitives match neither lexical polysemes nor homonyms, maybe because the pattern across the 4 conditions is different, and/or because mass and count containees behave in different ways.

An additional complication specific to pseudo-partitives is that because they are syntactically complex, the restrictive relative clause might attach at the level of the lower, containee noun, or at the level of the higher, container noun. If the relative clause attaches low (targeting the containee noun), the stimulus does not necessarily instantiate a case copredication: the relative clause predicates something of the containee, while the matrix clause predicates something of the container. But when the relative clause attaches high (targeting the container noun), the result is definitely a copredication structure of the requisite kind. In what follows, we will treat all four conditions for pseudo-partitives as copredicational, anticipating one of the results of our Experiment 2 indicating that relative clauses in pseudo-partitives do not seem to exhibit such syntactic freedom. But this additional complication leads us to select the container sense as our reference level when we do the statistical analysis of the Experiment 1 and Experiment 2 data.

3.2. Procedure and participants

36 native speakers of English participated in the study. All participants were undergraduate students at UCSC, and completed the study for course credit or extra-credit on a UCSC hosted installation of A. Drummond's IBEX platform (`https://code.google.com/p/webspr/`). They were instructed to rate the sentences presented in isolation on a 5-point Likert scale: -2 (very bad), -1 (fairly bad), 0 (neither good nor bad), 1 (fairly good), 2 (very good). The participants were rotated through the 4 lists of items described above. Each participant rated 138 stimuli (68 items + 70 fillers), the order of which was randomized for each participant.

3.3. Results and analysis

Graphical summaries of the data are provided in Figures 1 and 2. Since the response variable is ordered categorical, we used mixed-effects ordinal probit regression models to analyze the data. All models included the full fixed-effect structure (main effects and interactions) unless otherwise specified, and crossed random intercepts for subjects and items.[3]

[3] All the data summaries / plots / analyses in this paper have been generated / completed using `R` (R Core Team, 2013) and the packages `ggplot2` (Wickham, 2009) and `Ordinal` (Christensen, 2012).

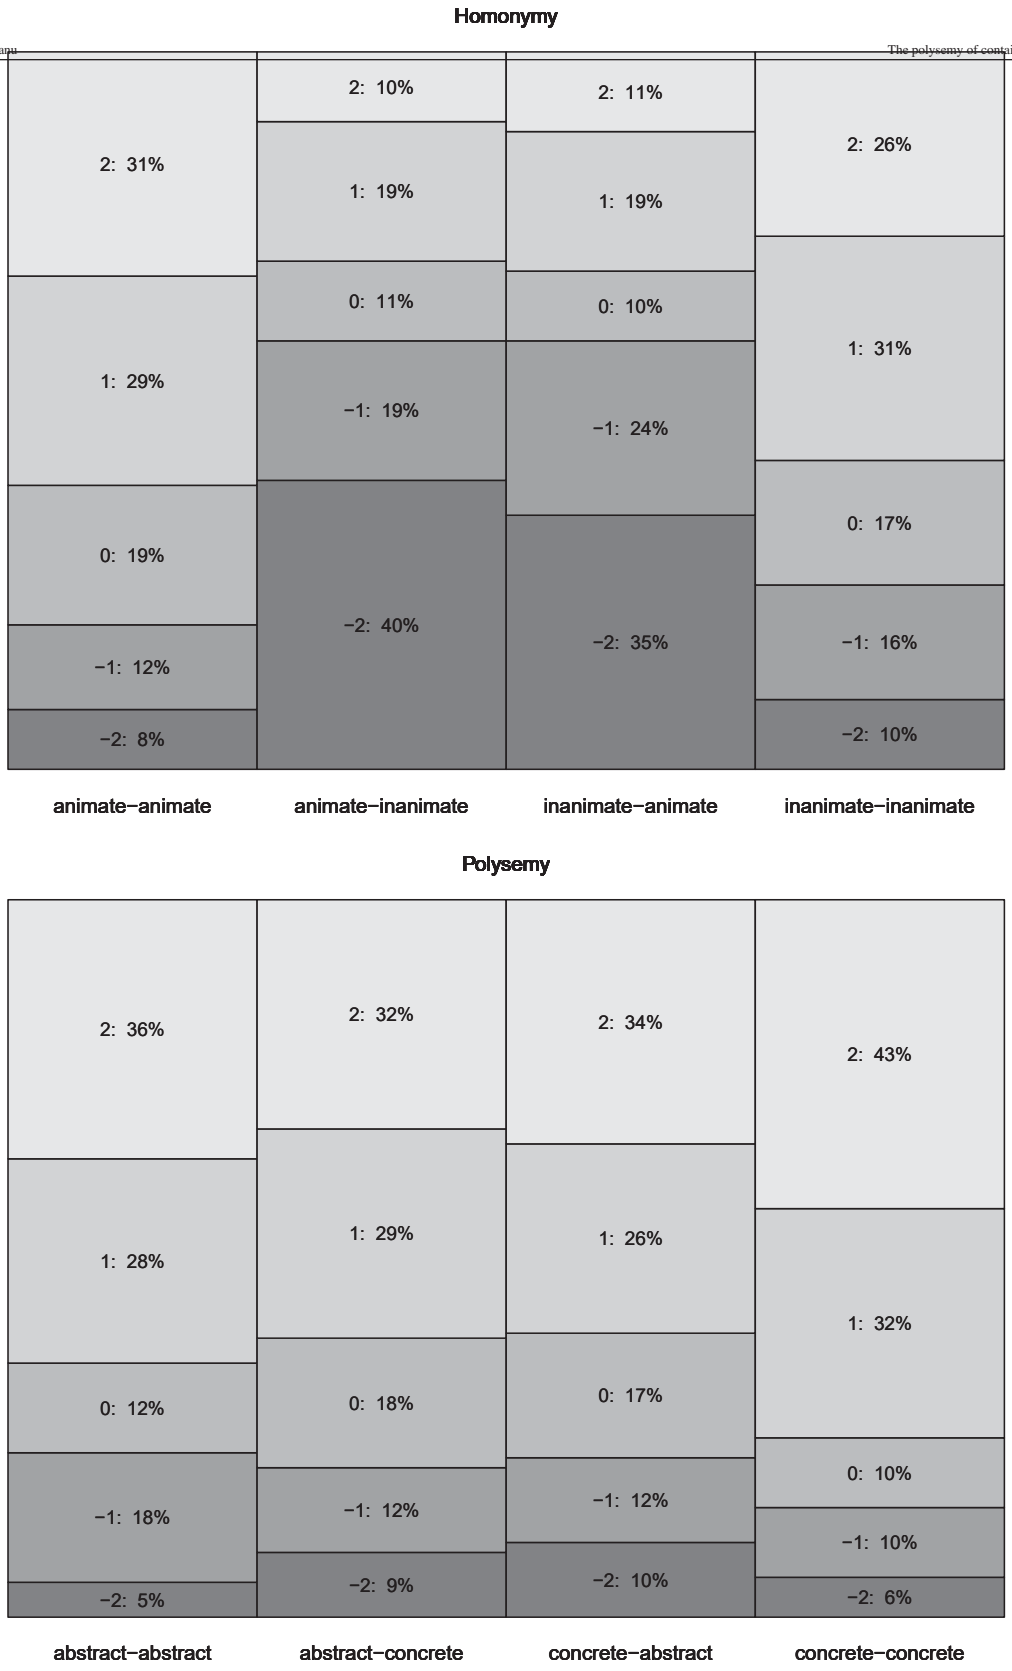

Figure 1: Experiment 1: Data summaries for the homonymy and polysemy conditions.

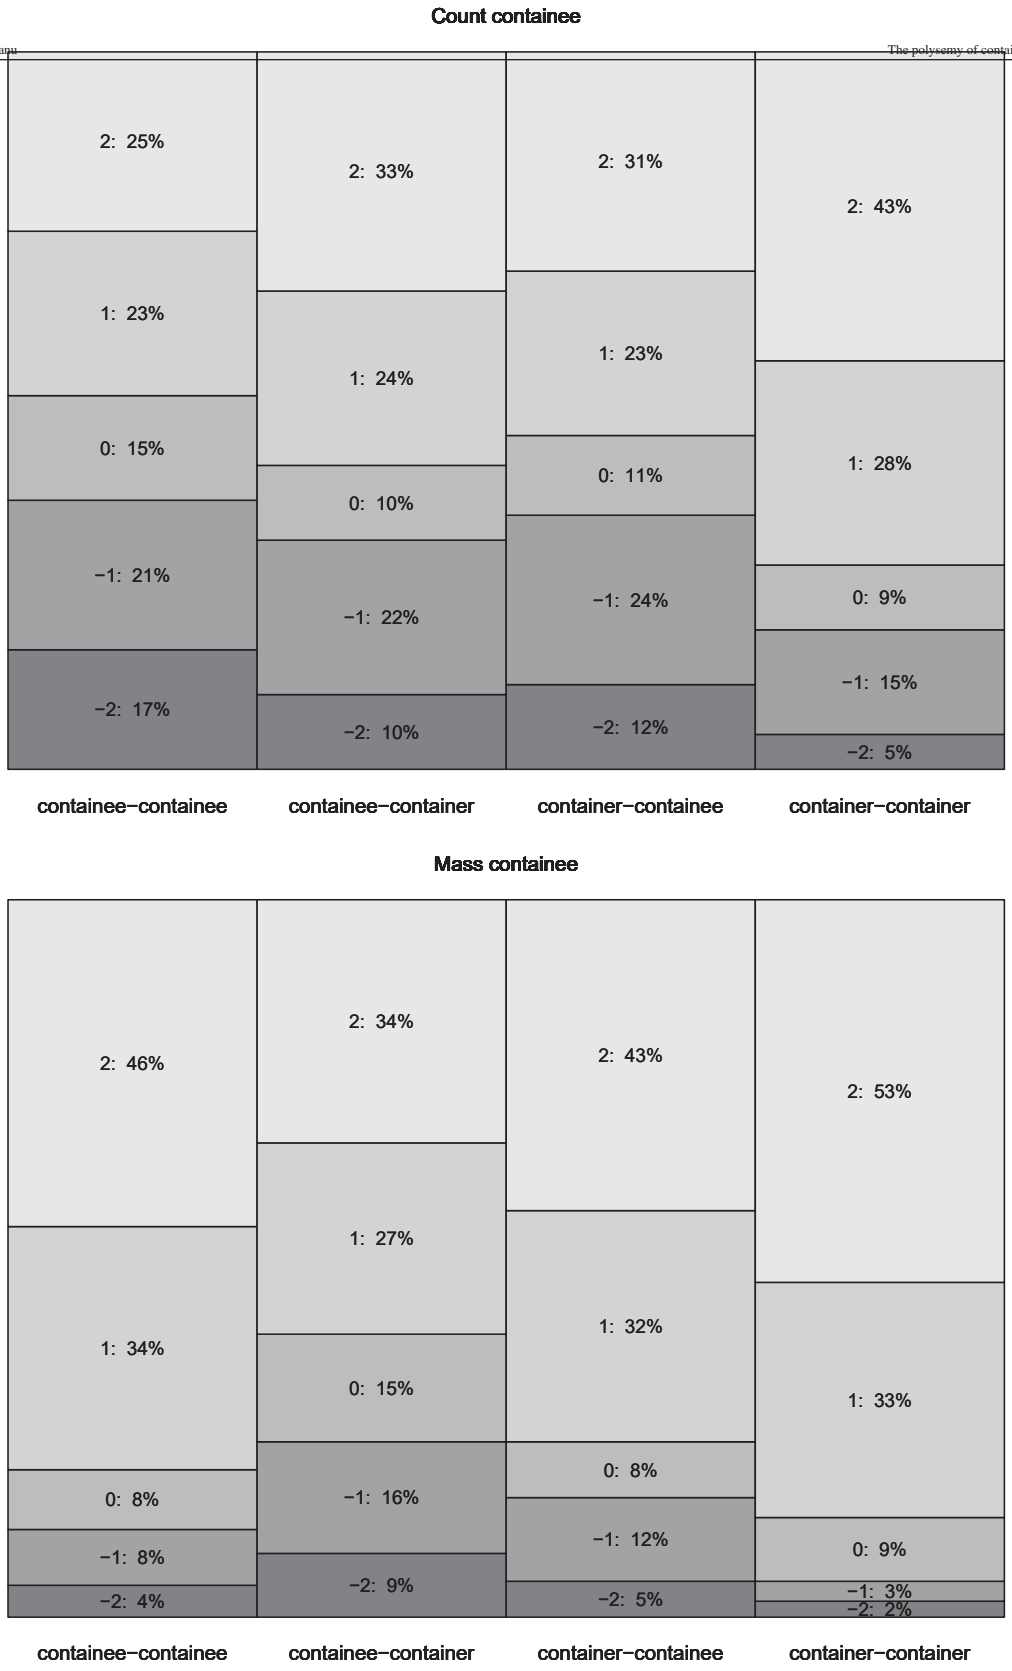

Figure 2: Experiment 1: Data summaries for the count and mass containee conditions.

As expected, homonyms were judged significantly worse in the mismatching conditions than in the matching ones. This can be observed by examining the top panel of Figure 1, e.g., the percentage of unacceptable -2 and -1 ratings given to the mismatching vs. matching conditions. This is confirmed by the statistical analysis of the data. The reference levels for both the first sense (the sense selected by the relative clause) and the second sense (the sense selected by the matrix VP) were set to ANIMATE. There was a main effect of INANIMATE for both the first sense ($\beta = -1.07, SE = 0.13, p = 0.00$) and the second sense ($\beta = -1.18, SE = 0.14, p = 0.00$), and a significant INANIMATE×INANIMATE interaction ($\beta = 2.11, SE = 0.19, p = 0.00$), which basically reversed the cumulative effect of the two main effects, thereby bringing the acceptability of the inanimate-inanimate condition back to the high level of the reference condition (animate-animate).

For polysemy, we set the reference levels for both the first and the second sense to ABSTRACT. The main effects of switching to CONCRETE were non-significant for both the first sense and the second sense (the estimates were almost identical in the two cases: $\beta = -0.12, SE = 0.13, p = 0.35$). That is, we detected no difference between the abstract-abstract (reference) condition and the mismatching abstract-concrete and concrete-abstract conditions. There was however a significant CONCRETE×CONCRETE interaction ($\beta = 0.51, SE = 0.19, p = 0.006$). This result, in conjunction with an inspection of the data summaries in Figure 1, indicates that mismatching conditions are slightly worse than matching conditions for polysemes also, but this difference is much smaller than for homonyms – as shown by the much smaller main & interaction effects for polysemy relative to homonymy.

The picture that seems to emerge is that mismatching conditions are worse than matching conditions across the board, but the differences in acceptability between these conditions fall on a gradient spectrum. Homonyms are at the high end of the spectrum, with large – and introspectively available – differences between matching and mismatching conditions. Polysemes are the low end of the spectrum, with small – and introspectively not (or less) available – differences between matching and mismatching conditions.

We can strengthen our confidence in this hypothesis by grouping the matching conditions together, and also the mismatching conditions, pooling the homonymy and polysemy data, and estimating a mixed-effects probit model with two fixed effects: (*i*) SAME-SENSE (reference level) vs. DIFFERENT-SENSE, and (*ii*) HOMONYMY (reference level) vs. POLYSEMY. All fixed effects (the main effects and their interaction) are significant in this model. There is a main effect of POLYSEMY ($\beta = 0.32, SE = 0.13, p = 0.01$) indicating that even in the matching conditions (the baseline), polysemes are relatively more acceptable. There is a large negative main effect of DIFFERENT-SENSE ($\beta = -1.01, SE = 0.10, p = 0.00$) indicating that zeugmatic sentences significantly decrease acceptability for homonyms, as expected. Finally, there is a significant positive interaction of POLYSEMY×DIFFERENT-SENSE ($\beta = 0.83, SE = 0.13, p = 4 \times 10^{-10}$), which almost – but not quite – reverses the negative main effect of DIFFERENT-SENSE observed with homonyms. These results are compatible with the hypothesis outlined above that the differences in acceptability between matching and mismatching conditions fall on a gradient spectrum, with

homonyms at the high end of the spectrum (large differences) and polysemes at the low end of the spectrum (small differences).

For all pseudo-partitives (both those with a count containee and those with a mass containee), we selected CONTAINER as the reference level for both the first predicate (the restrictive relative clause) and the second predicate (the matrix VP).

Mass-containee pseudo-partitives exhibit the same overall pattern as homonyms and polysemes: mismatching conditions are worse than matching conditions. And the difference between these classes of conditions falls roughly in the middle of the spectrum between homonyms and polysemes. This can be observed by comparing the data summaries for mass containees in the lower panel of Figure 2 with the corresponding summaries in Figure 1. The results of the statistical analysis are compatible with this: there is a main effect of CONTAINEE for both the first sense ($\beta = -0.79, \textit{SE} = 0.12, p = 1.9 \times 10^{-10}$) and the second sense ($\beta = -0.41, \textit{SE} = 0.12, p = 0.001$), and a significant CONTAINEE×CONTAINEE interaction ($\beta = 0.91, \textit{SE} = 0.17, p = 1.9 \times 10^{-7}$). All of these effects have the same direction as the corresponding homonymy and polysemy ones, and their magnitudes are intermediate between the corresponding homonymy and polysemy effects.

Count-containee pseudo-partitives exhibit a slightly different pattern: there is a main effect of CONTAINEE for both the first sense ($\beta = -0.36, \textit{SE} = 0.13, p = 0.006$) and the second sense ($\beta = -0.46, \textit{SE} = 0.13, p = 0.0004$), but no significant CONTAINEE×CONTAINEE interaction ($\beta = 0.19, \textit{SE} = 0.18, p = 0.30$); see also the top panel of Figure 2. This is compatible with count-containee pseudo-partitives exhibiting a polysemous behavior since zeugma / copredication does not lower acceptability. What we seem to observe is an across-the-board preference for the container sense, which is slightly more pronounced for the second predicate (the matrix VP) than for the first predicate (the restrictive relative clause).

The contrast between count and mass containees is further confirmed when we group the matching conditions together, and also the mismatching ones, pool the count and mass data, and estimate a mixed-effects probit model with two fixed effects: (*i*) SAME-SENSE (reference level) vs. DIFFERENT-SENSE, and (*ii*) COUNT (reference level) vs. MASS. We see that there is a positive main effect of MASS ($\beta = 0.65, \textit{SE} = 0.13, p = 1.1 \times 10^{-6}$), which might be due to independent issues like the pattern of agreement with the matrix verb (see the discussion subsection below). But most importantly, we see that the main effect for DIFFERENT-SENSE ($\beta = -0.09, \textit{SE} = 0.09, p = 0.32$) is non-significant, indicating that there is no difference between matching and mismatching conditions for count containees, while the interaction MASS×DIFFERENT-SENSE is significant ($\beta = -0.32, \textit{SE} = 0.13, p = 0.009$), indicating that there is such a difference for mass containees.

3.4. Discussion

Taken together, the results of Experiment 1 confirm the validity of the copredication test in distinguishing homonymy and polysemy. As expected, speakers judge copredication with homonyms

very poorly, and the contrast against matching conditions was quite sharp. Somewhat less expectedly, we still observe some cost of copredication for polysemes, even though the difference between regular predication (the match conditions) and copredication (the mismatch conditions) was much smaller than for homonyms.

The contrast between homonymy and polysemy is large enough to warrant their use as baselines against which to compare the behavior of container pseudo-partitives. However, the fact that copredication was not cost-free for lexical polysemes must qualify our interpretation of the behavior of pseudo-partitives as well. That is, if pseudo-partitives were to behave on a par with polysemes, the expectation should not be that mismatch conditions receive no penalty, but that the size of the effect would be relatively small.

Our results show that overall, container pseudo-partitives pattern more closely with polysemy than homonymy. For count containees, there was no significant interaction effect, i.e., sense matching vs. mismatching does not seem to affect acceptability; there is only an overall preference for the container sense. For mass containees, sense matching vs. mismatching has an effect on acceptability, and this effect seems to be intermediate between the one observed for polysemy and the one observed for homonymy. Generally, this indicates that the relevant readings of the container construction are available simultaneously. It suggests therefore that ambiguity-based accounts that assume distinct syntactic and semantic representations for those readings do not capture the full range of interpretations associated with pseudo-partitives.

However, the patterns of acceptability for pseudo-partitives were more diverse than initially expected. In particular, neither one of the approaches we considered above predicted a sensitivity to whether the containee noun was a mass or bare-plural count noun.

A few potential confounds must be addressed before we can suggest with confidence that the acceptability of copredication for container pseudo-partitives is indeed modulated by the mass/count status of the containee noun. First, while the set of containers used in both conditions overlapped, they were not identical. Second, there was a systematic number mismatch in the count-containee cases: container nouns were always singular and containees were always plural. This meant that the conditions in which the main predicate selected for the plural containee while bearing singular morphology exhibited a somewhat odd number-agreement pattern. For instance, in the sentence *The tupperware of cookies I baked this morning is all chocolate chip*, even though the main predication selects for *cookies*, the copula shows singular agreement morphology with the entire subject. It is possible that the different acceptability pattern exhibited by count-containee pseudo-partitives was a consequence of this confound.

The strength of selectional requirements is a third potentially confounding factor. Many of the predicates did not strongly select for one of the readings, but were more plausibly associated with either the containee or the container reading in the broader context of the sentence in which they appeared. For instance, the item *The pot of curry [that Chris carried] was very fragrant* was de-

signed to be in the condition in which the predicate in the relative clause selected for the container. But while it is clearly pragmatically odd to interpret *carry the pot of curry* as an event of carrying only curry, there is no grammatical clash between the selectional restrictions of *carry* and *curry*.

Finally, to preserve the naturalness of the stimuli, the mismatching conditions (container-containee vs. containee-container) were not mirror images of each other since the predicates were not preserved and simply flipped around. For instance, the reverse mismatching condition for *The pot of curry that Chris carried was very fragrant* was *The pot of curry Mary cooked fell on the floor and broke*. Although both *be fragrant* and *cook* select for the containee reading, it is possible that their selectional 'strength' is different, blurring the effect of selection itself.

4. Experiment 2

4.1. Design and materials

Experiment 2 tested the acceptability of copredications with pseudo-partitives using a different set of items that better controlled for the possible confounds listed above. First, the same container words were used throughout. Second, to avoid the issue of number agreement, containee-selecting predicates were never explicitly number marked (achieved by modalizing the predicate, for example). Third, when sense-biasing predicates were selected, we were more stringent and attempted to select predicates with a much higher bias for one sense or another independently of the particular sentence that the predicate occurred in. Finally, we ensured that for every item, the same predicates were used for the two mismatching conditions container-containee and containee-container, varying only the order in which they occurred. An example item is provided in (9).

(9) a. The container of {gold/diamonds} that was missing a lock has a velvet interior. *container-container*
b. The container of {gold/diamonds} that might have been mined in East Africa could be 24 carat. *containee-containee*
c. The container of {gold/diamonds} that was missing a lock might have been mined in East Africa. *container-containee*
d. The container of {gold/diamonds} that might have been mined in East Africa was missing a lock. *containee-container*

Procedure and participants. 42 native speakers of English recruited online participated, without compensation. As in the previous study, they were instructed to rate the sentences on a 5-point scale from -2 (very bad) to 2 (very good). The procedure was the same as in Experiment 1.

4.2. Results and analysis

Unlike in Experiment 1, the pattern of acceptability we observe in Experiment 2 is not affected by the count/mass status of the containee noun. In particular, we obtain different results when we group the matching conditions together, and also the mismatching ones, pool the count and mass data, and estimate a mixed-effects probit model with two fixed effects, (*i*) SAME-SENSE (reference level) vs. DIFFERENT-SENSE and (*ii*) COUNT (reference level) vs. MASS, just as we did for the Experiment 1 data. We see that the main effect of MASS ($\beta = 0.03, SE = 0.09, p = 0.73$) is non-significant, and so is the interaction MASS×DIFFERENT-SENSE ($\beta = 0.08, SE = 0.12, p = 0.50$). However, the main effect for DIFFERENT-SENSE ($\beta = -0.35, SE = 0.08, p = 2.6 \times 10^{-5}$) is significant, indicating (in conjunction with the fact that the other effects are non-significant) that there is a difference between matching and mismatching conditions for both mass and count containees.

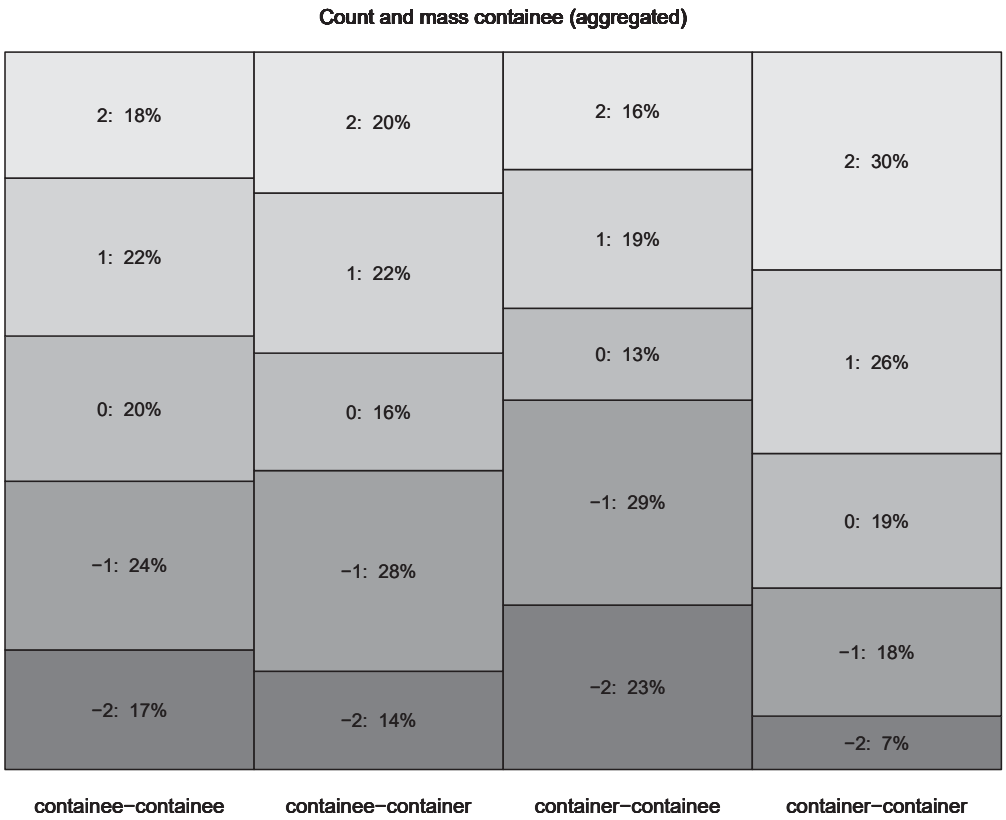

Figure 3: Experiment 2: Summary of the aggregated count and mass containee data.

We therefore aggregated over the count and mass data when we estimated the model examining the four experimental conditions for pseudo-partitives.[4] Just as before, we selected CONTAINER as the reference level for both the first predicate (the restrictive relative clause) and the second predicate

[4]The two models estimated for the count-containee subset only and for the mass-containee subset only yielded very similar results.

(the matrix VP). This time, the pseudo-partitives as a whole, not only the mass-containee ones, exhibited the same overall pattern as homonyms and polysemes, and fell roughly in the middle of the spectrum between them. There was a negative main effect of CONTAINEE for both the first sense ($\beta = -0.50, \mathit{SE} = 0.09, p = 5 \times 10^{-9}$) and the second sense ($\beta = -0.71, \mathit{SE} = 0.09, p = 0.00$), with a more pronounced main effect for the second sense, just as count containees exhibited in Experiment 1. Most importantly, there was a significant CONTAINEE×CONTAINEE interaction ($\beta = 0.64, \mathit{SE} = 0.12, p = 5.3 \times 10^{-8}$). All of these effects had the same direction as the corresponding homonymy and polysemy ones, and their magnitudes were intermediate between the corresponding homonymy and polysemy effects.

Importantly, the fact that there is a significant, negative main effect of CONTAINEE for the first sense, i.e., the fact that the container-container condition is significantly better than the containee-container condition, casts some doubt on the hypothesis that the restrictive relative clause (the first predicate) is free to attach itself to either the low (containee) noun or the high (container) noun in the pseudo-partitive. If such syntactic attachment freedom had been available, we would have expected to see no difference between the acceptability of the containee-container condition, where the relative clause would attach low, and the container-container condition, where the relative clause would attach high. Thus, it seems that relative clause attachment is fairly constrained in pseudo-partitives, eliminating one of the possible confounds associated with the fact that pseudo-partitives, unlike homonyms / polysemes, are syntactically complex.

When we compare the summaries in Figure 3 with the corresponding summaries in Figure 1, we see a slight difference in the acceptability profile associated with the four conditions. In particular, the difference between the container-containee and container-container conditions (conditions 3 and 4 in Figure 3) is as expected: the matching condition is significantly better than the mismatching condition (the *post hoc* comparison between these conditions is very highly significant: $\beta = -0.75, \mathit{SE} = 0.09, p = 0.00$). But there is no difference between the containee-containee and containee-container conditions (conditions 1 and 2 in Figure 3; the *post hoc* comparison is non-significant: $\beta = -0.05, \mathit{SE} = 0.08, p = 0.49$). We currently have no explanation for this difference in profile between pseudo-partitives and homonyms / polysemes.

4.3. Discussion

Broadly, the new items employed in Experiment 2 confirm the acceptability of copredication with container pseudo-partitives observed in Experiment 1. The rate of acceptance of copredication still stands somewhere in the middle of a continuum between homonymy and lexical polysemy. Speakers do not find that accessing multiple readings of pseudo-partitives leads to ungrammaticality, but they accept it less than when the same reading is accessed throughout the sentence. Crucially, the contrast between pseudo-partitives, homonyms and polysemes with respect to the acceptability of copredication does not seem to be categorical but a matter of degree.

5. General Discussion and Conclusion

The findings reported here, taken together, are quite intricate and complex, but they also point to a clear answer to our main empirical question, namely the acceptability of copredication with container pseudo-partitives. Throughout, pseudo-partitives failed to pattern exactly like homonymy. This is particularly relevant since it is one of the clearest theoretical predictions made by current analyses of container pseudo-partitives. Our results show that the container and containee (concrete portion) readings of the pseudo-partitive construction must be available simultaneously, albeit to a lesser extent than in the case of polysemes. These results lead to two desiderata.

The first is to provide a semantics for container pseudo-partitives that makes copredication licit, thus including it in the ranks of complex / dot-type expressions of the sort proposed by Asher (2011). The idea of constructing complex types outside of the lexicon is not entirely novel. Asher (2011) himself proposes that restricted predication of the sort illustrated in *Louise as a boss is strict* does precisely that. The *as*-phrase constructs the *boss* aspect of the expression in subject position and makes that aspect available for predication by *strict*. The complement of *as* will be a variable that will serve as argument for the main predication, and it will be of a type that can be constructed as an aspect of the term introduced by the subject. In other words, the subject's term is coerced into a complex type, such that one of its constituents is the type specified by the complement of *as*.

Similarly, the pseudo-partitive can be viewed as a complex type constructor, whose constituents are the types of the nominal phrases contained in it. A phrase like *bucket of water*, for instance, will contribute a variable with a complex type, roughly BUCKET • WATER, whose inhabitants are objects that have BUCKET and WATER as aspects associated through a containment relation. To make things concrete, assume a syntactic structure in which the pseudo-partitive contains a functional projection FP taking the container and containee NPs as arguments. The type constructor, contributed by the functional head F (maybe overtly realized as *of*), takes the container and containee NPs as arguments and returns a property of a complex type whose constituent types are contributed by the individual NPs. That is, the functional head F would contribute a function that extracts the most specific typing requirements associated with the properties contributed by the two NPs, and uses them to construct the complex type associated with the entire pseudo-partitive.

Overall, our experimental results indicate that it is worth exploring complex dot-type semantic representations beyond the lexicon. But we should note that the complex type approach (once properly fleshed out) predicts only a subset of our findings, namely that pseudo-partitives do not behave like homonyms. However, this approach in its simplest form cannot capture the fact that pseudo-partitives do not exhibit the exact same behavior as lexical polysemes. The proposal that pseudo-partitives and lexical polysemes make the same sort of complex, dot-type object available for copredication (and semantic composition more generally) fails to predict that in the pseudo-partitive case, the container aspect maintains some primacy. An immediate question is whether the asymmetrical availability of senses / aspects we observe with pseudo-partitives is specific to phrasal polysemy, or can be be observed in the lexical domain as well.

Another issue is identifying the conditions under which complex types can be compositionally constructed. In particular, our results naturally lead to the question of whether other binominal structures in English are associated with similar complex-type constructors, and whether they exhibit the same type of (possibly asymmetric) meaning multiplicity. For example, the position occupied by container words can also be filled by group nouns such as *committee (of administrators)* or *gang (of thieves)*, portion nouns such as *pile (of garbage/clothes)* and *bunch (of crap/roses)*, quantity / measure nouns such as *pound (of rice/beans)* and *ton (of gravel/beach pebbles)*, and classifying nouns such as *kind / sort / species (of cheese/squirrels)*.

It is known that pseudo-partitives do not exhibit meaning multiplicity only when the first noun is a container. For example, Brasoveanu (2008) and Rett (2014) argue that measure phrases in general are ambiguous: they can denote abstract measures or concrete individuals. The multifaceted nature of group nouns, which seem to allow reference to both an independent entity formed out of the group members and to the members themselves, has also received a great deal of attention in the literature (Barker 1992, Schwarzschild 1996, Winter 2002, Pearson 2011). The question emerging out of the present investigation is whether these senses can be made available simultaneously and if they can, whether there is a preference for one or another.

In sum, this paper has argued that the inventory of nominal meaning multiplicity in natural language is more diverse than that recognized by the current, relatively coarse-grained distinctions. We offer two main contributions. The first, more specific one is that container pseudo-partitives seem to demand a unique, complex representation in order to accommodate the acceptability of copredication. The second, broader contribution is providing experimental evidence in support of the idea that not all polysemy is created equal, and that there is a rich spectrum of polysemy to explore that includes both lexical and phrasal polysemy, and quite possibly distinct sub-varieties in each of these two major types of nominal polysemy.

References

Asher, N. (2011). *Lexical semantics in context: a web of words*. Oxford University Press.

Asher, N. and J. Pustejovsksy (2013). A type composition logic for generative lexicon. In J. Pustejovsky, P. Bouillon, H. Isahara, K. Kanzaki, and C. Lee (Eds.), *Advances in generative lexicon theory*, pp. 1–38. Springer Netherlands.

Barker, C. (1992). Group terms in English: representing groups as atoms. *Journal of Semantics 9.1*, 69–93.

Brasoveanu, A. (2008). Measure noun polysemy and monotonicity: Evidence from romanian pseudopartitives. *Proceedings of the 38th meeting of the North East Linguistic Society*, 139–150.

Christensen, R. H. B. (2012). ordinal—regression models for ordinal data. R package version 2012.09-11 http://www.cran.r-project.org/package=ordinal/.

Corver, N. (1998). Predicate movement in pseudopartitive constructions. In A. Alexiadou and C. Wilder (Eds.), *Possessors, Predicates, and Movement in the Determiner Phrase*, pp. 215–257. John Benjamins, Amsterdam.

Frazier, L. and K. Rayner (1989). Taking on semantic commitments. *Journal of Memory and Language 29*, 181–200.

Grimshaw, J. (2007). Boxes and piles and what's in them: Two extended projections or one. In *Architectures, Rules, and Preferences: Variations on Themes by Joan Bresnan*, pp. 199–206. Center for the Study of Language and Information Publications.

Landman, F. (2004). *Indefinites and the type of sets*. John Wiley and Sons.

Partee, B. and V. Borschev (2012). Sortal, relational and functional interpretations of nouns and russian container constructions. *Journal of Semantics 42*, 445–486.

Pearson, H. (2011). A new semantics for group nouns. *Proceedings of West Coast Conference on Formal Linguistics 28*, 160–168.

Pustejovsky, J. (1995). *The Generative Lexicon*. MIT Press.

R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rett, J. (2014). The polysemy of measurement. *Lingua 143*, 242–266.

Rothstein, S. (2009). Individuating and measure readings of classifier constructions. *Brill's Annual of Afrosiatic Languages and Linguistics 1*, 106–145.

Schwarzschild, R. (1996). *Pluralities*. Kluwer, The Netherlands.

Selkirk, E. (1977). Some remarks on noun phrase structure. In P. Culicover, T. Wasow, and A. Akmajian (Eds.), *Formal Syntax*, pp. 285–316. Academic Press.

Wickham, H. (2009). *ggplot2: elegant graphics for data analysis*. New York: Springer.

Winter, Y. (2002). Atoms and sets: A characterization of semantic number. *Linguistic Inquiry 33.3*, 493–505.