Bare Plurals, Multiplicity, and Homogeneity*

Manuel Križ
Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, CNRS, PSL Research University

Abstract This paper presents a novel view on the multiplicity implication of existential bare plurals. Deviating from the conventional wisdom that competition with the singular indefinite is involved, we argue that the facts can be conceived of in terms of the general trivalence of plural predication, which has been described for definite plurals under the name of *homogeneity*. This approach elegantly accounts for the behaviour of bare plurals using only independently motivated conceptual resources and overcomes problems that have been raised for the competition-based accounts.

Keywords Bare Plurals, Multiplicity, Homogeneity, Trivalence, Implicatures

1 Introduction

There is a question in formal semantics of whether the denotation of plural nouns like *zebras* in an existential bare plural noun phrase encompasses only pluralities of zebras (*exclusive reading*) or also atomic individuals (*inclusive reading*). At its simplest, the puzzle is that while (1a) implies that Mary saw multiple zebras, (1b) entails that she didn't see a single such animal. Thus, the plural noun seems to have an exclusive reading in (1a), but an inclusive reading in (1b).

(1) a. Mary saw zebras. ⇝ Mary saw more than one zebra.
b. Mary didn't see zebras. ⇝ Mary didn't see one or more zebras.

A different way of phrasing the point is to say that in (1a) there is a multiplicity implication that disappears in (1b).

Prior engagements with the phenomenon have mostly resulted in analyses of the multiplicity component of the meaning of (1a) as some kind of quantity implicature (Sauerland 2003, Sauerland et al. 2005, Spector 2007, Zweig 2008, 2009, Ivlieva 2013). That is to say, the meaning of the plural noun is inclusive, and so the literal meaning of (1a) is that Mary saw one or more zebras. But there is an intuition that if she had seen only one zebra, the speaker would

* I thank Daniel Büring, Lucas Champollion, Emmanuel Chemla, Scott Grimm, Natalia Ivlieva, Mora Maldonado, Clemens Mayr, Orin Percus, Philippe Schlenker, Roger Schwarzschild, and Benjamin Spector for valuable comments and discussion. All remaining errors are, of course, the author's. The research leading to these results has received funding from the following sources: European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n.313610; ANR-14-CE30-0010 TriLogMean; ANR-10-IDEX-0001-02 PSL*; and ANR-10-LABX-0087 IEC.

have used the singular, so that the multiplicity implication arises as a quantity implicature; the difference between the variants of the implicature approach lying in how they implement this intuition. In downward-entailing contexts such as under negation, this quantity implicature naturally disappears, and so there is an inclusive reading in (1b). An exception to this is the account by Farkas & de Zwart (2010), who assume that both an inclusive and an exclusive meaning are available and give an optimality-theoretic treatment of the choice between them.

The purpose of this paper is to propose an alternative perspective which uses quite different, but independently available conceptual resources: A view of multiplicity as likened to the homogeneity property of plural predication. The core datum illustrating the latter phenomenon is the following: A predicate applied to a plurality (denoted by a definite plural) is not true of this plurality when it is true of some, but not all of its parts; but neither is its negation true in such a case. Rather, both sentences have a third truth value, often called *undefined* for convenience.[1]

(2) Mary read the books.

true *iff Mary read all the books*
false *iff Mary read none of the books*
undef. *otherwise*

(3) Mary didn't read the books.

true *iff Mary read none of the books*
false *iff Mary read all of the books*
undef. *otherwise*

We proposed that in the same way as (2) is undefined when Mary read only half the books, (1a) is undefined when Mary saw exactly one zebra. From this assumption, substantive predictions about the behaviour of bare plurals follow immediately based on what is known independently about homogeneity.

We start by a brief discussion of the existing competition-based theories and the problems they face in section 2. In section 3, we show how the homogeneity property of plural predication naturally leads to an account of the behaviour of bare plurals in terms of trivalence. Section 4 discusses how both approaches deal with the context dependence of bare plurals, substantiated by recent experiments that have not had much reception in the theoretical literature so far. Section 5 finally sees our theory applied to the details of the behaviour of bare plurals by way of comparison with previous approaches.

2 The Competition-Based View

A number of analyses of existential bare plurals have been proposed, all of which build on the same underlying idea: the multiplicity implication of the bare plural arises because the existential bare plural is in competition with a singular

1 Cf. Schwarzschild 1994, Löbner 2000, Gajewski 2005, Križ 2015, *pace* Breheny (2005) and Magri (2014).

indefinite. The intuition behind it is quite simple and, in essence, Gricean: if Mary had seen only one zebra, the speaker would have said (4a). Thus, given that they said (4b), Mary must have seen several.[2]

(4) a. Mary saw a zebra.
 b. Mary saw zebras.

Several different implementations of this basic intuition have been given, all of which require some special-purpose assumptions. Zweig (2008) assumes local exhaustification of the verb phrase at the level of the event predicate; Mayr 2015 assumes local exhaustification at the level of the noun phrase; Sauerland (2003) and Sauerland et al. (2005) need to stipulate local computation of implicated presuppositions; Spector (2007) postulates recursive exhaustification and an intransitive relation of alternativehood; and Farkas & de Zwart (2010) give an optimality-theoretic meaning choice procedure specifically for the interpretation of number morphology. In the following, we will discuss puzzles for these kinds of approaches at a high level of abstraction. For the specifics of implementation, which are immaterial to our arguments, we refer the interested reader to the original publications.

2.1 Not a Matter of Meaning Choice

On most existing accounts, an instance of a bare plural is, in any given sentence, interpreted locally as either inclusive or exclusive. This is the result of the theories in Zweig 2008, 2009, Sauerland 2003, Sauerland et al. 2005, and Farkas & de Zwart 2010, despite the fact that the technical workings of those theories are quite different.

As long as the bare plural occurs in a logically monotonic context, this is unobjectionable. However, when a bare plural is embedded in a non-monotonic context, such as the scope of *exactly*, a problem arises. It is predicted, for example, that the meaning of (5) should be either (5a) or (5b). Which one it is depends on whether the competition-based rule is sensitive to strengthening or only to non-weakening; but it has to be one of them.

(5) Exactly one girl saw zebras.
 a. Exactly one girl saw one or more zebra.
 b. Exactly one girl saw two or more zebras.

However, as pointed out by Spector (2007) and Ivlieva (2014), and acknowledged as a problem for their own theory by Farkas & de Zwart (2010), one actually obtains an overall meaning that is stronger than either (5a) or (5b):

(6) Exactly one girl saw zebras. ⇝ One girl saw multiple zebras and the other girls saw none.

2 Here and later on, we use *several* as a sloppy but convenient shorthand for *more than one*, even though we recognise that in ordinary usage the word is usually reserved for slightly larger quantifies and two is, at best, a borderline case of *several*.

It appears that the meaning of (5) is, in fact, the *conjunction* of (5a) and (5b). Thus, in a way, the sentence makes use of both the inclusive and the exclusive meaning of the bare plural: the exclusive meaning features in the upward-monotonic part of the meaning of *exactly one* ("at least one"), and the inclusive meaning features in the downward-monotonic part of the meaning ("not more than one"). Theories that determine the meaning of the bare plural locally cannot account for this fact.

2.2 Recursive Exhaustification: Spector 2007

Spector's (2007) theory derives the multiplicity implication as a global *higher-order* implicature. We present the theory semi-formally here and refer the reader to the original paper for a more rigorous presentation. To see the basic principle, consider this. A statement with a singular indefinite, such as (7a), is taken to generate an implicature based on the alternatives with plural numerals in (7b).

(7) a. Mary saw a zebra. ⇝ Mary saw only one zebra.
b. Mary saw two/three/... zebras.

In virtue of its literal meaning, the sentence (8) has the same truth conditions as (7a): it is true if Mary saw one or more zebras. However, it has (7a) as an alternative, and an utterance of (7a) convey, in virtue of its own implicature, that Mary saw exactly one zebra. Hence, the fact that the speaker didn't choose the singular alternative, but instead uttered (8), gives rise to an implicature that the speaker didn't see exactly one zebra. That is to say, implicature computation takes into account not the literal meaning, but the communicated meaning of the alternative: we compute higher-order implicatures.

(8) Mary saw zebras.

Despite the fact that *Mary saw a zebra* and *Mary saw zebras* have the same literal meaning, they have different implicatures because they have different alternatives. This is the crucial stipulation in Spector's theory: alternativehood is intransitive. While *two zebras* is an alternative of *a zebra*, and *a zebra* is an alternative of *zebras*, the numeral-headed *two zebras* is not an alternative of the bare plural *zebras*.

(9) a. $\text{Alt}(\textit{zebras}) = \{\textit{a zebra}\}$
b. $\text{Alt}(\textit{a zebra}) = \{\textit{two zebras, three zebras, } \ldots\}$

The original paper contains a detailed presentation of the computation of these higher-order implicatures in various sentences, which we will not repeat here.

This theory, uniquely among the competition-based theories, derives the correct predictions for non-monotonic contexts if it is assumed that exhaustification consists in negating non-weaker alternatives (as opposed to stronger alternatives). Consider (10a) and its alternatives in (10b).[3]

[3] Note that sentences with numerals must not be understood under an *exactly*-reading here. Thus, *Exactly one girl saw two zebras* means that that exactly one girl saw two or more zebras, not that exactly one girl saw exactly two zebras.

(10) a. Exactly one girl saw a zebra.
b. Exactly one girl saw two/three/dots zebras.

These alternatives are not stronger than the original (10a), but they are also no weaker, so they are negated. What we obtain as the strengthened meaning of (10a) is (11).

(11) EXH (exactly 1... a) =
(exactly 1... a) and not (exactly 1... 2) and not (exactly 1 ... 3) ...
'Exactly one girl saw a zebra and she saw only one.'

Now the sentence with the bare plural, is exhaustified with respect to the alternative (11); that is to say, it has as its implicature the negation of (11).

(12) EXH (EXH (exactly 1... pl)) =
(exactly 1... a) and not EXH (exactly 1... a) =
(exactly 1... a) and ((exactly 1... 2) or (exactly 1... 3) or ...)
'Exactly one girl saw one or more zebras and exactly one girl saw two or more zebras.' = 'Exactly one girl saw more than one zebra and the others saw none.'

The negation of (12) — the strengthened meaning of the singular alternative — is then the implicature of the sentence (13a) with the bare plural. This implicature is given in (13b). Taken together with the literal meaning of the original assertion (12a), this yields the overall meaning (13c): that one girl saw multiple zebras and all the others saw none.

(13) a. Exactly one girl saw zebras.
'Exactly one girl saw one or more zebras.'
b. Either (not (exactly 1... a)) or (exactly 1... 2) or (exactly 1... 3)...
c. Exactly one girl saw one or more zebras and exactly one girl saw two or more zebras.

As we will see later in section 5, the predictions of this theory for complex sentences are extremely similar to those of our own, with only subtle differences to be discussed in sections 5.1 and 5.3.

2.3 A Lack of Alternatives

Magri (2011) points out that a multiplicity seems to be present with a least some mass nouns, such as *change*:

(14) a. John has change in his pockets. $\rightsquigarrow$ He has several pieces of change.
b. John doesn't have change in his pockets. $\rightsquigarrow$ He doesn't have even a single piece of change.

He presents this as a challenge for competition-based theories of the multiplicity implication, since it is not clear what alternative *change* is supposed to compete with. An approach that builds on competition between singular and plural *morphology* (Sauerland 2003; Sauerland et al. 2005; Mayr 2015; Farkas & de Zwart 2010) is particularly out of luck, since *change* is already morphologically singular and has no other form. But even when competition between whole noun phrases is considered, there is a problem. The semantically singular alternative to *change* would seem to be *a piece of change*. This noun phrase, however, is markedly more complex, and so it is doubtful whether it can serve as an alternative in the sense required for implicature computation (Katzir 2007). In the particular case of *change*, one might attempt to work off the lexical replacement with *a coin*; but in a country with 1$-banknotes, are we certain that all change is coins?

One may note, though, that the behaviour described by Magri does not appear to be even remotely universal among mass nouns. Certainly my *luggage* can be a single suitcase, today's *mail* can be a single letter, and a cardinal's *jewellery* may consist in a single ring. It is thus doubtful to what extent the odd behaviour of *change* in particular should be viewed as a general problem.[4]

A stronger version of the same argument has been raised by Ivlieva & Sudo (2015): the multiplicity inference is also, and more consistently, present with *pluralia tantum* of a particular kind, such as *clothes, belongings, possessions, goods, movables, valuables*, and *eatables*, which refer to collections of well-individuated objects.

(15) a. This bag has clothes in it. ⇝ It contains multiple articles of clothing.
b. This bad doesn't have clothes in it. ⇝ It doesn't contain a single article of clothing.

In some cases, such as with *clothes*, there may be lexical replacement (*a garment*) or circumlocutions available (*an article of clothing*) to generate a potential singular alternative. For most of these nouns, however, not even those are easily found. In addition, it is simply not possible to use *pluralia tantum* of this kind to refer to a single object: while one can point at a pair of scissors (a different type of plurale tantum) and utter (16a), one cannot hold up a single shirt and say (16b).[5]

(16) a. *Pointing at a pair of scissors.*
These are scissors.
b. *Pointing at a single shirt.*
#These are clothes./ #This is clothes.

We thus appear to have genuinely semantically plural nouns with no singular alternative. This is a particular problem for Spector's theory, which requires competition between whole noun phrases. The approaches that rely on local

4 Benjamin Spector (p.c.) suggests that *change* is a vague term and that in many contexts at least, a single coin would be a borderline case of change (and, consequently, also a borderline case of non-change, hence the behaviour of negation).

5 On this diagnostic, singular mass nouns appear to mostly pattern with *scissors*. It is quite possible, for example, to point at a single table and utter (i). (We thank an anonymous reviewer for pointing this out.)

(i) *Pointing at a table.*
This is furniture.

implicatures or implicated presuppositions based on the meaning of number morphology,[6] however, may have a way out. These accounts postulate number operators as nodes in the logical form (above the determiner for Sauerland et al., above the noun but below the determiner for Mayr), whose presence is in some way regulated by number morphology. Thus, the plural *clothes* has a logical form like (17).

(17) [*cloth-* PL]

One could assume now that the alternatives at play are actually alternative logical forms, obtained by replacing the PL node with a SG node. This is perfectly possible on the logical form in (17) and gives us (18). It just so happens that there is no pronounceable phonological form associated with this logical form, but this does not have to mean that the logical form is in and of itself ill-formed.

(18) [*cloth-* SG]

Thus, the alternative needed for the computation of the plurality inference on these accounts would, in fact, exist, despite the fact that it cannot be pronounced.

2.4 Interim Summary

In this section, we discussed challenges for the existing accounts of the multiplicity implication in terms of competition between the bare plural and the singular indefinite. These challenges are of three kinds.

First, all of these accounts need to make special-purpose assumptions to account for bare plurals. If a theory could be found that does not require such assumptions, it would therefore, *ceteris paribus*, have an advantage in plausibility.

Second, with the sole exception of Spector's (2007), existing accounts make incorrect predictions for sentences where a bare plural is embedded in the scope of a non-monotonic quantifier such as *exactly one*. Spector's (2007)'s theory, in turn, faces an objection pointed out by Ivlieva & Sudo (2015): it cannot account for the existence of a multiplicity implication with *pluralia tantum*, which do not have a singular alternative to compete with the plural.

In the next section, we will introduce our own account, which rests on a fundamentally different conceptual basis that allows it to overcome these limitations.

3 Homogeneity: From Definite to Bare Plurals

A well-known type of sentences where the truth conditions of an affirmative and its negation are not complementary are sentences with definite plural noun phrases.

(19) a. Mary bought the books. ⇝ She bought all of the books.

6 Sauerland 2003; Sauerland et al. 2005; Mayr 2015. It is possible to rephrase Zweig 2009 in such terms as well.

b. Mary didn't buy the books. ⇝ She bought none of the books.

This phenomenon, known as the *homogeneity effect*, has generally received an analysis in terms of logical trivalence. This means that the sentences in (19) are sometimes true, sometimes false, and sometimes they have a third truth value. We will call sentences that have this truth value *undefined*.[7]

(20) Mary bought the books.

true *iff Mary bought all the books*
false *iff Mary bought none of the books*
undef. *otherwise*

Crucially, the falsity conditions of a sentence are identified with the truth conditions of its negation. In other words, negation does what is usual for it to do in a trivalent logic: it switches truth and falsity, but leaves undefinedness alone.

(21) Mary didn't buy the books.

true *iff Mary bought none of the books*
false *iff Mary bought all of the books*
undef. *otherwise*

Thus one captures the gap between the truth conditions of (19a) and those of (19b): in those situations where Mary bought some, but not all of the books, both sentences are undefined. Note that at this point, the third truth value — undefinedness — has been introduced only as a theoretical concept that serves the purpose of explaining the truth conditions of negated sentences. How this third truth value translates to intuitive judgements of sentences is, in principle, a separate question.[8]

It is now easy to see the picture that emerges when one looks at sentences with bare plurals through the lens of trivalence: such sentences are undefined in a situation where there is only a single witness.

(22) Mary saw zebras.

true *iff Mary saw multiple zebras*
false *iff Mary saw no zebra*
undef. *iff Mary saw exactly one zebra*

7 Schwarzschild 1994, Löbner 2000, Gajewski 2005, Križ 2016. *Pace* Magri 2014, who propose approaches to homogeneity that is explicitly not based on logical trivalence.

8 Experimental findings indicate that, even if the third truth value may not always be introspectively accessible, undefined sentences have a particular signature that sets them apart from true as well as false sentences. Schwarz (2013) found that in a truth-value judgement task with the two answer options *true* and *false*, undefined affirmative sentences received varying responses. (Negated sentences did not feature in the experiment.) Križ & Chemla (2015) presented both affirmative and negative sentences, with the three answer options *completely true*, *completely false*, and *neither*. For undefined sentences, answers were predominantly split between *completely false* and *neither*, with only a small proportion of *completely true* responses, whereas sentences that were either true or false received very consistent responses. We can thus say at the very least that the theoretical status of undefinedness is associated with some variability in judgements. For some further discussion of the role of context in steering this variability, see section 4 below.

(23) Mary didn't see zebras.

true *iff Mary saw no zebra*
false *iff Mary saw multiple zebras*
undef. *iff Mary saw exactly one zebra*

In the following sections, we will demonstrate how the trivalence of sentences with bare plurals can be derived from the very same principles that have been posited independently on the basis of sentences with definite plurals.

3.1 The Homogeneity Constraint

The trivalence of sentences with bare plurals is generally seen as rooted in the predicate: most predicates in natural language are, when applied to a plurality,[9] undefined under certain conditions.[10] Much of the literature has restricted its attention to distributive predicates. In this case, the generalisation is quite easy to see: a predicate is undefined of a plurality if it is true of some parts of it and false of others. Križ (2016) extends the picture to collective predicates, which he argues are also sometimes undefined, and postulates a constraint of which that for distributive predicates is a special case. This constraint is taken to apply to lexical predicates in natural language.[11]

(24) **Generalised Homogeneity**
A homogeneous predicate P is undefined of a plurality a if it is not true but there is a plurality b that overlaps with a (i. e. has constituent individuals in common) such that P is true of b. (Križ 2016)

Križ assumes that the positive extension of a predicate — the set of entities, including pluralities, that they are true of — is given as a matter of lexical meaning, and that predicates are false whenever the homogeneity constraint allows them to be false; otherwise, they are undefined. Using these principles, we can thus derive the falsity conditions of a predicate from its truth conditions and do not have to postulate them separately.

The homogeneity constraint, as applied to verbal predicates, explains the behaviour of sentences with definite plurals in argument position. Take, for example, the predicate *smiled*. In virtue of its lexical meaning, this predicate is true of a plurality if all of its members smiled; that is just what it means for a plurality to smile. Now we look at all pluralities x which are not such that all of their members smiled. If x contains any part y that smiled, then homogeneity forbids the predicate from being false of x (since the containment of y in x is a special case of overlap). If x contains no part that smiled, then homogeneity

9 The underlying view of plural predication is the usual from Link 1983.

10 Exceptions are Breheny 2005, Büring & Križ 2013, and Magri 2014, who locate the source of the homogeneity effect in the meaning of definite plural noun phrases. See Križ 2015 and Križ & Spector 2017 for arguments that this is mistaken.

11 Exceptions to this generalisation exist, but appear to be limited to a class of predicates that, intuitively speaking, have something to do with measuring a quantity, such as *numerous* and *few in number*.

permits the predicate *smiled* to be false of x, and so it is. Therefore, (25), where *smiled* is predicated of the plurality consisting of the professors, is true if and only if all professors smiled, and false if and only if none of them did. Otherwise, it is undefined.

(25) The professors smiled.
true *iff all professors smiled*
false *iff none of the professors smiled*
undef. *otherwise, i.e. if some but not all of them smiled*

In order to understand our original example *Mary bought the books*, where the other argument position of the transitive verb is filled by an atomic individual, it is, for current purposes, sufficient to note that the one-place predicate *Mary bought* behaves like a lexical predicate and obeys the generalised homogeneity constraint. This is so in general for one-place predicates that are obtained by filling the other arguments of a relation with individuals (cf. Križ 2015: §1 and §2 for a more detailed discussion).

We submit that the key to explaining the trivalence of sentences with bare plurals is to recognise that the homogeneity constraint also applies to nominal predicates, and in particular plural nouns. Assume that the plural noun *zebras* is true of pluralities of zebras, but not true of atomic zebras. By the homogeneity constraint, it then cannot be false of atomic zebras, and so it has to be undefined of them (along with pluralities that consist only partly of zebras).

(26) $$[\![\text{zebras}]\!] = \lambda x. \begin{cases} 1 & \text{iff } x \text{ is a plurality of zebras} \\ 0 & \text{iff } x \text{ doesn't contain any zebras} \\ \star & \text{otherwise, e.g. if } x \text{ is an atomic zebra} \end{cases}$$

This alone, of course, does not yet tell us what the truth and falsity conditions of an entire sentence with an existential bare plural are. It is standardly assumed that sentences with an existential bare plural involve a silent existential quantifier that is applied to the plural noun. Thus, (27a) has the logical structure in (27b).

(27) a. Mary saw zebras.
b. $\mathbf{exists}(zebras, \lambda x.saw(m, x))$

We therefore need to understand what happens when a generalised quantifier, such as this silent existential quantifier, is applied to a trivalent predicate. This is the topic of the next sections.

3.2 Homogeneity and Quantification

We will explore the interaction of quantification with trivalent predicates in three steps. First, we will discuss the behaviour of one-place quantifiers that range over atoms when they are applied to a trivalent scope predicate. Then we will make the step to unary quantifiers over pluralities. Finally, we will generalise our approach to binary quantifiers, that is, determiners.

3.2.1 Quantification over Atoms

Let us begin in a place far away from existential bare plurals: by looking at quantifiers that quantify over atoms when they are applied to a trivalent scope predicate, such as a verb phrase that contains a definite plural.

(28) $\left.\begin{array}{l}\text{Every student}\\ \text{Exactly one student}\\ \text{No student}\end{array}\right\}$ bought the books.[12]

Here, the quantifiers are being applied to the trivalent predicate *bought the books*. Since the quantification is only over atomic individuals, we can restrict our attention to the truth value that this predicate assumes for atomic arguments:

(29) $[\![\text{bought the books}]\!] = \lambda x. \begin{cases} 1 & \text{iff } x \text{ bought all of the books} \\ 0 & \text{iff } x \text{ didn't buy any of the books} \\ \star & \text{iff } x \text{ bought only some of the books} \end{cases}$

This type of case was experimentally investigated by Križ & Chemla (2015). In a truth-value judgement task with the three answer options *completely true*, *completely false*, and *neither*, the presence of a significant proportion of *neither*-responses could be identified as a diagnostic for undefinedness. This was used to detect a consistent pattern across the various quantifiers tested. This pattern can be described in terms of the trivalent predicate and the familiar *bivalent* meaning of the quantifier; we thus do not have to postulate separately described falsity conditions for each quantifier. The application of a unary quantifier $\mathcal{Q}$ (such as *every student*) to a trivalent predicate P is captured by the following rule:

(30)
1. Let P^1 be that predicate which is just like P except that it is true of all the individuals where P is undefined, and P^0 that predicate which is just like P except that it is false of all the individuals where P is undefined.[13]
2. If $\mathcal{Q}(P^1)$ and $\mathcal{Q}(P^0)$ have the same truth value, then $\mathcal{Q}(P)$ has that truth value.
3. Otherwise, $\mathcal{Q}(P)$ is undefined.

Intuitively, we look at the two ways of resolving the undefined cases — resolving them either to truth or to falsity — and then ask whether the quantifier is true of both resulting predicates or false of them. Let us see how this applies to (31).

(31) Every student bought the books.

12 In this example and the ones to follow, we sometimes equivocate between concrete, physical books, which can only be bought by one person, and books as abstract entities, where to buy them means either to buy a material instance of them or to buy the rights to publish them. These difference are entirely inconsequential for our argument, since pluralities of abstract books behave the same as pluralities of physical books for the purposes of homogeneity.

13 The bivalent P^1 and P^0 will, of course, in general violate homogeneity.

The first step is to find the two ways of resolving the undefined cases, that is, to form $[\![\text{bought the books}]\!]^1$ and $[\![\text{bought the books}]\!]^0$. These are equivalent to the meanings of *bought at least some of the books* and *bought all of the books*, respectively.

(32) a. $[\![\text{bought the books}]\!]^1 = \lambda x. \begin{cases} 1 & \text{iff } x \text{ bought at least some of the books} \\ 0 & \text{iff } x \text{ didn't buy any of the books} \end{cases}$
'bought at least some of the books'

b. $[\![\text{bought the books}]\!]^0 = \lambda x. \begin{cases} 1 & \text{iff } x \text{ bought all of the books} \\ 0 & \text{iff } x \text{ didn't buy all of the books} \end{cases}$
'bought all of the books'

Now we apply the well-understood bivalent quantifier *every student* to both of these predicates in turn:

(33) a. Every student bought at least some of the books.
b. Every student bought all of the books.

The original trivalent sentence (31) is now true if and only if both of (33a) and (33b) are true, and false if and only if both of them are false. It is undefined when one is true and the other is false, such as when some of the students read all of the books and the other students read only some—in which case (33a) is true and (33b) is false.

(34) Every student bought the books.
true *iff every student bought all the books*
false *iff at least one student bought none of the books*
undef. *otherwise*

3.2.2 Quantification over Pluralities

Plural quantifiers in natural language are compatible with (at least some) collective predicates, and for that reason we want to be able to quantify over pluralities. (35), for example, on its most natural reading, does not entail that there are two girls such that each of them lifted the piano. Rather, it needs to be conceptually analysed along the lines of (35b): as asserting the existence of a plurality with certain properties.

(35) a. Two girls lifted the piano.
b. 'There is a duality (a plurality with two members) of girls such that that duality lifted the piano.'

Thus, we have to analyse *two girls* as a quantifier quantifier over pluralities that is true of a predicate just in case the predicate's extension contains a duality of girls (as opposed to containing two atomic girls).

We would furthermore like to assume that there is quantification over pluralities even when a plural quantifier is applied to a distributive predicate, since otherwise we would have to assume that plural quantifiers are systematically ambiguous. In other words, we want (36a) to be analysed, informally, as (36b).

(36) a. Mary saw two zebras.
b. 'There is a duality of zebras such that Mary saw that duality.'

In this sentence, we are thus applying the quantifier *two zebras*, which is, by assumption, true of all predicates whose extension contains a duality of zebras, to a trivalent predicate of pluralities:[14]

(37) $$[\![\text{Mary saw}]\!] = \lambda x. \begin{cases} 1 & \text{iff Mary saw all of } x \\ 0 & \text{iff Mary saw no part of } x \\ \star & \text{iff Mary saw only a part of } x \end{cases}$$

If we employ without modification the algorithm (30), then we first have to form $[\![\text{Mary saw}]\!]^1$ and $[\![\text{Mary saw}]\!]^0$, the result of which is easy to grasp intuitively:

(38) a. $[\![\text{Mary saw}]\!]^1(x)$ is true iff Mary saw all of x.
b. $[\![\text{Mary saw}]\!]^0(x)$ is true iff Mary saw at least a part of x.

Thus, according to the rule, (36) is true if both (39a) and (39b) are true, and false if both are false.

(39) a. There is a duality of zebras such that Mary saw both of them.
b. There is a duality of zebras such that Mary saw at least one of them.

Now assume that Mary saw exactly one zebra z_1. Then (39a) is false, but as long as there is at least one other zebra z_2 in the universe, (39b) is true, because there is the duality $z_1 \oplus z_2$, of which Mary saw a part. More generally, assuming there are at least two zebras in the universe, we obtain the following result:

(40) Mary saw two zebras.

true *iff Mary saw at least two zebras*
false *iff Mary saw no zebras*
undef. *iff Mary saw exactly one zebra*

This is obviously undesirable, since it is intuitively quite clear that (40) is simply false whenever Mary saw only one zebra, and its negation is true.

(41) Situation: *Mary saw exactly one zebra.*
a. #Mary saw two zebras. (false)
b. ✓Mary didn't see two zebras. (true)

The following modification of the algorithm in (30) solves this problem:

(42) 1. Let P^* be that predicate which is just like P except that if P is undefined of x and there is an individual y which overlaps with x and of which P is false, then P^* is false of x.
2. Apply rule (30) to $\mathcal{Q}$ and P^*. The result is the truth value of $\mathcal{Q}$ for P.

14 Again, we are setting aside the upper-bounded reading of bare numerals, on which *two* means, effectively, *exactly two*. Our eventual system is set up in such a way as to be defined for this case as well, but it is immaterial to the point we are making here.

We have added here an additional step at the beginning: the predicate is first changed so that it is false of a plurality as soon as it is false of any part of the plurality. Only then are the two possible ways of "resolving" (the remaining) undefined cases determined.

For example, we turn ⟦Mary saw⟧ in (37) into $⟦\text{Mary saw}⟧^*$ in (43). Since ⟦Mary saw⟧ is a distributive predicate, it is defined of all atoms and undefined only of pluralities which contain of a mixture of individuals that Mary saw and individuals that Mary didn't see. It is precisely those individuals for which the * function turns the predicate false, resulting in a predicate that is bivalent.

(43) $$[\![\text{Mary saw}]\!]^* = \lambda x. \begin{cases} 1 & \text{iff Mary saw all of } x \\ 0 & \text{iff Mary didn't see all of } x \\ \star & \text{otherwise, i.e. never} \end{cases}$$

Since $[\![\text{Mary saw}]\!]^*$ is already bivalent, it is identical to $[\![\text{Mary saw}]\!]^{*0}$ and $[\![\text{Mary saw}]\!]^{*1}$. The overall truth value of the sentence we are interested in is therefore the truth value of *two zebras* as applied to $[\![\text{Mary saw}]\!]^*$, which is truth. Thus, (36) is correctly predicted to be bivalent, and equivalent to (39a):

(44) Mary saw two zebras.
true *iff Mary saw at least two zebras*
false *iff Mary saw at most one zebra*
undef. *never*

We now also capture the behaviour of *all* as a homogeneity remover that was pointed out by Löbner (2000). What me mean by this is the following: if a sentence with a definite plural is undefined because the predicate is true of some individuals in the plurality denoted by the definite plural and false of others, then the corresponding sentence with *all* added to the definite plural is just false. Compare, for example, (45) and (46).

(45) Mary bought the books.
true *iff Mary bought all of the books*
false *iff Mary bought none of the books*
undef. *otherwise, i.e. iff she bought some but not all of the books*

(46) Mary bought all of the books.
true *iff Mary bought all of the books*
false *iff Mary didn't buy all of the books*
undef. *never*

This falls out of our rules if we assume the following bivalent meaning for *all the books*:[15]

(47) $[\![\text{all the books}]\!](P)$ is true iff P is true of the maximal plurality of books.

15 The reason for assuming this meaning, rather than a universal quantifier over atoms, is, of course, the fact that *all* is compatible with collective predicates, whereas *every*, being a true universal quantifier over atoms, is not.

In order to apply this quantifier to $[\![\text{Mary read}]\!]$, we form $[\![\text{Mary read}]\!]^*$, analogously to (43):

(48) $$[\![\text{Mary bought}]\!]^* = \lambda x. \begin{cases} 1 & \text{iff Mary bought all of } x \\ 0 & \text{iff Mary didn't buy all of } x \\ \star & \text{otherwise, i.e. never} \end{cases}$$

Since this predicate is already bivalent, it is identical to $[\![\text{Mary bought}]\!]^{*0}$ and $[\![\text{Mary bought}]\!]^{*1}$. We therefore simply apply our quantifier to it and obtain truth or falsity. In this case, this amounts to just taking the truth value that (48) assigns to the plurality of all books: truth if Mary bought all of them, falsity if she didn't. What is responsible for the homogeneity-removing property of *all* is thus actually the move from an individual to a quantifier. Its nature as a quantifier triggers the application of our algorithm, which involved the application of the *-operator to the scope predicate. It is this operator that actually does the work of removing homogeneity in the requisite way.

Our improved algorithm still works if the predicate we are dealing with is undefined not only of pluralities, but also of some atoms. In particular, when quantification is only over atomic individuals, it reproduces the results of (30) from the previous section.[16] For an example, consider the sentence in (49), on its distributive reading.

(49) Two students bought the books.

With respect to atoms, *bought the books* (again, conceived of as a distributive predicate; we ignore the possibility of collective acquisitions resulting in shared ownership) is trivalent in the way shown in (29). When it comes to pluralities, the predicate is true of a plurality only if it is true of all members of the plurality, since we have stipulated that it is to be understood distributively. By homogeneity, it is then false of the plurality only if it is false of all members of the plurality.

(50) $$[\![\text{bought the books}]\!] = \lambda x. \begin{cases} 1 & \text{iff all of } x \text{ bought all of the books} \\ 0 & \text{iff none of } x \text{ bought any of the books} \\ \star & \text{otherwise} \end{cases}$$

The first step of applying the algorithm in (42) is to find P^*, that is, in this case, $[\![\text{bought the books}]\!]^*$. This is obtained by switching the predicate from undefined to false for all pluralities where it is false of some part of that plurality.

(51) $$[\![\text{bought the books}]\!]^* = \lambda x. \begin{cases} 1 & \text{iff all of } x \text{ bought all of the books} \\ 0 & \text{iff at least one of } x \text{ didn't buy any of the books} \\ \star & \text{otherwise} \end{cases}$$

16 It is easy to see why the latter is the case. As we move from P to P^*, we switch the predicate from $\star$ to 0 for individuals that have proper parts for which P is false. Since atoms have no proper parts, this change can never apply, so as far as atomic arguments are concerned, P^* is necessarily identical to P.

This predicate predicate is undefined only of atoms that bought some but not all of the books, and of pluralities such that all of their members bought at least some of the books and at least one member bought only some of the books. Starting from this, we now look at the two ways of resolving the undefined cases:

(52) a. $[\![\text{bought the books}]\!]^{*1} = \lambda x. \begin{cases} 1 & \text{iff all of } x \text{ bought at least some of the books} \\ 0 & \text{iff at least one of } x \text{ didn't buy any of the books} \end{cases}$

b. $[\![\text{bought the books}]\!]^{*0} = \lambda x. \begin{cases} 1 & \text{iff all of } x \text{ bought all of the books} \\ 0 & \text{iff at least one of } x \text{ didn't buy all of the books} \end{cases}$

Our original sentence is then true if both (53a) and (53b) are true, and false if both are false.

(53) a. There is a duality x of students such that both of x bought all of the books.
b. There is a duality x of students such that both of x bought at least some of the books.

Putting everything together, we obtain the following overall truth and falsity conditions:

(54) Two students bought the books.
true *iff at least two students bought all of the books*
false *iff at most one student bought any of the books*
undef. *otherwise*

3.3 Determiners

The last step we need to take in order to be able to analyse sentences with existential bare plurals, and to understand the role that the trivalence of the plural noun plays, is to generalise the rule for applying quantifiers from unary quantifiers (such as *every student*) to binary quantifiers, that is, determiners (such as *every*, or the silent existential determiner that comes with bare plurals). This is easily done:

(55) A determiner $\mathcal{D}$ applied to a pair of restrictor and scope $\langle P, Q\rangle$ is true/false iff it is true/false for all sequences $\langle P^{*i}, Q^{*j}\rangle$ with $i, j \in \{0, 1\}$, and otherwise undefined.

Consider now a sentence with a bare plural with its logical structure:

(56) Mary saw zebras.
$\mathbf{exists}(zebra, \lambda x.saw(m, x))$

Here we are applying the quantifier **exists**, which is true of P and Q if there is an individual x such that $P(x) = Q(x) = 1$, to two trivalent predicates:

(57) a. $[\![\text{zebras}]\!] = \lambda x. \begin{cases} 1 & \text{iff } x \text{ is a plurality of zebras} \\ 0 & \text{iff } x \text{ doesn't contain any zebras} \\ \star & \text{otherwise, e.g. if } x \text{ is an atomic zebra} \end{cases}$

b. $[\![\text{Mary saw}]\!] = \lambda x. \begin{cases} 1 & \text{iff Mary saw all of } x \\ 0 & \text{iff Mary saw none of } x \\ \star & \text{otherwise, i.e. iff Mary saw only some of } x \end{cases}$

First we form the starred versions of these predicates. For *zebras*, it is obtained by making the predicate false of any plurality that contains at least one non-zebra. For *Mary saw*, it is obtained by making the predicate false of any plurality at least part of which Mary didn't see. Since *Mary saw* is distributive and always defined of atomic individuals, this incidentally renders the scope predicate bivalent.[17] This makes it apparent that it is really the trivalence of the restrictor predicate *zebras* that is at work here.

(58) a. $[\![\text{zebras}]\!]^* = \lambda x. \begin{cases} 1 & \text{iff } x \text{ is a plurality of zebras} \\ 0 & \text{iff } x \text{ contains at least one non-zebra} \\ \star & \text{iff } x \text{ is an atomic zebra} \end{cases}$

b. $[\![\text{Mary saw}]\!]^* = \lambda x. \begin{cases} 1 & \text{iff Mary saw all of } x \\ 0 & \text{iff Mary didn't see all of } x \end{cases}$

Now we look at the ways of resolving undefinedness. For $[\![\text{Mary saw}]\!]^*$, this changes nothing since that predicate is already bivalent. For $[\![\text{zebras}]\!]^*$, we obtain the following:

(59) a. $[\![\text{zebras}]\!]^{*1} = \lambda x. \begin{cases} 1 & \text{iff } x \text{ consists solely of one or more zebras} \\ 0 & \text{otherwise, i.e. iff } x \text{ contains at least one non-zebra} \end{cases}$
'one or more zebras'

b. $[\![\text{zebras}]\!]^{*0} = \lambda x. \begin{cases} 1 & \text{iff } x \text{ is a plurality of zebras} \\ 0 & \text{otherwise} \end{cases}$
'multiple zebras'

Hence, (56) is true iff both (60a) and (60b) are true, and false iff both are false. Otherwise, that is, when Mary saw exactly one zebra, it is undefined. This is exactly the result we want.

(60) a. Mary saw one or more zebras.
b. Mary saw multiple zebras.

This concludes our demonstration of how the trivalent meaning for sentences with existential bare plurals can be derived on the basis of general principles that are at work in plural predication.

We would like to point out that our approach is compatible with both a determiner analysis and an adjectival analysis of *two* and other numerals. On a

17 We are assuming here for simplicity that one has seen a zebra as soon as one has seen a part of it.

quantificational analysis, the underlying bivalent meaning of *two* has to be as in (61).

(61) $[\![\text{two}]\!](P, Q)$ is true iff there is a duality x such that $P(x) = Q(x) = 1$.

Since atomic zebras are, in any case, not dualities, the fact that the restrictor predicate *zebras* is undefined of them cannot have any influence on the overall meaning of a sentence like *Mary saw two zebras*: the sentence is false no matter how these undefined cases are resolved.

Alternatively, one might view *two* as a modifier of the noun that returns another predicate (Link 1987 and many since). In this case, the predicate that it returns has to be assumed to be a non-homogeneous predicate that is true of things that are dualities (of the required kind) and false of everything that isn't a duality.[18]

(62) $$[\![\text{two zebras}]\!] = \lambda x. \begin{cases} 1 & \text{iff } x \text{ is a duality of zebras} \\ 0 & \text{iff } x \text{ is not a duality or contains no zebras} \\ \star & \text{iff } x \text{ is a duality containing only one zebra} \end{cases}$$

This predicate violates homogeneity, but it is not lexical, and so homogeneity does not apply to it. One the other hand, one might want to assume that *two* itself is simply a predicate that is combined with *zebras* by intersection. In this case, *two* would seem to be a lexical predicate that violates homogeneity. There is, however, precedent for this: Križ (2015) points out that lexical predicates systematically violate homogeneity when they involve some kind of measure function. These include predicates such as *numerous*, and *two*, being concerned with measuring how many parts there are in a plurality, would fall into the same exempted class.

The predicate *two zebras*, on the adjectival view, then combines with the same silent existential quantifier that is also at play in existential bare plurals. Since, however, this restrictor predicate is already false of everything that is not a duality, including atomic zebras, the resulting quantification is straightforwardly false if only one zebra fulfils the scope predicate.

3.4 Excursus: Homogeneity and Collective Predicates

The homogeneity constraint, as stated by Križ (2016), applies to collective predicates as well as distributive predicates, and this is what is required for the case of bare plural nouns, which we have analysed as collective predicates: they are true of pluralities without being true of their constituent atoms. In this section, we would like to set aside for a moment our main line of argument, which is concerned with bare plurals, and take a look at collective *verbal* predicates and their interaction with quantifiers, which the very general rules we have set forth also cover.

18 (62) is obtained by conjoining via Strong Kleene conjunction the trivalent meaning of *zebras* with a predicate that is true of dualities and false of non-dualities. In the case of a duality that contains only one zebra, *zebras* is undefined and the duality-predicate is true, hence undefinedness is obtained for their conjunction.

Križ points out that a collective predicate is, in accordance with the generalised homogeneity constraint, false only of those individuals who did not even participate in any collective action described by the predicate. For example, *perform Hamlet* is false only of those who did not participate in a performance of the play.[19] We therefore end up with the following overall truth and falsity conditions for the sentence in (63):

(63) The boys performed *Hamlet*.

true *iff all the boys together performed the play*
false *iff no boy participated in a performance of the play*
undef. *otherwise*

This is most readily seen when considering the negated counterpart of the sentence, which, with neutral intonation, is most naturally understood as saying that the boys had no part in any performance.[20]

(64) The boys didn't perform *Hamlet*. $\rightsquigarrow$ No boy participated in any performance.

What this means is that there are three types of situations where the sentence is undefined:

(65) *The boys performed Hamlet* is undefined iff

a. the boys are a proper superset of the performers, i.e. the play was performed by only some of the boys; or
b. the boys are properly included in the set of performers, e.g. the play was performed by the boys together with the girls; or
c. there is overlap between the boys and the performers without inclusion, e.g. when the play was performed by only some of the boys together with some girls.

What is interesting for our purposes is how homogeneity with collective predicates interacts with quantification. We saw in section 3.2.2 above that with distributive predicates, *all* removes homogeneity with respect to the plurality that it is associated with. In the case of collective predicates, it does so only partially; that is to say, it renders the sentence false in only some of these three types of scenarios. In particular, the sentence with *all* is still undefined in scenario (65b): there appears to be no difference in status between (66a) and (66b), or between their negations in (67), in such a context.

(66) Situation: *The boys together with the girls performed Hamlet.*

a. #The boys performed Hamlet. (undef.)
b. #All the boys performed Hamlet. (undef.)

(67) Situation: *The boys together with the girls performed Hamlet.*

---

19 The collective predicate here must be distinguished from the distributive participatory predicate *perform **in** Hamlet*, which has the same falsity conditions, but is true also of atomic individuals and not only of pluralities of actors.

20 For more detailed empirical arguments, partly repeated from Križ 2015: §1, see appendix A.

a. #The boys didn't perform Hamlet. (undef.)
b. #Not all the boys performed Hamlet. (undef.)

The undefined status of these sentences comes into sharper view when we compare them to the *all*-sentences with the corresponding participatory predicate:

(68) Situation: *The boys together with the girls performed Hamlet.*
a. ✓All the boys performed in Hamlet. (true)
b. #Not all the boys performed in Hamlet. (false!)

Thus, the truth and falsity conditions for the *all*-sentence with the collective predicates can be summarised as follows:

(69) All the boys performed *Hamlet*.

true *iff the boys and only the boys performed the play*
false *iff at least one boy didn't participate in a performance of the play*
undef. *otherwise*

This is exactly what our rules for quantification predict. Recall that the bivalent meaning of *all the boys* is just the Montagovian individual corresponding to the plurality of all boys:

(70) $[\![\text{all the boys}]\!](P) = 1$ iff P is true of the maximal plurality of boys

Assume that the scope predicate P is such that it is false of at least one boy. Then P^* is, by definition, false of the plurality of all boys, and so are, consequently, P^{*0} and P^{*1}. So $[\![\text{all the boys}]\!]$ is false of both P^{*1} and P^{*0} and our application rule returns overall falsity. We therefore correctly predict that *all the boys* is false of any predicate that is false of at least one boy.

Now consider a situation where P is true (only) of a superplurality of the boys, say, the plurality that contains the boys and the girls together. By homogeneity, it then has to be undefined of all atomic boys and pluralities of boys. As we move to P^*, nothing changes in this respect. So P^{*0} is false of all boys and pluralities of boys, and P^{*1} is true of the same. Hence $[\![\text{all the boys}]\!]$ is false of P^{*0} and true of P^{*1}, and the result of our application procedure is undefinedness.

3.5 Intermediate Summary

In this section we discussed how to derive the trivalent behaviour of bare plurals from the independently established properties of plural predication. We started from the observation that predicates, as applied to pluralities, are trivalent according to a certain pattern (Križ 2016). If we consider a plural noun *zebras* as a predicate of individuals in this light, we find that it is true of pluralities of zebras, false of non-zebras, but has the third truth value of atomic zebras. In a sense, its meaning is therefore not really exclusive or inclusive, but a superposition of the two. In the case of existential bare plurals, this trivalent predicate is the restrictor of an existential quantifier. To capture the empirical generalisations from Križ & Chemla 2015 about how quantification interacts with the trivalence of plural predication, we defined an algorithm that tells us how to apply a quantifier, defined by its classical bivalent meaning, to a trivalent predicate. When this rule is applied to the case of bare plurals, it follows immediately that *Mary saw zebras* has the third truth value when Mary saw exactly one zebra.

4 Context Dependence

Experimental data form Grimm 2013, 2014 indicate that whether a bare plural is interpreted inclusively or exclusively is dependent on contextual factors. Subjects were presented with questions and asked which answer, *yes* or *no*, is most appropriate in a scenario where there is only one witness.[21] There were two types of questions. The first type was about a concrete physical configuration, and the scenario was presented visually as a picture. For these questions, subjects preferred the answer *no* (68%). To give a concrete example, in a situation where there was only one computer next to the mug, subjects judged the negative answer to (71) to be more appropriate.

(71) Is the mug in this picture sitting next to computers?

The second type of questions concerned more abstract contexts that had to do with rules and regulations, where the scenario was presented as a verbal vignette. Here, subjects were more accepting of — in fact, showed a preference for — *yes* (78%). Such was the case for (72), where subjects predominantly responded that one should answer *yes* if one's team terminated even one project.

(72) Did your team terminate projects this fiscal quarter?

This particular example makes it easy to see a plausible hypothesis as to what is behind this variability: when the difference between one and several is contextually irrelevant, the bare plural receives an inclusive interpretation. The question in (72) would most plausibly be asked in a context where it matters only whether any project was terminated, not how many, and so the *yes*-answer is most appropriate. This is, in fact, the exact intuition that one of Grimm's subjects reported in a comment on the experiment:

> "Even though the question uses the plural of the word 'projects', the intent seems to want the employee to disclose whether any projects were terminated, not just whether more than one project was terminated."

These facts, it turns out, can be accommodated both on our theory and on an implicature-based view of multiplicity. In both cases, the explanations make use of independently established principles governing governing how context interacts with the phenomena in question, and the two explanations are deeply analogous.

21 In addition, there were the obvious control conditions with multiple or no objects.

4.1 Plural Predication and Non-Maximality

In order to show how the context-dependent behaviour of bare plurals can be accommodated on our view, let us again start by looking at sentences with definite plurals. It is well-known that in some contexts, such plural predications are tolerant to exceptions, that is to say, we accept the sentence even when the predicate is not actually true of the whole plurality (Brisson 1998; Lasersohn 1999; Malamud 2012; Križ 2016). Following Brisson's terminology, we call this phenomenon *non-maximality*. The literature is unanimous about the intuitive characterisation of this phenomenon: exceptions can be ignored when they are irrelevant for the current purposes of the conversation. This is well illustrated by examples from Lasersohn 1999.

In (73), our interest in the conversation is presumed to be just a general picture in broad strokes of what's going on in the town. In this case, isolated individuals who are not, in fact, asleep will not matter as long as they aren't doing anything interesting, and so the sentence *The townspeople are asleep* can be felicitously used even if there are some such exceptions.

(73) *What's going on in the town?*

The townspeople are asleep.
⇝ More or less all of them are, and the ones still awake aren't doing anything of note.

(74), on the other hand, has a context where it is clear that it is very relevant whether all or only almost all of the subjects are asleep, and in this case, the sentence with the definite plural receives a strict interpretation.

(74) *We are waiting for all subjects to be asleep to start our sleep study.*

The subjects are asleep.
⇝ All of them are asleep and we can start.

Križ 2016 gives an account of this phenomenon that goes with the trivalent conception of plural predication, where non-maximality is a consequence of how pragmatic principles govern the use of trivalent sentences. The interaction of the principles he postulates amounts to, essentially, the following: an undefined sentence, but not a false sentence, can be used when the situation described is, for current purposes, equivalent to a situation where the sentence is literally true.[22]

Sentences with bare plurals are trivalent in the same way as those with definite plurals, and so the same pragmatic principles apply. Let us consider how this can make sense of Grimm's experimental data. The first thing to note is that in the experimental items, the bare plural is inside a question, not inside a declarative. While Križ's theory gives us a way of interpreting a trivalent proposition against the backdrop of the purposes of the conversation, the theory is not immediately applicable to questions. While such an extension may be possible and ultimately

22 This paraphrase is actually not fully equivalent to the formal statement given in Križ 2016, where there is an additional condition. This does not, however, matter for any of the applications we are going to discuss, and so we refrain from repeating the details here.

desirable, it seems to us that we can do without it at this point by making a simple and plausible assumption: when *yes* is uttered in response to a *yes/no*-question $?p$, it amounts to an assertion of p. In our case, p happens to be a trivalent proposition, and so we can treat the *yes*-answer as the assertion of a trivalent proposition to which Križ's theory is applicable.

When evaluating a *yes*-response to (72), we are thus effectively dealing with an assertion that *the team terminated projects this fiscal quarter*. This assertion is felicitous as long as (i) the team terminated at least one project, and (ii) the situation is, for current purposes, equivalent to one in which the team terminated more than one project. What the current purposes are, is of course unclear to the subject, so they have to imagine different contexts in which someone might ask (72) and take a guess.

If the subject concludes that the asker likely is just interested in knowing whether any projects were terminated, then the *yes*-response will be felicitous. It is, effectively, an assertion of *the team terminated projects* that is evaluated against the rather coarse-grained purposes of the asker and, by non-maximality, ends up with the overall meaning that the team terminated one or more projects.

But what if the subject decides that the context is likely not such that an utterance of *yes* is felicitous? Then they are left with the alternative of *no*. It suggests itself to assume, in analogy to what we have done for *yes*, than a reply of *no* in response to a question $?p$ amounts to an assertion of the negation of p. In this case, this would be an assertion that *the team didn't terminate projects this fiscal quarter*. Now according to Križ's rules, this assertion can only be made if the situation is, for current purposes, equivalent to one where the team terminated no projects. It seems very implausible that a subject who recommends a *no*-response necessarily thinks the context supports this equivalence. We therefore have to assume that in such a forced-choice paradigm, the *no*-response is not, in fact, to be treated as analogous to the *yes*-response, that is, it does not amount to an assertion that *the team didn't terminate projects*. Rather, *no* functions as a more generic negative response whose significance is really just to indicate that the affirmative response *yes* is not appropriate.

Finally, to complete our account of Grimm's data, we have to assume that when, as in (71), concrete objects rather than instances of rather abstract administrative categories are concerned, subjects find it less plausible that the context of the question is such that it is irrelevant whether there is only one object or several, and so they are less inclined to choose the *yes*-response.

4.2 Questions, Answers, and Implicatures

Implicatures are usually regarded as context dependent: when the distinction between two alternatives is irrelevant for the purposes of the conversation, then the corresponding implicatures is not drawn (cf. e.g. Magri 2009). This makes essentially the same prediction as our above analysis in terms of non-maximality: the multiplicity implication is absent when the difference between one and several is irrelevant for current purposes.

Let us now consider how this could be applied to Grimm's data. The first option is to locate the action in the interpretation of the question. A question $?p$ would have to be interpreted as "Is p together with its implicatures true?", which would mean that exhaustification happens locally inside the question. There are good reasons to think that this is impossible. First, in the absence of elaborate prosodic focus on the scalar item, local exhaustification is generally taken to be impossible in downward-entailing contexts, and questions appear to count as downward-entailing for the purposes of linguistic phenomena that are sensitive to logical monotonicity, such as NPI. Second, local exhaustification is clearly impossible in embedded questions: the sentence in (75) cannot receive the interpretation indicated.

(75) John knows whether Mary ate some of the cookies.
*'John knows whether Mary ate some, but not all of the cookies.'

The natural understanding of the sentence is that John knows whether Mary ate some of the cookies, but does not know whether she ate all of them (Sudo & Spector to appear). But it is precisely in such a situation that the interpretation in (75) would be false, rather than true.

An explanation of Grimm's data on the implicature view therefore has to have the same structure as the explanation we proposed in terms of our own theory: it must have to do with the interpretation of answers. Again, we might assume that an answer *yes* in response to a question $?p$ behaves like an assertion of p and therefore also triggers the implicatures that p would have in the context. This is not altogether implausible as a general rule. For example, it seems quite easy to understand B's reply in (76) as implicating that Mary didn't eat all of the cookies.

(76) A: Did Mary eat some of the cookies?
B: Yes. $\rightsquigarrow$ Not all of them.

Furthermore, if B wants to add that Mary ate all of the cookies, they most naturally continue in a way that is typical for implicature cancellation, such as with *in fact*.

(77) A: Did Mary eat some of the cookies?
B: Yes. All of them, in fact.

Our interpretation of the *no*-responses then has to be the same as previously on the homogeneity-based theory: at least in the forced choice context of the experiment, *no* functions as an indicator that the *yes*-response would be inappropriate.

4.3 Intermediate Summary

Experiments by Grimm (2013, 2014) indicate that the presence of the multiplicity inference is dependent on context and subject matter. Since implicatures, and considerations of competition, are naturally context sensitive as well, these approaches can deal with the facts quite naturally. There is, however, *prima facie* a puzzle for our approach, which postulates an essentially inclusive reading, so that the sentence is undefined when the multiplicity implication is not true. We have argued that the phenomenon can be unified under with so-called non-maximal readings known for sentences with definite plurals, where sentences are also acceptable under certain conditions despite being, in terms of their strict semantics, undefined rather than true. Again, the conceptual resources from the analysis of definite plurals are naturally carried over to the case of bare plurals.

5 Applications

This section presents and application of our theory to a number of concrete examples, such as bare plurals in various embedded environments, with an empirical evaluation of the predictions and comparison with alternative approaches.

5.1 Negation and Undefinedness

A distinctive prediction of our approach concerns the interaction of bare plurals with negation. Negation has its usual meaning in trivalent logic: it switches truth and falsity, but leaves undefinedness the same. This, of course, was our initial motivation for assuming trivalence in the first place: it enables us to explain how an affirmative sentence and its negation have non-complementary truth conditions. This makes a crucial prediction: both (78a) and (78b) have the same undefined status in a situation where Mary saw exactly one zebra.

(78) Situation: *Mary saw exactly one zebra.*
   a. #Mary saw zebras. (undef.)
   b. #Mary didn't see zebras. (undef.)

This contrasts with the implicature approach, which does not predict this equal status. In the situation in question, the affirmative sentence is literally true, but has a false implicature. This plausibly gives rise to an intuition that the sentence has some sort of intermediate status, and indeed literally true sentences with false implicatures have been found in at least one experimental paradigm to behave similarly to sentences that are undefined due to homogeneity (Križ & Chemla 2015).

(79) Situation: *Mary saw exactly one zebra.*
   #Mary saw zebras.
   (lit) Mary saw one or more zebras. (true)
   (impl) Mary saw more than one zebra. (false)

Things are different for the negated sentence, however. Its literal meaning is plainly false in the situation in question, and there is no implicature. Thus, the sentence is predicted not to have an intermediate status of some sort, but to be simply false.

(80) Situation: *Mary saw exactly one zebra.*
#Mary didn't see zebras.
(lit) Mary didn't see one or more zebras. (false)
(impl) —

Thus, the implicature approach predicts that negative sentences with bare plurals should behave differently from affirmative sentences on diagnostics that distinguish falsity from undefinedness.

We find the difference between falsity and undefinedness difficult enough to access introspectively to be concerned that theory-based biases are likely to influence the judgement, and therefore profess agnosticism as to the correctness of either prediction. One may be hopeful, however, that an experimental test with naive subjects can be devised, perhaps along the very lines of Križ & Chemla 2015.

5.2 Non-Monotonic Contexts

Recall from section 2.1 that sentences with bare plurals in the scope of a non-monotonic quantifier pose a problem for most analyses of bare plurals in the literature, with the sole exception of Spector's (2007).

(81) Exactly one girl saw zebras.
'One girl saw multiple zebras and the others saw none.'

According to the homogeneity theory, *saw zebras* is undefined of an individual that saw only one zebra in the same way as *bought the books* is undefined of an individual that bought only some of the books, and so we predict that the two should behave analogously when embedded under quantifiers. The behaviour of definite plurals in the scope of quantifiers is independently known from experiments by Križ & Chemla (2015): when *exactly one* has a scope predicate containing a definite plural, the sentence is true if and only if the scope predicate is true of one restrictor individual and false of all the other restrictor individuals.

(82) Exactly one girl bought the books.
'One girl bought all the books and the others bought none of them.'

This is, indeed, analogous to (81), and so the homogeneity approach's basic prediction is borne out.

Let us now consider what the general rules for quantification that we have set forth in 3.2 have to say about this case. These rules worked by taking as given a (bivalent) generalised quantifier that is defined for bivalent predicates and telling us how to apply it to trivalent predicates. The first step, then, is to note the bivalent meaning of *exactly one girl*:

(83) $[\![\text{exactly one girl}]\!](P) = 1$ iff
a. there is a girl x such that $P(x) = 1$, and
b. for all atomic girls $y \neq x$, $P(y) = 0$.

No other bivalent meaning would do, since it would make obviously wrong predictions for bivalent scope predicates (or rather, scope predicates that are defined for all atoms). The trivalent truth conditions we have noted then follow straightforwardly. By our rule, (84) is true iff both (84a) and (84b) are true. (84a) means that P is true of exactly one girl, (84b) means that P is true or undefined of exactly one girl. Taken together, they entail that P is true of one girl and false of all the others.

(84) Exactly one girl is P.

a. Exactly one girl is P^{*0}.

b. Exactly one girl is P^{*1}.

One might think that not much is explained here, since we have just stipulated truth conditions: one can take the definition in (84) as it is and apply it to trivalent predicates to obtain the right result. This, however, is an illusion: we could equivalently have stated clause (84b) as in (85) while defining the very same bivalent quantifier:

(85) for all atomic girls $y \neq x$, $P(y) \neq 1$.

This formulation is equivalent to (84b) for bivalent Q, but not for trivalent Q: the change to (85) yields truth for a wider range of trivalent scope predicates. It is thus clear that the bivalent definition of the quantifier underdetermines its trivalent truth conditions, and the work that our rule for quantification does is to reduce this underdetermination so that a unique trivalent quantifier is associated with every bivalent one.

The undefinedness and falsity conditions we predict for *exactly one girl* are somewhat involved and best stated disjunctively. (84) is undefined if (84a) is true and (84b) is false or *vice versa*, that is to say:

(86) *Exactly one girl is P* is undefined iff

a. P is true of one girl and undefined of at least one other girl; or

b. P is undefined of one girl and false of all other girls.

Falsity, finally, is obtained if both (84a) and (84b) are false.

(87) *Exactly one girl is P* is false iff

a. P is false of all girls; or

b. P is true of more than one girl; or

c. P is true of no girl and undefined of more than one.[23]

These various types of situations were tested by Križ & Chemla (2015) for verb phrases containing a definite plural, and the pattern they found is what we have just described. Our theory predicts the same pattern for predicates containing a bare plural.

23 This case is noteworthy because it constitutes a deviation from the Strong Kleene semantics for the first-order definition of *exactly one*, according to which the sentence would be undefined, rather than false, here.

5.3 Embedding under Universals

In accordance with section 3.2.1, our predictions for sentences with bare plurals in the scope of a universal quantifier is as in (88).[24]

(88) Every girl saw zebras.
true *iff every girl saw multiple zebras*
false *iff at least one girl saw no zebra*
undef. *otherwise*

Note, however, that, taking into account the phenomenon of non-maximality discussed in section 4.1, we do not predict these very strong truth conditions to surface under all circumstances. Rather, what we predict is that sentence can be used to describe situations where every girl saw at least one zebra (so that the sentence is not false) and which are, for current purposes, equivalent to a situation where every girl saw multiple zebras. It is quite possible, for example, that in some context, we may not care about how many girls, exactly, saw multiple zebras as long as some of them did, in which case the overall interpretation would be that *every girl saw one or more zebras and some saw more than one*.

Our strong truth conditions in (88) are also what Spector's (2007) implicature theory predicts based on the competition between (89a) and (89b).[25]

(89) a. Every girl saw zebras.
b. Every girl saw a zebra.

The exhaustified meaning of (89a) is as in (90).

(90) EXH (every. . . a) = (every . . . a) and not (every. . . 2) and not (every. . . 3) . . .
'Every girl saw a zebra and at least one girl saw only one.'

Now (89a) has the negation of (90) as its implicature:

(91) EXH (EXH (every. . . pl)) = (every. . . a) and not EXH (every. . . a) =
(every. . . a) and ((every. . . 2) or (every. . . 3) or. . .) = every. . . 2
'Every girls saw more than one zebra.'

Based on the assumption that implicatures are calculated only when relevant, the sentence of interest then has two possible interpretations:

(92) a. Every girl saw one or more zebras.
b. Every girl saw more than one zebras.

24 Note that it is important to use the universal quantifier over atoms *every*, since *all* allows for so-called *dependent plurals*. For example, in (i) (from Ivlieva 2013), there is no implication that any one of the boys attends more than one school; rather, there is only an implication that the boys attend more than one school overall, i.e. that they don't all go to the same school.

(i) All the boys attend good schools.

25 This is so also for the other competition-based theories in the literature. The discussion of intermediate readings which follows below, however, is specific to Spector's.

This is in contrast to our theory, which allows for finer gradations, in that the overall reading, taking non-maximality into account, is, so to say, always as strong as is relevant, and can fall in-between (92a) and (92b). It could, for example, be that what is necessary for a situation to be equivalent, for current purposes, to one where every girl saw multiple zebras is that every girl saw at least one zebra and *most* of the girls saw more than one zebra.

However, Spector (2007) points out that at least one additional reading is predicted if one takes into account that the sentence contains two scalar expressions: the bare plural and the quantifier *every*, the latter of which stands in competition with *some*. Then the alternatives with respect to which (93a) is exhaustified is as in (93b), and the overall result is (94).[26]

(93) a. Every girl saw zebras.
b. Every girl saw two/three/... zebras, some girl saw two/three/... zebras.

(94) EXH (every... a) = (every... a) and not (every... 2) and not (every... 3) ... and not (some... 2) and not (some... 3)...
'Every girl saw exactly one zebra.'

Now the implicature of the sentence with the bare plural is the negation of (94):

(95) EXH (EXH (every... pl)) = (every... a) and not EXH (every... a) = (every... a) and (some... 2)
'Every girl saw at least one zebra and at least one girl saw more than one.'

In order to derive this reading of intermediate strength, it is crucial that one can simultaneously replace *every* by *some* and *a* with a numeral when generating the alternatives of *every girl saw a zebra*.[27] The possibility of such multiple replacements in implicature derivation has, indeed, frequently been denied,[28] and so the prediction of the existence of (95) is not set in stone. But let us assume for the moment that multiple replacement is possible. Then an interesting question can be asked: does the implicature theory predict that the weaker reading (95) arises under the same contextual conditions under which our theory would predict this reading as an instance of non-maximality?

Our theory predicts that the interpretation in (95) arises from non-maximality under the following conditions: it is relevant, for the purposes of the conversation, whether any girl saw more than one zebra, but it is not relevant how many. That is to say, *some girl saw multiple zebras* is relevant but, for example, *every girl saw multiple zebras* is not.

An explicit view that has been advanced on how the relevance of alternatives comes into play in the computation of implicatures is as follows (Magri 2009):

26 For the detailed step-by-step derivation, the reader is again requested to consult the original publication.

27 Additional readings of intermediate strength are possible on the implicature theory if instead of *some*, we take the alternative to *every* to be *many* or *most*. Note that it does not do to have a scale ⟨*every, many, some*⟩: the weak reading arises as soon as *some* is on the scale at all, no matter which other items are.

28 See Fox 2007, Magri 2009, Chemla 2009, Romoli 2012, Trinh & Haida 2015. See Chemla & Romoli 2015 for a new perspective.

whenever EXH is applied somewhere, the set of alternatives relative to which exhaustification is performed is filtered by contextual relevance. That is to say, one generates all the formally possibly alternatives by replacing various scalar items, and then removes those alternatives which express propositions that are not relevant for the purposes of the conversation. Going with this picture, we find that there is only one fact that matters for the derivation of the intermediate reading: whether or not *some girl saw two (or more) zebras* is relevant. If it is relevant, then it is negated as we compute the strengthened meaning of *every girl saw a zebra*, as in (94). Since the negation of *some girl saw two zebras* entails the negation of *every girl saw two zebras*, it does not matter whether the latter is actually in the set of alternatives.

Thus, we find a difference here between the two approaches: the implicature theory predicts that the intermediate reading arises as soon as it is relevant whether any girl saw more than one zebra, regardless of whether it is relevant whether *every* girl saw more than one zebra. Our theory, on the other hand, predicts that an intermediate reading is only possible when the latter question is specifically *ir*relevant. As for the strong reading, the implicature theory predicts that it arises only if (i) it is relevant whether every girl saw more than one zebra and (ii) it is irrelevant, assuming that not every girl saw multiple zebras, whether any girls saw multiple zebras. Our theory, in contrast, predicts that the strong reading is obtained whenever it is relevant whether all girls saw more than one zebra, no matter what else is relevant or irrelevant. These predictions are summarised in Table 1.

What is relevant		Predicted reading	
some … several	*every… several*	implicature theory	homogeneity theory
✓	✓	intermediate	strong
✓	✗	intermediate	intermediate
✗	✓	strong	strong
✗	✗	(weak)	(weak)

Table 1: Predicted readings depending on what is relevant in the conversation

We are inclined to think that the strong reading is, indeed, forced as soon as *every… several* is relevant, but the question is certainly difficult to test in any decisive way. This is, in particular, because what is taken to be relevant for the purposes of the conversation is vexingly resistant to explicit and direct manipulation.[29] Still, granting the correctness of our intuition, there is *prima facie* a problem for Spector's theory: either one assumes that multiple replacement is not possible in implicature calculation, but then one predicts that the intermediate reading is never available; or one assumes that it is possible, in which case the intermediate reading is predicted to be available too easily.

Spector (p.c.) suggests that this dilemma may be evaded by making a different assumption about how relevance enters into implicature computation. Instead of filtering the set of alternatives by relevance during each application of EXH,

29 Cf. the discussion in Križ 2016 on how the purposes of the conversation cannot simply be set to a certain value by an explicit question.

one proceeds thus: one computes all possible overall meanings of the utterance based on the various possible sets of alternatives, and then chooses among these meanings the logically strongest one that is relevant.[30] More concretely, the set of readings generated for *every girl saw zebras*, depending on various scales of alternatives, is as in (96).

(96) Every girl saw at least one zebra and...

a. — $\varnothing$
b. every girl saw multiple zebras. $\langle$a, two$\rangle$
c. most girls saw multiple zebras. $\langle$every, most$\rangle$, $\langle$a, two$\rangle$
d. some girls saw multiple zebras. $\langle$every, some$\rangle$, $\langle$a, two$\rangle$
e. ...

These readings have various logical strengths, and the one observed in a given context is the logically strongest one that is relevant. This result is indistinguishable from the predictions of our theory; in particular, the strong reading will be chosen over any of the intermediate readings whenever *every... several* is relevant. A comparison therefore hinges entirely on conceptual plausibility rather than empirical differences. It seems to us that Spector's proposed interpretation procedure suffers a plausibility cost for being rather ad hoc and no longer traceable to anything resembling Gricean reasoning or a natural conventionalisation or grammaticalisation thereof. To be sure, it uses similar technical tools and notions, in that possible readings are obtained by taking alternatives and negating them, but the overall reading that is generated by the procedure is not at all related to any consideration of what the speaker could have said instead. Our own theory, in contrast, uses exclusively conceptual resources that are independently motivated from the analysis of non-maximality with definite plurals.

5.4 *well*-Answers

One feature of homogeneity is the kind of answers that are appropriate to give to sentences that have the third truth value. These sentences, being intuitively somehow half-true, license a response with a somewhat drawn-out, pensive *well*.[31]

(97) A: Mary bought the books.
B: Weeell... She bought HALF of them.

For sentences that are literally true, but have a false implicature (and are therefore misleading), this response pattern is not particularly natural.

(98) A: Mary ate some of the cookies.
B: ?#Weeell... She ate ALL of them.

30 This procedure is inspired by the recent Fox & Spector to appear, in which the application of an economy condition also requires computing all possible overall meanings of an exhaustified sentence based on all possible sets of alternatives.

31 In the author's native German, the particle *naja* plays the exact same role. However, the existence of a particle with this function appears to be not even close to universal cross-linguistically.

If *well* is to be used here at all (our preferred particle would be *why*), we find that it needs to be short, and while there is naturally a prosodic boundary between it and the following correction, there must not be an actual pause. This is clearly a distinct use of *well* from the drawn-out, pensive kind in (97).

Conversely, *indeed* or *in fact*, which naturally introduce implicature cancellations, do not go well with homogeneity violations.

(99) A: Mary ate some of the cookies.
B: Indeed / In fact, she ate ALL of them.

(100) A: Mary bought the books.
B: ?#Indeed / In fact, she bought HALF of them.

It appears to us that a violation of the multiplicity inference of bare plurals patterns with homogeneity in this respect and not with implicatures, in the that hesitant, somewhat conceding *well*-response is quite possible when there is only one witness to the existential quantification, whereas *indeed* is rather odd.

(101) A: Mary went on safari yesterday. She saw zebras.
B: Weeell... She saw ONE.

(102) A: Mary went on safari yesterday. She saw zebras.
B: ?#Indeed / In fact, she saw ONE.

We have thus identified a further area in which sentences with bare plurals behave like other sentences that are trivalent due to homogeneity, and unlike standard cases of scalar implicatures. The onus is therefore on the implicature theorist to distinguish the multiplicity implicature on principled grounds from other cases of scalar implicature and show how the relevant behaviour can be sensitive to this difference.

5.5 Pluralia Tantum

Recall from section 2.3 that *pluralia tantum*, and possibly some mass nouns, were pointed out as a problem for competition-based theories of the multiplicity implication, because they have a multiplicity implication even though there is no plausible singular alternative.

This is not a problem for our theory, where a singular alternative plays no role at all. All that matters is that the meaning of the noun is trivalent. To be sure, it would be perfectly possible, on our theory, for *clothes* (and *belongings*, etc.) to have meanings that are true of both single items of clothing (and property, etc.) and of pluralities of such items, as these meanings would not violate homogeneity. We can, therefore, not claim to have *predicted* a plurality inference for these words. But we are not prevented, either, from assuming that these words have meanings that are only true of pluralities, in which case homogeneity forces them to be undefined of atomic items. And so, unlike the competition-based theories, ours can at least accommodate the fact that there is a multiplicity implication here.

5.6 Infelicitous Questions

Orin Percus (p.c.) observes that a theory of bare plurals ought to explain why questions such as those in (103) are incoherent rather than having the reading here indicated.

(103) a. #Which of these mothers have children?
'Which of these mothers have more than one child?'
b. #Which of these novelists have written books?
'Which of these novelists have written more than one book?'

If the bare plural is interpreted inclusively, the explanation is immediate: the questions are tautologous, since by definition, every mother has one or more children and every novelist has written one or more books.[32] The implicature theory can simply avail itself of this reasoning. Questions appear to count as downward-entailing environments for both the licensing of negative polarity items and the computation of local implicatures, so the literal inclusive meaning would not be locally strengthened to an exclusive meaning, and the interpretation of the question ends up defective in the way we just described.

Farkas & de Zwart's (2010) theory, while being competition-based just like the implicature theories, has more trouble in dealing with this case. On this theory, we reason on the basis of the competition between the plural and the singular about how to locally interpret the bare plural, and considerations of the logical strength and relevance of the overall meaning that would result on each interpretation play a role. It is not very clear why saving a sentence from infelicity should not be sufficient to push us to interpret the bare plural exclusively here.

Our own theory is, *prima facie*, even worse off because in terms of its truth conditions, our interpretation of the bare plural already *is* exclusive, and so the questions (and their possible answers) are not tautologous. There is, however, still a way in which the possible answers to these questions are defective: they are trivalent propositions that can, given the contextual presuppositions, only be true or undefined, but never false.

We take it that the possible answers, in the standard sense of Hamblin (1973), are, informally, obtained by replacing the *wh*-phrase with an individual in the restrictor of the *wh*-phrase. Then the answers to (103a), for example, are propositions of the form *x has children* where *x* is one or several of the mothers that are being referred to by the phrase *these mothers*. Assuming that (103a) was uttered by a speaker pointing at Anne, Mary, and Sue, (104) is one of these possible answers.

(104) Anne has children.

On our theory, this answer is a trivalent proposition that is false only if Anne doesn't have a child. But since there is a contextual presupposition that Anne is a mother, we know that this cannot be the case: either she has only one child, in which case (104) is undefined, or she has more than one child, in which case it is

32 This is actually a simplification in the latter case, which we will return to below.

true. Let us call such a proposition — one that cannot be false, but only true or undefined — a *quasi-tautology*. It is easy to see that all possible answers to (103a) are quasi-tautologies.

It is natural to assume that there is something defective about asserting contextual tautologies, and therefore that asking a question all of whose answers are contextual tautologies is infelicitous. In order to account for the facts, all we need to do is to extend this prohibition to quasi-tautologies. To be sure, it is not as intuitively obvious why quasi-tautologies should be forbidden; after all, they could conceivably be used to convey information because they distinguish certain worlds where they are true from certain other worlds where they are undefined. As such, our explanation is slightly more stipulative.

There is, however, reason to think that even on the implicature theory of multiplicity we would need a prohibition on quasi-tautologies (or, more narrowly, questions that have them as answers) to fully explain the facts. Once we take a closer look at our second example (103b) from above, it turns out that the implicature theory's explanation for this case hinges crucially on our simplifying assumption that being a novelist entails having written at least one book. It is, to be sure, exceedingly rare for novels to be the result of collaboration between two authors, but it is not impossible.[33] Yet it seems to us that the two co-authors of a novel would both appropriately be called *novelists*, while we could not say of either of them individually that that person has written a novel. The latter is so because homogeneity exists not only with respect to the parthood relation between pluralities and their constituent atoms, but also with respect to other sorts of parthood relations such as those between a book and its chapters (Löbner 2000).

(105) Niven wrote a novel.

true *iff Niven wrote a whole novel.*
undef. *iff Niven wrote only part of a novel.*
false *iff Niven didn't contribute to any novel.*

Consequently, (106b) is undefined in the scenario in question, unlike the nominal predication (106a) and the participatory predication (106c).

(106) Context: *Niven and Pournelle collaboratively wrote a novel.*

a. ✓Niven is a novelist.
b. ?#Niven has written a novel.
c. ✓Niven has co-written a novel.

To be sure, we do not wish to claim that (106b) is always unacceptable. It is thinkable that *write a book* could be coerced into a predicate of participation, or that by way of non-maximality, the sentence would be counted as true because for current purposes, all that matters is any authorial activity, so that having written part of a published book is as good as having written a whole book. Nevertheless, one can equally ealisy imagine contexts where one would be hesitant to assert, or

33 I thank Viola Schmitt (p.c.) for bringing to my attention the example of Larry Niven and Jerry Pournelle's *The Mote in God's Eye*.

assent to, (106b), whereas (106a) just strikes us as plainly true. If these judgements are correct, then the implicature theory's explanation for the infelicity of (103b) turns out to be insufficient. For now it is no longer the case that all answers to the question are tautologies, since being a novelist does not entail having written a book (only part of a book). The answer (107), for example, is, even on an inclusive reading of the bare plural, only a quasi-tautology: it is undefined if Smith only wrote part of a book.

(107) Smith has written books.

Thus, once we take into account the fact that homogeneity also exists with respect to other parthood relations, we find independent support for the generalisation that questions whose answers are all quasi-tautologous are infelicitous.

5.7 Intermediate Summary

In this section, we have discussed how our theory deals with a number of details in the behaviour of bare plurals. In some cases, it does so equally well as existing theories, while in others, we find it to be superior.

- There is a hard-to-evaluate differential prediction between our theory and the implicature-based accounts for negated sentences, which we predict to have an intermediate status in some cases where the implicature approaches predict them to be plainly false (section 5.1).
- Our theory, like Spector's (2007), but unlike the other competition-based theories, correctly predict the meaning of sentences with bare plurals in the scope of a non-monotonic quantifier like *exactly* (section 5.2).
- Our theory's predictions sentences with bare plurals in the scope of *every* are subtly different from and superior to those of Spector's (2007). These predictions can be replicated in principle by Spector's approach, but only at the cost of assumptions of questionable plausibility (section 5.3).
- When it comes to natural reaction in discourse to sentences which, according to our theory, are undefined, we predict parallelism with the case of definite plurals, which appears to be borne out empirically. The implicature approach, on the other hand, incorrectly predicts parallelism with scalar implicatures rather than definite plurals.
- Our theory has no trouble accounting for *pluralia tantum*, which we identified at the outset as a problem for competition-based approaches (section 5.5).
- Finally, we discussed how our theory can account for the infelicity of a certain class of questions with bare plurals. These are unproblematic on the implicature approach, but the additional principle we have to postulate to explain them on our theory is independently supported.

6 Conclusion

The theory we have presented analyses sentences with existential bare plurals as logically trivalent in the same way as sentences with definite plurals. *Mary saw zebras* has the third truth value when Mary saw exactly one zebra, just as *Mary bought the books* has the third truth value when Mary bought some, but not all of the books. We have argued that both instances of trivalence can be traced to the way in which natural language predicates are systematically trivalent when applied to pluralities. In the case of bare plurals, it is the plural noun that is trivalent: it is true of pluralities of zebras, but has the third truth value of atomic zebras. The behaviour of sentences with bare plurals is then a consequence of the very general principles that govern how (existential and other) quantifiers apply to trivalent predicates.

Our theory thus uses only independently established conceptual resources. It differs fundamentally from existing theories, which are based, in one way or another, on the idea of competition between the singular and the plural, but all require some special-purpose assumptions. Furthermore, our theory allows us to overcome some limitations of these theories: in particular, it correctly predicts the behaviour of existential bare plurals in the scope of quantifiers, and it can deal with plural nouns that have no singular form (and for which no competition with the singular is therefore possible).

References

Breheny, Richard. 2005. Exhaustivity, homogeneity, and definiteness. In Paul Dekker & Michael Franke (eds.), *Proceedings of the fifth Amsterdam Colloquium*, 59–65.

Brisson, Christine. 1998. *Distributivity, Maximality, and Floating Quantifiers*: Rutgers dissertation.

Büring, Daniel & Manuel Križ. 2013. It's that and that's it! Exhaustivity and homogeneity presuppositions in clefts (and definites). *Semantics & Pragmatics* 6.

Champollion, Lucas. 2010. *Parts of a whole: Distributivity as a bridge between aspect and measurement*: University of Pennsylvania dissertation.

Chemla, Emmanuel. 2009. Similarity: towards a unified account of scalar implicatures, free choice permission and presupposition projection. Manuscript, LSCP.

Chemla, Emmanuel & Jacopo Romoli. 2015. The role of probability in the accessibility of scalar inferences. Ms. LSCP and Ulster University.

Dowty, David. 1979. *Word meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, David. 1987. A note on collective predicates, distributive predicates and *all*. In *Proceedings of escol 86*, 97–115.

Farkas, Donka & Henriëtte de Zwart. 2010. The semantics and pragmatics of plurals. *Semantics & Pragmatics* 3(6). 1–54.

Fox, Danny. 2007. Free choice and the theory of scalar implicatures. Ms. MIT.

Fox, Danny & Benjamin Spector. to appear. Economy and embedded exhaustification. *Natural Language Semantics* .

Gajewski, Jon. 2005. *Neg-raising: Polarity and Presupposition*: MIT dissertation.

Grimm, Scott. 2013. Plurality is distinct from number neutrality. In Yelena Fainleib, Nicholas LaCara & Yangsook Park (eds.), *Proceedings of NELS 41*, .

Grimm, Scott. 2014. Inclusive and exclusive plurality: Experimental investigations. Presented at New Directions in Experimental Semantics, The Graduate Center, CUNY.

Hamblin, Charles L. 1973. Questions in Montague English. *Foundations of Language* 10(1). 41–53.

Ivlieva, Natalia. 2013. *Scalar implicatures and the grammar of plurality and disjunction*: MIT dissertation.

Ivlieva, Natalia. 2014. Multiplicity and non-monotonic environments. In Luka Crnič & Uli Sauerland (eds.), *The art and craft of semantics: A festschrift for Irene Heim*, vol. 1, 245–251. Cambridge, MA: MITWPL.

Ivlieva, Natalia & Yasutada Sudo. 2015. Another problem for alternative-based theories of plurality inferences: The case of mass plurals. *Snippets* 29. 7–8.

Katzir, Roni. 2007. Structurally defined alternatives. *Linguistics and Philosophy* 30. 669–690.

Križ, Manuel. 2015. *Aspects of homogeneity in the semantics of natural language*: University of Vienna dissertation.

Križ, Manuel. 2016. Homogeneity, non-maximality, and 'all'. *Journal of Semantics* 33(3). 493–539. doi:10.1093/jos/ffv006.

Križ, Manuel & Emmanuel Chemla. 2015. Two methods to find truth-value gaps and their application to the projection problem of homogeneity. *Natural Language Semantics* 23(3). 205–248. doi:10.1007/s11050-015-9114-z.

Križ, Manuel & Benjamin Spector. 2017. Interpreting plural predication: Homogeneity and non-maximality. Ms. IJN.

Lasersohn, Peter. 1999. Pragmatic halos. *Language* 3(75). 522–551.

Link, Godehard. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In R. Baeuerle, C. Schwarze & Arnim von Stechow (eds.), *Meaning, use and interpretation of language*, de Gruyter.

Link, Godehard. 1987. Generalized quantifiers and plurals. In Peter Gärdenfors (ed.), *Generalized quantifiers: Linguistic and logical approaches*, Reidel.

Löbner, Sebastian. 2000. Polarity in natural language: predication, quantification and negation in particular and characterizing sentences. *Linguistics and Philosophy* 23. 213–308.

Magri, Giorgio. 2009. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics* 17(3). 245–297.

Magri, Giorgio. 2011. The plurality inference of object mass nouns. *Snippets* 24. 9–10.

Magri, Giorgio. 2014. An account for the homogeneity effects triggered by plural definites and conjunction based on double strengthening. In Salvatore Pistoia Reda (ed.), *Pragmatics, semantics and the case of scalar implicatures* Palgrave Studies in Pragmatics, Language and Cognition, 99–145. Houndsmills: Palgrave Macmillan.

Malamud, Sophia. 2012. The meaning of plural definites: A decision-theoretic approach. *Semantics & Pragmatics* 5. 1–58.

Mayr, Clemens. 2015. Plural definite NPs presuppose multiplicity via embedded exhaustification. In Sarah D'Antonio, Marry Moroney & Carol Rose Little (eds.), *Proceedings of SALT 25*, 204–225.

Romoli, Jacopo. 2012. *Soft, but strong. neg-raising, soft triggers, and exhaustification*: Harvard University dissertation.

Sauerland, Uli. 2003. A new semantics for number. In Rob Young & Yuping Zhou (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) 13*, 258–275. Ithaca, NY: CLC Publications, Cornell University.

Sauerland, Uli, Jan Anderse & Kazuko Yatsushiro. 2005. The plural is semantically unmarked. In Stephan Kepser & Marga Reis (eds.), *Linguistic evidence*, 413–434. Berlin: Mouton de Gruyter.

Schwarz, Florian. 2013. Maximality and definite plurals – experimental evidence. In Emmanuel Chemla, Vincent Homer & Grégoire Winterstein (eds.), *Proceedings of Sinn und Bedeutung 17*, 509–526.

Schwarzschild, Roger. 1994. Plurals, presuppositions and the sources of distributivity. *Natural Language Semantics* 2(3). 201–248.

Spector, Benjamin. 2007. Aspects of the pragmatics of plural morphology: On higher-order implicatures. In Uli Sauerland & Penka Stateva (eds.), *Presupposition and implicature in compositional semantics* Palgrave Studies in Pragmatics, Language and Cognition, chap. 9, 243–281. Palgrave Macmillan.

Sudo, Yasutada & Benjamin Spector. to appear. Presupposed ignorance and exhaustification: How scalar implicature and presupposition interact. *Linguistics and Philosophy* .

Trinh, Tue & Andreas Haida. 2015. Constraining the derivation of alternatives. *Natural Language Semantics* 23(4). 249–270.

Winter, Yoad. 2001. *Flexibility principles in boolean semantics*. Cambridge, MA: MIT Press.

Zweig, Eytan. 2008. *Dependent plurals and plural meaning*: NYU dissertation.

Zweig, Eytan. 2009. Number-neutral bare plurals and the multiplicity implicature. *Linguistics and Philosophy* 32. 353–407.

A Homogeneity and Collectives: Empirical Arguments

In this section, we would like to substantiate the assumptions about homogeneity with collective predicates which form the basis of section 3.4. While no experimental investigation of this issue has been performed, we believe that the case can be made on the basis of introspective judgements using, essentially, two diagnostics. The first diagnostic is the simple acceptability of both the affirmative and the negated sentence: when neither seems appropriate to describe the situation, we have good reason to suspect homogeneity. The second diagnostic is how one might naturally react to an inappropriate assertion of the kind in question. In both cases, it is useful to compare a putatively undefined sentence to a non-homogeneous sentence which is clearly true or false.[34]

Let us consider in turn the three types of situations where a collective predication is, according to Generalised Homogeneity, undefined.

(108) Scenario 1: *Only a subgroup of the boys is staging the performance.*
  a. #The boys are performing Hamlet (together).[35,36]
  b. #The boys aren't performing Hamlet (together).

There is a clear contrast between (108b) and the corresponding *all*-sentence, which is uncontroversially true (indicating that its affirmative counterpart is false).

(109) Scenario 1: *Only a subgroup of the boys is staging the performance.*
  a. #All the boys are performing Hamlet (together).[37]
  b. Not all the boys aren't performing Hamlet (together).

This concludes the application of the first diagnostic, plain acceptability, and we can move on to the second, appropriate responses. We submit that the

34 All of the judgements reported in this section are mirrored in the author's native German with the particle *naja* functioning like *well*. It should be noted, however, that the existence of a particle with this function does not appear to be universal across languages.

35 To the extent that (108a) may ever seem acceptable in such a scenario, it is only in a situation where all the boys together are being credited as a team, while it is irrelevant how many members of the group actually contributed. This possibility has been observed for various collective predicates and is commonly referred to as *team credit* after Dowty 1987. Križ 2016 has argued that this is, in fact, just another instance of non-maximality.

36 For a brief discussion of the bracketed *together*, see the end of this section.

most natural response to both sentences in (108) is a *well*-answer followed by a correction. Replies with *yes* and *no* both appear possible to some extent, but are more marked.

(110) Scenario 1: *Only a subgroup of the boys is staging the performance.*
A: The boys are performing Hamlet (together).
B: Well, SOME of them are.
B′: ??Yes, some of them are.
B″: ??No, only some of them are.

Again we find a stark contrast with the corresponding *all*-sentence: here, both *no* and *well* are possible, while *yes* is out of the question. This indicates that the unacceptable homogeneous sentence really has a different status from the unacceptable non-homogeneous sentence.

(111) Scenario 1: *Only a subgroup of the boys is staging the performance.*
A: All the boys are performing Hamlet (together).
B: Well, SOME of them are.
B′: #Yes, some of them are.
B″: No, only some of them are.

Let us now move on to the second type of scenario.

(112) Scenario 2: *The boys and the girls together are performing Hamlet.*
a. #The boys are performing Hamlet (together).
b. #The boys aren't performing Hamlet (together).

What is interesting here is that for this scenario, the false counterpart of (112a) is not the corresponding *all*-sentence, which is likewise unacceptable:

(113) Scenario 2: *The boys and the girls together are performing Hamlet.*
a. #All the boys are performing Hamlet (together).
b. #Not all the boys are performing Hamlet (together).

Rather, the false counterpart is (114a), whose negation (114b) is uncontroversially true.

(114) Scenario 2: *The boys and the girls together are performing Hamlet.*
a. #The boys are performing Hamlet alone.
b. The boys aren't performing Hamlet alone.

In terms of preferred answers, we do not see much of a difference between the definite plural and the *all*-phrase. This is crucial for motivating our treatment of the interaction of quantifiers, homogeneity, and collective predication in section 3.4.

(115) Scenario 2: *The boys and the girls together are performing Hamlet.*
A: (All) the boys are performing Hamlet (together).
B: Well, they are doing it together with the girls.

B′: ??Yes, together with the girls.
B″: ??No, they are doing it together with the girls.

This is in contrast to the *alone*-variant of the sentence, which is plainly false.

(116) Scenario 2: *The boys and the girls together are performing Hamlet.*
A: (All) the boys are performing Hamlet alone.
B: Well, they are doing it together with the girls.
B′: #Yes, together with the girls.
B″: No, they are doing it together with the girls.

Finally, consider the third type of scenario where generalised homogeneity predicts undefinedness.

(117) Scenario 3: *Some of the boys together with some girls are giving the performance.*
a. #The boys are performing Hamlet (together).
b. #The boys aren't performing Hamlet (together).

Here the situation is slightly more complicated because it is not clear what variant of the sentence is plainly false in this scenario. To be sure, the *all*-variant of (117a) is unacceptable in this situation, but its negation also does not strike us as entirely natural.

(118) Scenario 3: *Some of the boys together with some girls are giving the performance.*
a. #All the boys are performing Hamlet (together).
b. ??Not all the boys are performing Hamlet (together).

One might take this to indicate that this is the same situation as in Scenario 2, where both sentences with plain *the* and those with *all* are undefined, in which case our analysis from section 3.4 is slightly off the mark, since it predicts (118a) to be false (and (118b) to be true). It would be hasty to draw this conclusion, however, before considering the second diagnostic.

(119) Scenario 3: *Some of the boys together with some girls are giving the performance.*
A: The boys are performing Hamlet (together).
B: Well, some of them are doing it together with some girls...
B′: ?#Yes, together with the girls.[38]
B″: ??No, some of them are doing it together with some girls.

The judgement is subtle, but it seems to us that a *no*-answer is more natural for the corresponding *all*-sentence in (120) than it is in (119).

(120) Scenario 3: *Some of the boys together with some girls are giving the performance.*
A: All the boys are performing Hamlet (together).
B: Well, some of them are doing it together with some girls...

38 We find it very difficult to imagine a context in which this *yes*-answer would be acceptable. This is natural if its limited acceptability is due to team credit, since it is hard to see how the boys, as a group, could every be credited with the performance when only some of the boys are participating in it together with somebody else.

B′: Yes, together with the girls.
B″: No, some of them are doing it together with some girls.

If this is correct, then what of the relative oddness of (118b), if it is indeed correctly perceived? We submit the following explanation for our intuition about this example. While (118b) is, indeed, literally true in the scenario in question, it triggers an indirect scalar implicature which is not true, and that is responsible for the feeling of degradedness. In particular, the scalar implicature is (121), which is undefined in the scenario in question.

(121) Some of the boys are performing Hamlet (together).

We would like to close with a note on the bracketed *together* that we have been carrying through this section. It was already noted by Dowty (1979) that speakers appear to disagree about the acceptability of *all* with collective action predicates such as *perform Hamlet*, a disagreement also reflected in the linguistic literature: while Dowty and Brisson (1998) report (122a) as acceptable — a judgement that aligns with our own — , Winter (2001) and Champollion (2010) fit it degraded. The variant (122b) with *together*, however, is uncontroversially acceptable.

(122) a. %All the boys are performing Hamlet.
b. All the boys are performing Hamlet together.

It is not entirely clear to us what, if any, is the difference between these two versions, for speakers who accept both. We have discovered one very specific type of scenario in which a difference appears to exist.

(123) Scenario: *The boys split into two groups, each group is performing Hamlet separately.*
a. #All the boys are performing Hamlet.
b. #All the boys are performing Hamlet together.

There is some indication that while (123a) is undefined in this situation, (123b) is actually false:

(124) Scenario: *The boys split into two groups, each group is performing Hamlet separately.*
A: All the boys are performing Hamlet.
B: Well, they're doing it in two groups.
B′: ??Yes, in two groups.
B″: ??No, they're split into two groups.

(125) Scenario: *The boys split into two groups, each group is performing Hamlet separately.*
A: All the boys are performing Hamlet together.
B: Well, they're doing it in two groups.
B′: #Yes, in two groups.
B″: No, they're split into two groups.