Collective Nouns and the Distribution Problem

David Nicolas[1†] and Jonathan D. Payton[2†]

[1]Institut Nicod, ENS-EHESS-CNRS, 29 rue d'Ulm, Paris, 75005, France.
[2]Department of Philosophy, Bilkent University, 1598 Caddesi, Ankara, TR-06800, Çankaya, Türkiye.

Contributing authors: david.nicolas@ens.psl.eu; jonathanpayton@bilkent.edu.tr;
[†]The two authors contributed equally to this work.

Abstract

Intuitively, collective nouns are pseudo-singular: a collection of things (a pair of people, a flock of birds, etc.) just *is* the things that make 'it' up. But certain facts about natural language seem to count against this view. In short, distributive predicates and numerals interact with collective nouns in ways that they seemingly shouldn't if those nouns are pseudo-singular. We call this set of issues 'the distribution problem'. To solve it, we propose a modification to cover-based semantics. On this semantics, the interpretation of distributive predicates and numerals depends on a cover, where the choice of cover is strongly semantically constrained by the noun with which they interact.

Keywords: collective, cover, distributive, number, plural, pseudo-singular

1 Introduction

Collective nouns stand for collections of objects. They include numerically specific nouns like 'pair', 'trio', and 'quartet', as well as numerically non-specific ones like 'flock', 'group', and 'team'. On some views, a collection is an individual, distinct from the individuals which make it up: a pair of people, for instance, is an individual which is composed, or constituted, by two people. We favor an opposing view according to which a collection just *is* the individuals which make it up: the pair *is* the people.[1]

1. See also Black (1971), Korman (2015, 139–150), Oliver and Smiley (2016, 138–140, 303–307), and Horden and López de Sa (2021).

Collective noun phrases are *pseudo-singular* (Oliver and Smiley, 2016, 306): despite their grammatical singularity, they stand for pluralities rather than individuals.

This view faces a problem. Consider:

(1) The pairs (of people) lifted the piano.

(2) The people lifted the piano.

As is well-known, each of these sentences admits a variety of readings. When (1) is read *collectively*, it's true just in case the pairs lifted the piano *all together*; likewise, (2) is true on its collective reading just in case the people lifted the piano all together. Our view predicts that these readings are truth-conditionally equivalent, and that prediction appears to be borne out: the only way for the pairs to do something together is for all the people to do it together.

But now consider *pure distributive* readings. When (1) is read purely distributively, it's true just in case each *pair* lifted a piano on their own, but when (2) is read that way, it's true just in case each *person* lifted a piano on their own. In each case, the predicate is required to distribute over different things, leading to different truth-conditions.[2]

Both sentences likewise admit *intermediate distributive* readings. When (1) is read this way, the pairs are divided into groups of *pairs* and each group is said to have lifted the piano on their own. For instance, suppose that there are three pairs: Alice and Beth, Cathy and Diane, and Eve and Fiona. Then there's a reading of (1) on which it's true if Alice and Beth cooperated with Cathy and Diane to lift the piano, and also did so with Eve and Fiona. But there's no reading of the sentence on which it's true if Alice and Cathy cooperated with Diane and Eve to lift the piano, and also did so with Beth and Fiona. By contrast, when (2) is read intermediately distributively, the people are divided into groups of *people*, and each of *these* groups is said to have lifted the piano on their own. So, there *is* a reading of (2) on which it's true if Alice and Cathy cooperated with Diane and Eve to lift the piano, and also did so with Beth and Fiona.

As stressed by Landman (2020, sec. 4.7) — who focuses on ordinary plurals rather than collective nouns — these issues are intimately linked to counting. In particular, in a pure distributive reading, the predicate distributes over the very same 'things' that are counted.

(3) The pairs (of people) are two.

The count expressed by (3) is one in which each pair (and not each person) is counted 'as one'. There's no reading of the sentence on which it says that the people are two.

None of this would be a problem if a pair were an individual distinct from the people making it up. For then, the pairs would be distinct from the people, and so to say that each of the former (or some groups of them) lifted the piano wouldn't be

2. An anonymous referee objects that (1) *does* have a reading on which the predicate distributes down to each person, and indeed that *any* plural predicate which is predicated of some collections can be read so as to distribute down to the individuals making up those collections. Linguistic intuitions are, of course, difficult to adjudicate. But consider a variant of (1):

(i) The pair lifted the piano.

Intuitively, this sentence is true just in case the pair *collectively* lifted the piano; it has no reading on which it's true just in case each *member* of the pair lifted the piano (Ritchie, 2017, 466). But if the predicate could distribute all the way down to individual people in (1), it would be mysterious that it can't do so in this variant.

to say that the latter (or some groups of them) did, and to say that the former are two wouldn't be to say that the latter are. By contrast, on our view the pairs *are* the people; 'the people' and 'the pairs' are *co-denoting*. But then, if each of the pairs lifted the piano, doesn't it follow that each of the people did, and *vice-versa*? And if the pairs are two, doesn't it follow that the people are two, and *vice-versa*? Inspired by Landman, we'll refer to this collection of issues as *the distribution problem*.

A solution is to adopt a semantics which includes *covers* (Gillon, 1992; Schwarzschild, 1996; Champollion, 2017). Informally, a cover is a way of 'carving up' a plurality into some of its subpluralities. The basic thought behind cover-based semantics is that whether some things satisfy the application-conditions of a predicate depends not only on what those things are like but also on how they're carved up.

Keeping things informal — and focusing for the moment on pure distributive readings — suppose we say that a predicate, interpreted distributively, is true of some things just in case it's true of each subplurality of a cover of those things. Then, we might say, (1) is interpreted with respect to a cover which carves the plurality of people who make up the pairs into each of those pairs; thus, 'lifted the piano' distributes over the pairs and the sentence says that each of *them* lifted the piano. By contrast, (2) is interpreted with respect to a cover which carves the plurality of people who make up the pairs into each of those people; thus, 'lifted the piano' distributes over the people and the sentence says that each of *them* lifted the piano. The pure distributive readings of the sentences are different, despite the fact that 'the pairs of people' and 'the people' are co-denoting.

This solution is satisfying only if we have an account of how the relevant covers get selected. What makes it the case that (1) and (2) are interpreted with respect to the covers which yield the right truth-conditions? The common answer is that cover-selection is a matter of pragmatics (Gillon (1990, 482-483) and Schwarzschild (1996, 92-98)). Drawing on Kennedy (2007), Grimau (2021b, 179) suggests that, all other things being equal, speakers aim to 'maximize the contribution of the conventional meaning of the elements of a sentence to the contribution of its truth-conditions thereby avoiding indeterminacy'. So we might say that when the speaker has available to her two co-denoting descriptions — 'the pairs' and 'the people', say — she should favor the one whose constituent noun phrase picks out the kinds of things over which she intends the predicate to distribute.

This appeal to pragmatics explains why (1) and (2) have certain readings, namely their pure distributive readings. But it doesn't explain why they *lack* certain readings. In particular, there's no reading of (1) on which 'lifted the piano' distributes over individual people rather than pairs. But if cover-selection were entirely a matter of pragmatics, then such a reading would be merely unusual; it wouldn't be ruled out.

In this article, we propose an alternative account. Cover-selection isn't an entirely pragmatic matter; it's *semantically* constrained. The idea is that certain elements of natural language — count nouns and plural definites for instance — are associated with special covers, and that some predicates are constrained so as to be interpreted with respect to *these* covers. This account deals with all facets of the distribution problem and does so in a unified way.

2 The framework

In what follows we'll be concerned with both *analysis* and *compositional semantics*. The former requires a formal language in which we can precisely state the truth-conditions of the sentences we're interested in. The latter requires an account of how the truth-conditions of such sentences are derived from the meanings of their components. In both respects, our approach is non-standard.

2.1 The language $\mathcal{CPL}$

Our formal language is an extension of the language $\mathcal{CPL}$, developed in Payton (nd). $\mathcal{CPL}$ is a first-order plural language[3] containing both singular terms ('a', 'b', 'c', ...) and plural ones ('aa', 'bb', 'cc', ...). While a singular term can only denote one individual, a plural term can denote more than one. Likewise, it contains both singular variables ('x', 'y', 'z', ...) and plural ones ('xx', 'yy', 'zz', ...). While a singular variable can only take a singular term as substitution, a plural variable can take either a singular term or a plural one. Existential and universal quantifiers ('$\exists$', '$\forall$') can bind both singular and plural variables.

Predicates ('F', 'G', ...) can take either singular or plural terms at any argument-place. For any predicate F we can define a distributive variant, F^D. (The details will be given below.)

$\mathcal{CPL}$ contains two logical predicates familiar from other plural languages. The first of these is the inclusion predicate ('$\leq$'), which denotes an inclusion relation between entities in the domain. This relation determines how certain individuals or pluralities, so to speak, 'build' a plurality of things. For instance, a pair of people *includes* each of those people. Inclusion is reflexive, anti-symmetric, and transitive. The second is the identity predicate ('$=$'). Identity is understood as mutual inclusion. This means that pluralities are individuated extensionally: aa and bb are identical just in case they include all and only the same things.

$\mathcal{CPL}$ differs from other plural languages in its use of *covers*, which we'll denote using the Latin letters 'i', 'j', 'k', etc. Formally, a cover is a multi-valued partial function whose input is a plurality aa and whose outputs are one or more subpluralities of aa.[4]

$$\begin{array}{ccl} & & \longrightarrow bb_1 \\ aa & & \qquad\vdots \\ & & \longrightarrow bb_n \end{array}$$

We define these functions as follows:

Covers
i covers aa just in case:
(i) for any bb, if bb are a value of $i(aa)$, then bb are included in aa and
(ii) for any x, if x is included in aa, then there are some bb such that bb are a value of $i(aa)$ and x is included in bb.

3. For more on plural languages, see e.g., Rayo (2002), Yi (2005, 2006), McKay (2006), and Oliver and Smiley (2016).
4. Nicolas and Payton (nd) treat a cover as a function from a plurality and an index to a subplurality. Here, we use the multi-valued function approach for simplicity. Compare also the more traditional set-theoretic approach as found in, e.g., Gillon (1987) and Schwarzschild (1996).

When $i(aa)$ is defined, then it either divides aa into subpluralities, or it 'does nothing', returning aa themselves. Of course, if a is an individual then the only thing included in a is itself, and so a cover can only map a to itself.

Crucially, an output of $i(aa)$, bb, can themselves be a plurality for which $i(bb)$ is defined. That is, a cover can carve aa into certain subpluralities and then carve each of *those* into further subpluralities. For example, a single cover might carve six people into three pairs and carve each pair into two individual people.

$$\begin{array}{lllll} & & & \longrightarrow & a \\ & \longrightarrow & ab & \longrightarrow & b \\ & & & \longrightarrow & c \\ abcdef & \longrightarrow & cd & \longrightarrow & d \\ & & & \longrightarrow & e \\ & \longrightarrow & ef & \longrightarrow & f \end{array}$$

This will be important in Sections 2.3 and 3.1.

In $\mathcal{CPL}$, a term can only occur as an argument of a predicate if it's indexed to a cover: 'Faa' isn't a formula, but 'Faa^i' is. Different terms, and different occurrences of the same term, may be indexed to different covers in the same formula. Likewise for variables: '$\exists xx(Fxx)$' isn't a formula, but '$\exists xx(Fxx^i)$' is.

The application-conditions of a predicate are allowed to be sensitive, not merely to the denotations of the terms (or the values of the variables) which fill its argument-places, but also to the covers to which those terms (and variables) are indexed. Thus, whether 'F' is true of aa can depend on our choice of cover. This isn't true for all predicates. The inclusion predicate, for instance, is cover-insensitive: '$aa^i \leq bb^j$' is true for all choices of i and j if it's true for any of them. (Likewise, of course, for identity.) However, $\mathcal{CPL}$ contains a cover-sensitive analogue of inclusion, c-inclusion ('$\precsim$').

c-inclusion
'$aa^i \precsim bb^j$' is true iff aa are a value of $j(bb)$.

The claim that aa are c-included in bb, with respect to i and j respectively, is an object-language variant of the meta-language claim that aa are a value of $j(bb)$, and so its truth depends on our choice of cover. (Note that '$\precsim$' is only cover-sensitive at its second argument-place: whether aa are c-included in bb depends on how bb are carved up, but not on how aa are.)

In plural logic, distributive interpretations of predicates are traditionally understood in terms of inclusion: a distributive predicate distributes over every individual included in a plurality. In $\mathcal{CPL}$, distribution is understood in terms of c-inclusion. Using the notation 'F^D' as an abbreviation for the distributive reading of 'F', 'F^D' is true of aa with respect to i just in case 'F' is true of every value of $i(aa)$, with respect to that same cover:

Distributive interpretations
$F^D(aa^i) \equiv_{\text{df}} \forall xx(xx^i \precsim aa^i \rightarrow Fxx^i)$

Thus, whether 'F^D' is true of aa depends on how aa have been carved up.

Finally, we extend $\mathcal{CPL}$ by allowing ourselves, not merely to refer to covers in the object language, but to quantify over them.

2.2 Compositional semantics

So, in our semantics, the truth-conditions of sentences will be stated in a first-order plural language in which terms and variables are indexed to covers, and these truth-conditions will be sensitive to those covers. How do we ensure that sentences are assigned such truth-conditions?

Our approach is a modification of the familiar type-theoretic approach exemplified by Heim and Kratzer (1998). Our basic semantic types are pluralities (type e), the truth-values 1 and 0 (type t), and covers (type c). Other semantic types are constructed recursively from these; for instance, the extensions of one-place verb phrases are functions from pluralities to truth-values (type <e,t>).

Ours is an intensional semantics. The denotation of an item in the lexicon is an *intension*, i.e. a function from one or more indices to some entity in the type-theoretic hierarchy. A complete intensional semantics would include worlds and times (and perhaps much else) as indices. For our purposes the only indices we'll need are covers, but the story of how these figure into the semantics is, in all essential respects, the same as for worlds, times, and any other indices we may eventually need. Thus, the denotation of an item in the lexicon is always a function which takes one or more covers as inputs and outputs some other entity in the type-theoretic hierarchy. We use '$[\![\alpha]\!]$' to refer to the denotation of the lexical item α. We use superscripts to indicate that one or more argument-places for covers have been saturated, i.e. that one or more covers have been input into the function $[\![\alpha]\!]$: $[\![\alpha]\!]^i$ is the result of saturating one of these argument-places with the cover i, $[\![\alpha]\!]^{i,j}$ is the result of saturating one of these argument-places with the cover i and another with the cover j, etc. When we want to refer to ordered sequences of covers, we use capital letters 'I', 'J', 'K', etc. $[\![\alpha]\!]^I$ is the result of saturating one or more of these argument-places, in order, with the covers in the sequence I.

The denotation of a term is an intension of type <c, e>, a function from a cover to a plurality:

$$[\![\text{aa}]\!] = \lambda i[aa]$$

Note that this is a constant function: the extension of a term doesn't shift with the choice of cover.

The denotation of a verb phrase is an intension of type <c,<e,t>>, a function from a cover to a property. For example, the denotation of the one-place verb phrase 'is a person' is:

$$[\![\text{is a person}]\!] = \lambda i \lambda xx[\textsc{Person}(xx^i)]$$

In keeping with $\mathcal{CPL}$, the occurrence of 'xx' which is bound by 'λxx' is indexed to the cover-variable 'i'. By supplying a cover, we get the *extension* of this verb phrase with respect to i: the property of being a person with respect to i.[5]

5. Notice, we don't say that the extension of 'is a person' is '$\lambda xx^i.[\textsc{Person}(xx^i)]$', with the first instance of 'xx' indexed to a cover. That's because the extension of 'is a person' is a function of type <e,t>: it operates on pluralities, not on pluralities-indexed-to-covers.

$[\![\text{is a person}]\!]^i = \lambda xx[\textsc{Person}(xx^i)]$

Saturating 'xx' with a constant yields a sentence of $\mathcal{CPL}$: '$\textsc{Person}(aa^i)$'. Since 'is a person' is semantically singular, the sentence can only be true if aa is a single individual. The denotation of the plural verb phrase 'are people', with respect to a cover i, is obtained by application of a distribution operator. We'll see in a moment how the compositional story goes in more detail, but the end result can be stated simply in $\mathcal{CPL}$:

$[\![\text{are people}]\!] = \lambda i \lambda xx[\textsc{Person}^D(xx^i)]$

Now, saturating 'xx' with a constant yields the sentence, '$\textsc{Person}^D(aa^i)$'. Since 'are people' is semantically plural, there's no longer a requirement that aa be a single individual. But the sentence is only true with respect to a cover i which carves aa into individual people.

In general, the extension of a verb phrase shifts with our choice of cover(s); in general, $[\![\text{VP}]\!]^I \neq [\![\text{VP}]\!]^J$. This means that whether a verb phrase is true of some things can depend on how those things are 'carved up'. Notice, however, that the application-conditions of a verb phrase can sometimes be *constant* in the way that terms are: sometimes, if $[\![\text{VP}]\!]^I$ is true of aa, then $[\![\text{VP}]\!]^J$ is true of aa, for any I and J. The verb phrase 'is a person' is like this: since every person is an individual, and a cover can only map an individual to itself, there's no way for an individual to count as a person with respect to one cover but not another; whether an individual is a person doesn't (and *can't*) depend on how that individual is 'carved up'.

To derive the truth-conditions of sentences, we use three compositional rules. The first is a variant of Functional Application (Heim and Kratzer, 1998, 44):

Functional application (FA)
If α is a branching node and β and γ are its daughters, and $[\![\beta]\!]^I$ is a function whose domain includes $[\![\gamma]\!]^J$, then $[\![\alpha]\!]^{I,J} = [\![\beta]\!]^I([\![\gamma]\!]^J)$.

As a simple example, we derive the truth-conditions of 'Alice is a person' as follows:

$$\begin{aligned}[\![\text{Alice is a person}]\!]^{i,j} &= [\![\text{is a person}]\!]^i([\![\text{Alice}]\!]^j)\\ &= \lambda xx[\textsc{Person}(xx^i)](a)\\ &= \textsc{Person}(a^i)\end{aligned}$$

Of course, since the denotation of a singular term is a constant function to an individual, our choice of indices plays no role, here. $[\![\text{Alice is a person}]\!]^{i,j} = [\![\text{Alice is a person}]\!]^{i,k}$, for any i, j and k. Indices will, however, play a role when a predicate is applied to a plural term.

Our second rule is a variant of Intensional Functional Application (Heim and Kratzer, 1998, 308):

Intensional functional application (IFA)
If α is a branching node, and β and γ are its daughters, then for any I, if $[\![\beta]\!]^I$ is a function whose domain includes $\lambda j([\![\gamma]\!]^{j,K})$, then $[\![\alpha]\!]^{I,K} = [\![\beta]\!]^I(\lambda j([\![\gamma]\!]^{j,K})$.

Here, $\lambda j([\![\gamma]\!]^{j,K}])$ is a function from some cover j to the denotation of γ with respect to that j and some arbitrary sequence of covers, K, which have saturated other argument-places of $[\![\gamma]\!]$. In the limit case, where no other argument-places of $[\![\gamma]\!]$ have been saturated, reference to K is inert: $\lambda j(\llbracket\gamma\rrbracket^{j,K}) = \lambda j(\llbracket\gamma\rrbracket^{j})$, and likewise $\llbracket\alpha\rrbracket^{I,K} = \llbracket\alpha\rrbracket^{I}$. IFA allows an item in the lexicon to operate, not only on the extensions of other items (relative to some cover), but on their intensions, as well.[6]

IFA is used to generate the plural variants of singular noun phrases. We introduce a distribution operator which, effectively, converts a predicate into its distributive variant:[7]

$$\begin{aligned}\llbracket\text{DIST}\rrbracket &= \lambda i\lambda\phi\lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \phi(i)(yy^i))]\\ &= \lambda i\lambda\phi\lambda xx[\phi^D(xx^i)]\end{aligned}$$

Now, for example, if we let the denotation of 'is a person' be as before:

$$\llbracket\text{is a person}\rrbracket = \lambda i\lambda xx[\textsc{Person}(xx^i)]$$

then we can derive the denotation of 'are people', with respect to a cover i, as follows:

$$\begin{aligned}\llbracket\text{are people}\rrbracket^i &= \llbracket\text{DIST}\rrbracket^i(\llbracket\text{is a person}\rrbracket)\\ &= \lambda\phi\lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \phi(i)(xx^i))](\lambda j\lambda zz[\textsc{Person}(zz^j)])\\ &= \lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \lambda j\lambda zz[\textsc{Person}(zz^j)](i)(xx^i))]\\ &= \lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \textsc{Person}(xx^i))]\\ &= \lambda xx[\textsc{Person}^D(xx^i)]\end{aligned}$$

Our final rule is a variant of Predicate Modification (Heim and Kratzer, 1998, 65):

Predicate modification (PM)
If α is a branching node whose daughters are a head β and a modifier γ, both of type ¡e,t¿, then $\llbracket\alpha\rrbracket^{I,J} = \lambda xx(\llbracket\beta\rrbracket^I(xx)$ & $\llbracket\gamma\rrbracket^J(xx))$.

For instance, suppose our $\llbracket\beta\rrbracket^I$ and $\llbracket\gamma\rrbracket^J$ are:

$$\begin{aligned}\llbracket\text{is a person}\rrbracket^i &= \lambda xx[\textsc{Person}(xx^i)]\\ \llbracket\text{is tall}\rrbracket^j &= \lambda xx[\textsc{Tall}(xx^j)]\end{aligned}$$

Then, the result of combining these by PM is:

$$\llbracket\text{is a tall person}\rrbracket^{i,j} = \lambda xx[\textsc{Person}(xx^i) \;\&\; \textsc{Tall}(xx^j)]$$

Applying this to $\llbracket\text{Alice}\rrbracket^k$ by FA, we get:

$$\begin{aligned}\llbracket\text{Alice is a tall person}\rrbracket^{i,j,k} &= \lambda xx[\textsc{Person}(xx^i) \;\&\; \textsc{Tall}(xx^j)](a)\\ &= \textsc{Person}(a^i) \;\&\; \textsc{Tall}(a^j)\end{aligned}$$

Here, as before, the cover to which $\llbracket\text{Alice}\rrbracket$ is originally indexed gets 'swallowed up': for any l, $\llbracket\text{Alice is a tall person}\rrbracket^{i,j,k} = \llbracket\text{Alice is a tall person}\rrbracket^{i,j,l}$.

2.3 Collective nouns

It should be agreed on all sides — even by our opponents, who take collective nouns to stand for special kinds of individuals — that *bare* collective nouns, on their own, don't stand *either* for individuals *or* for pluralities. They need to be supplemented

6. Note that, because j hasn't been saturated, the result of IFA is $\llbracket\alpha\rrbracket^{I,K}$ (or $\llbracket\alpha\rrbracket^{I}$), not $\llbracket\alpha\rrbracket^{I,j,K}$ (or $\llbracket\alpha\rrbracket^{I,j}$).
7. Compare Schwarzschild (1996, 68–71).

by other count nouns. The denotation of 'pair', for instance, should take us from the denotation of 'people' to the property *being a pair of people*, from the denotation of 'birds' to the property *being a pair of birds*, and so on. The question is how we should represent this in our framework. What conditions do some things have to meet, to be a pair of Fs, and how can we generate those conditions compositionally?

You might think that some things are a pair of Fs just in case they're two individual Fs — i.e. just in case they include some x and some y, x is distinct from y, and nothing else is included in them. But that seemingly natural thought doesn't fit with the view that collective nouns are pseudo-singular. For, bare collective nouns can modify other collective nouns: just as we can have a pair of individual people, we can have a pair of pairs, a pair of trios, and so on. So it's not true in general that a pair of Fs is a pair of *individuals* which are Fs.

$\mathcal{CPL}$ provides the resources we need. Recall that in $\mathcal{CPL}$ the application-conditions of a predicate can be sensitive to our choice of cover. Assuming that collective nouns are pseudo-singular, we get a better account of their meaning by treating their meanings as cover-sensitive. Whether some things are a pair of Fs, for instance, depends on how they're carved up. Specifically, aa are a pair of Fs with respect to i just in case i carves aa into two Fs. Here, 'F' can be a plural predicate satisfied by more than one individual. So, for instance, four people count as two pairs, with respect to a cover i, if i carves those four people into two pairs.

Putting some meat on these bones — and continuing to focus on 'pair' for the moment — we first define numerically specific quantifier phrases as follows (where the cover i is inert):

Numerically specific quantifier phrases

(i) $\exists_1 xxF(xx^i) \equiv_{\text{df}} \exists xx(F(xx^i) \mathbin{\&} \forall yy(F(yy^i) \rightarrow yy^i = xx^i))$

(ii) For any $n \geq 2$,
$\exists_n xxF(xx^i) \equiv_{\text{df}} \exists xx_1 \ldots \exists xx_n(F(xx_1^i) \ldots \& \ldots F(xx_n^i) \mathbin{\&} xx_1^i \neq xx_2^i \ldots \& \ldots xx_{n-1}^i \neq xx_n^i \mathbin{\&} \forall yy(F(yy^i) \rightarrow yy^i = xx_1^i \ldots \vee \ldots yy^i = xx_n^i))$

These phrases allow us to say that there are exactly n pluralities satisfying a predicate 'F'. Substituting '$\lambda xx[xx^i \precsim aa^i]$' for '$F$' allows us to say that $i(aa)$ has exactly n values, and hence to define predicates which correspond to numerically specific collective nouns. For instance, we can define a predicate 'Pair' as follows:

$$\text{P}\textsc{air}(aa^i) \equiv_{\text{df}} \exists_2 xx(xx^i \precsim aa^i)$$

That is, '$\text{P}\textsc{air}(aa^i)$' is true just in case $i(aa)$ has exactly two values.

Next, we give the denotation of the English word 'pair' as a function from a property ϕ to the property of being ϕ and being carved into two subpluralities:

$$[\![\text{pair}]\!] = \lambda i \lambda \phi \lambda xx[\text{P}\textsc{air}(xx^i) \mathbin{\&} \phi(i)(xx^i)]$$

Finally, we can derive the denotation of 'pair of people' with respect to a cover i, from $[\![\text{pair}]\!]^i$ and $[\![\text{people}]\!]$ by our rule IFA:

$$\begin{aligned}[\![\text{pair of people}]\!]^i &= [\![\text{pair}]\!]^i([\![\text{people}]\!]) \\ &= \lambda\phi\lambda xx[\textsc{Pair}(xx^i) \mathbin{\&} \phi(i)(xx^i)](\lambda j\lambda yy[\textsc{Person}^D(yy^j)]) \\ &= \lambda xx[\textsc{Pair}(xx^i) \mathbin{\&} \textsc{Person}^D(xx^i)]^8\end{aligned}$$

aa are a pair of people with respect to i, then, just in case $i(aa)$ has exactly two values and '$\textsc{Person}^D$' distributes over those values — that is, just in case i carves aa into two people. For instance:

$$\begin{aligned}[\![\text{Alice and Beth are a pair of people}]\!]^{i,j} &= [\![\text{are a pair of people}]\!]^i([\![\text{Alice and Beth}]\!]^j) \\ &= \lambda xx[\textsc{Pair}(xx^i) \mathbin{\&} \textsc{Person}^D(xx^i)](ab) \\ &= \textsc{Pair}(ab^i) \mathbin{\&} \textsc{Person}^D(ab^i)\end{aligned}$$

This formula comes out true only with respect to a cover whose only values are Alice and Beth:

$$ab \begin{array}{l}\longrightarrow a \\ \longrightarrow b\end{array}$$

The plural form, 'pairs of people', is obtained by application of the distributive operator. To make this (and what follows) more reader-friendly, we first introduce the following abbreviation:

$$\textsc{PairPeo}(aa^i) \equiv_{\text{df}} \textsc{Pair}(aa^i) \mathbin{\&} \textsc{Person}^D(aa^i)$$

So, we can give the denotation of 'pair of people' more simply as:

$$[\![\text{pair of people}]\!] = \lambda i\lambda xx[\textsc{PairPeo}(xx^i)]$$

Applying the distributive operator, we get:

$$\begin{aligned}[\![\text{pairs of people}]\!]^i &= [\![\text{DIST}]\!]^i([\![\text{pair of people}]\!]) \\ &= \lambda\phi\lambda xx[\forall yy(yy^i \precsim xx^i \to \phi(i)(yy^i))](\lambda j\lambda zz[\textsc{PairPeo}(zz^j)]) \\ &= \lambda xx[\forall yy(yy^i \precsim xx^i \to \textsc{PairPeo}(yy^i))] \\ &= \lambda xx[\textsc{PairPeo}^D(xx^i)]^9\end{aligned}$$

aa are some pairs of people with respect to i, then, just in case each value of $i(aa)$, bb, is such that $i(bb)$ has two values and '$\textsc{Person}^D$' distributes over *those* — that is, just in case i carves aa into pairs of people. For instance, suppose that the demonstrative 'they' denotes Alice, Beth, Cathy, Diane, Eve, and Fiona. Then:

---

8. Dropping all our abbreviations:

$$\begin{aligned}[\![\text{pair of people}]\!]^i =& \lambda xx[\exists yy\exists zz(yy^i \precsim xx^i \mathbin{\&} zz^i \precsim xx^i \mathbin{\&} \forall vv(vv^i \precsim xx^i \to (vv^i = yy^i \lor vv^i = zz^i))) \mathbin{\&} \\ & \forall uu(uu^i \precsim xx^i \to \textsc{Person}(uu^i))]\end{aligned}$$

9. Dropping all our abbreviations:

$$\begin{aligned}[\![\text{pairs of people}]\!]^i &= \lambda xx[\forall yy(yy^i \precsim xx^i \to (\textsc{Pair}(yy^i) \mathbin{\&} \textsc{Person}^D(yy^i)))] \\ &= \lambda xx[\forall yy(yy^i \precsim xx^i \to \exists zz\exists vv(zz^i \precsim yy^i \mathbin{\&} vv^i \precsim yy^i \mathbin{\&} \forall uu(uu^i \precsim yy^i \to \\ & \quad (uu^i = zz^i \lor uu^i = vv^i))) \mathbin{\&} \forall ww(ww^i \precsim yy^i \precsim \textsc{Person}(ww^i)))]\end{aligned}$$

$$
\begin{aligned}
[\![\text{They are pairs of people}]\!]^{i,j} &= [\![\text{are pairs of people}]\!]^{i}([\![\text{they}]\!]^{j}) \\
&= \lambda xx[\textsc{PairPeo}^{D}(xx^{i})](abcdef) \\
&= \textsc{PairPeo}^{D}(abcdef^{i})
\end{aligned}
$$

This formula comes out true only with respect to a cover like the one described in Section 2.1, which first carves the six women into three pairs and then carves each pair into individual people:

$$
\begin{array}{lll}
 & & \longrightarrow a \\
 & \longrightarrow ab & \longrightarrow b \\
 & & \longrightarrow c \\
abcdef & \longrightarrow cd & \longrightarrow d \\
 & & \longrightarrow e \\
 & \longrightarrow ef & \longrightarrow f
\end{array}
$$

Compare this to a cover which instead carves the women into two trios:

$$
\begin{array}{lll}
 & & \longrightarrow a \\
 & \longrightarrow abc & \longrightarrow b \\
 & & \longrightarrow c \\
abcdef & & \longrightarrow d \\
 & \longrightarrow def & \longrightarrow e \\
 & & \longrightarrow f
\end{array}
$$

The formula is *false* with respect to this cover, because each value of $i(abcdef)$ has *three* values, and so fails to satisfy ‘Pair’. So, unlike $[\![\text{is a person}]\!]$, $[\![\text{are pairs (of people)}]\!]$ is *non-constant*: whether some people count as pairs depends on how they’re carved up.

In presenting our approach to collective nouns we’ve focused on ‘pair’, but the approach generalizes. The extension to other numerically specific collectiven nouns is perhaps the most straightforward. For instance, we can define a predicate ‘Trio’,

$$\textsc{Trio}(aa^{i}) \equiv_{\mathrm{df}} \exists_{3} xx(xx^{i} \precsim aa^{i})$$

and then give the denotation of the English word ‘trio’ as follows:

$$[\![\text{trio}]\!] = \lambda i \lambda \phi \lambda xx[\textsc{Trio}(xx^{i}) \ \&\ \phi(i)(xx^{i})]$$

The extension to numerically non-specific collective nouns like ‘flock’ is less straightforward, if only because we don’t know how to define a predicate ‘Flock’—that is, we don’t know how to specify, in informative terms, the conditions under which some things constitute a flock. But that doesn’t obstruct the basic idea. The denotation of ‘flock’ can be given as follows:

$$[\![\text{flock}]\!] = \lambda i \lambda \phi \lambda xx[\textsc{Flock}(xx^{i}) \ \&\ \phi(i)(xx^{i})]$$

and so the denotation of, for instance, ‘flock of birds’ with respect to a cover i is:

$$
\begin{aligned}
[\![\text{flock of birds}]\!]^{i} &= [\![\text{flock}]\!]^{i}([\![\text{birds}]\!]) \\
&= \lambda \phi \lambda xx[\textsc{Flock}(xx^{i}) \ \&\ \phi(i)(xx^{i})](\lambda j \lambda yy[\textsc{Bird}^{D}(yy^{j})]) \\
&= \lambda xx[\textsc{Flock}(xx^{i}) \ \&\ \textsc{Bird}^{D}(xx^{i})]
\end{aligned}
$$

aa are a flock of birds with respect to i just in case the values of $i(aa)$ are birds — i.e. i carves aa into individual birds — and those birds are (collectively) a flock. Likewise, letting '$\textsc{FlockBir}(aa^i)$' serve as an abbreviation for '$\textsc{Flock}(aa^i)$ & $\textsc{Bird}^D(aa^i)$',

$$\begin{aligned}
[\![\text{flocks of birds}]\!]^i &= [\![\text{DIST}]\!]^i([\![\text{flock of birds}]\!]) \\
&= \lambda\phi\lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \phi(i)(yy^i))](\lambda j\lambda zz[\textsc{FlockBir}(zz^j)]) \\
&= \lambda xx[\forall yy(yy^i \precsim xx^i \rightarrow \textsc{FlockBir}(yy^i))]
\end{aligned}$$

aa are some flocks of birds with respect to i just in case each value of $i(aa)$, bb, are (collectively) a flock, and each value of $i(bb)$ is a bird:

$$\begin{array}{ccccc}
 & & & \longrightarrow & b_1 \\
 & \longrightarrow & bb_1 & & \vdots \\
 & & & \longrightarrow & b_m \\
aa & & \vdots & & \\
 & & & \longrightarrow & b_{m+1} \\
 & \longrightarrow & bb_n & & \vdots \\
 & & & \longrightarrow & b_l
\end{array}$$

2.4 Plural definites

Finally, we need to give a semantics for plural definites like 'the people' and 'the pairs'.

In general, 'the Fs' denotes the largest plurality of Fs.[10] This is often taken to mean that 'the Fs' denotes that plurality which includes all the Fs and nothing else.[11] This works fine if we only consider plural definites involving genuinely singular noun phrases; for instance, 'the people' denotes that plurality which includes all the people and nothing else. Problems arise when we consider pseudo-singular nouns. On the standard approach, 'the pairs' denotes that plurality which includes all the pairs of people and nothing else. But given that a pair just *is* two people, there can be no such plurality. Any plurality which includes even one pair of people must include each of the individual people making it up, and so must include some things which aren't pairs. The standard approach leaves 'the pairs' without a denotation.

We want 'the pairs (of people)' to be capable of denoting a plurality in the domain. Moreover, as defenders of pseudo-singularity, we want it to be capable of denoting the *same* plurality as 'the people'. To allow for this, we say that 'the Fs' denotes the largest plurality which can be carved into Fs.[12]

Plural definites schema
'The Fs' denotes the unique xx, if there are such, such that:
$\exists i(\phi(xx^i))$ & $\forall yy\forall j(\phi(yy^j) \rightarrow yy^j \leq xx^j)$

On this approach, 'the pairs' not only gets a denotation, it co-denotes with 'the people'. Suppose there are only four people: Alice, Beth, Cathy, and Diane. 'The people'

10. Or rather, the largest plurality of contextually salient Fs. We'll suppress this qualification in what follows.
11. See, e.g., McKay (2006, 164–165), Sider (2014, 215–216), and Oliver and Smiley (2016, 95–96, 130–136).
12. See also Payton (2021).

denotes the four women together, since they're the largest plurality which can be carved into individual people. Moreover, 'the pairs (of people)' *also* denotes the four women together, since they're the largest plurality which can be carved into pairs of people — i.e. the largest plurality which can be carved into subpluralities each of which is carved into two people.

To derive these denotation-conditions compositionally, we assign the following denotation to the definite article.

$$\begin{aligned}[\![\text{the}]\!] &= \lambda i\lambda\phi[\iota xx(\exists j(\phi(j)(xx)) \,\&\, \forall yy\forall k(\phi(k)(yy) \to yy^k \leq xx^k))] \\ &= \lambda i\lambda\phi[\iota xx(\exists j(\phi(xx^j)) \,\&\, \forall yy\forall k(\phi(yy^k) \to yy^k \leq xx^k))]\end{aligned}$$

Fudging slightly, the extension of 'the' is a function from a property ϕ to the property of being carveable into all the ϕ-ers there are. Notice that the cover i is inert; no terms or variables are required to be indexed to it. Thus, for any i, the denotation of 'the people' is:

$$\begin{aligned}&[\![\text{the people}]\!]^i = [\![\text{the}]\!]^i([\![\text{people}]\!]) \\ &= \lambda\phi[\iota xx(\exists j(\phi(j)(xx^j)) \,\&\, \forall yy\forall k(\phi(k)(yy^k) \to yy^k \leq xx^k))](\lambda l\lambda zz[\textsc{Person}^D(zz^l)]) \\ &= \iota xx(\exists j(\textsc{Person}^D(xx^j)) \,\&\, \forall yy\forall k(\textsc{Person}^D(yy^k) \to yy^k \leq xx^k))\end{aligned}$$

Likewise, for any i, the denotation of 'the pairs (of people)' is:

$$\begin{aligned}&[\![\text{the pairs (of people)}]\!]^i = [\![\text{the}]\!]^i([\![\text{pairs of people}]\!]) \\ &= \lambda\phi[\iota xx(\exists j(\phi(j)(xx^j)) \,\&\, \forall yy\forall k(\phi(k)(yy^k) \to yy^k \leq xx^k))](\lambda l\lambda zz[\textsc{PairPeo}^D(zz^l))]) \\ &= \iota xx(\exists j(\textsc{PairPeo}^D(xx^j) \,\&\, \forall yy\forall k(\textsc{PairPeo}^D(yy^k)) \to yy^k \leq xx^k))\end{aligned}$$

3 The distribution problem solved

3.1 Distributive predication

Returning now to the distribution problem, recall the sentences:

(1) The pairs (of people) lifted the piano.

(2) The people lifted the piano.

We've seen how, on our view, 'the pairs' can co-denote with 'the people', so that some pairs just *are* the people making them up. But intuitively, there's no reading of (1) on which 'lifted the piano' distributes over people rather than pairs. Given what we've said so far, this is difficult to explain. Simplifying slightly, the truth-conditions of (1) can be derived as follows:

$$\begin{aligned}&[\![\text{the pairs lifted the piano}]\!]^{i,j} = [\![\text{lifted the piano}]\!]^i([\![\text{the pairs}]\!]^j) \\ &= \lambda xx[\textsc{Lifted}^D(xx^i)](\iota xx(\exists k(\textsc{PairPeo}^D(xx^k)) \,\&\, \forall yy\forall l(\textsc{PairPeo}^D(yy^l) \to yy^l \leq xx^l))) \\ &= \textsc{Lifted}^D(\iota xx(\exists k(\textsc{PairPeo}^D(xx^k)) \,\&\, \forall yy\forall l(\textsc{PairPeo}^D(yy^l) \to yy^l \leq xx^l))^i)\end{aligned}$$

Note that, while the description includes two bound cover-variables, 'k' and 'l', the description is *indexed* to i, and so the interpretation of 'lifted the piano' depends on our choice of i. What prevents us from selecting a cover which maps the pairs to each individual person, thereby getting the unavailable reading?

As noted in Section 1, our view is that count nouns and plural definites are associated with special covers and that some predicates are semantically constrained to be interpreted with respect to them. We first define a range of *fitting covers* for a count noun 'N' contributing a predicate 'F'. A cover i fits the noun, i.e. fits 'F', if any value of the cover consists of *one or more* Fs:

Fitting covers
A cover i fits 'F' if it satisfies these two conditions:
- For any xx, $i(xx)$ is defined whenever there is some cover j of xx such that: $\mathrm{F}^D(xx^j)$.
- When $i(xx)$ is defined, its values are such that: $\forall yy(yy^i \precsim xx^i \rightarrow \mathrm{F}^D(yy^i))$.

That is, when $i(xx)$ is defined, i carves xx into Fs and/or arbitrary combinations thereof, where the latter are subpluralities of xx which i carves into Fs, thereby allowing 'F' to distribute over them.

Among the covers that fit 'F', one of them is *basic*: when defined for some xx, it carves xx into all and only the Fs that there are among xx:

Basic cover
The basic cover $\star^{\mathrm{F}}$ is the cover that fits 'F' and whose values, whenever it's defined for some xx, satisfy: $\forall yy(yy^{\star^{\mathrm{F}}} \precsim xx^{\star^{\mathrm{F}}} \leftrightarrow (yy^{\star^{\mathrm{F}}} \leq xx^{\star^{\mathrm{F}}} \ \&\ \mathrm{F}(yy^{\star^{\mathrm{F}}})))$.

The left-to-right direction of the biconditional ensures that every value of $\star^{\mathrm{F}}(xx)$ is an F; arbitrary combinations of Fs are disallowed. The right-to-left direction ensures that no Fs included in xx are left out; any such thing(s) is/are a value of $\star^{\mathrm{F}}(xx)$.

Nominal expressions headed by the noun inherit the fitting and basic covers of the noun. For instance, 'person', 'people' and 'the people' have the basic cover which we'll label '$\star^{\mathrm{Per}}$' and which, when defined for some xx, carves them into individual people. Likewise, 'pair of people', 'pairs of people', and 'the pairs of people' have the basic cover which we'll label '$\star^{\mathrm{Pair}}$' and which, when defined for some xx, carves them into pairs of people.

With the notions of fitting and basic covers in place, the empirical observations from Section 1 can be explained as follows. A sentence must be understood with respect to a cover that fits the predicate denoted by the noun. Thus, in (2), the cover used to interpret 'lifted the piano' must fit the predicate '$\lambda xx[\textsc{Person}]$'. This ensures that the predicate distributes over individual people, and so the sentence is true just in case the piano was (distributively) lifted by one or more people in the denotation of 'the people'. The pure distributive reading, on which *each* person lifted the piano, is given by the basic cover $\star^{\mathrm{Per}}$:

$$\textsc{Lifted}^D(\iota xx(\exists k(\textsc{Person}^D(xx^k)) \ \&\ \forall yy \forall l(\textsc{Person}^D(yy^l) \rightarrow yy^l \leq xx^l))^{\star^{\mathrm{Per}}})$$

Intermediate distributive readings, on which the piano was lifted by combinations of people working in tandem, are given by various fitting covers.

By contrast, in (1), the cover used to interpret 'lifted the piano' must fit the predicate '$\lambda xx[\textsc{PeoPair}^D(xx^i)]$'. This ensures that the predicate distributes over pairs of people, and so the sentence is true just in case the piano was (distributively) lifted by one or more pairs of people in the denotation of 'the pairs'. The pure distributive reading, on which *each* pair lifted the piano, is given by the basic cover $\star^{\mathrm{Pair}}$:

$$\textsc{Lifted}^D(\iota xx(\exists k(\textsc{PeoPair}^D(xx^k)) \ \&\ \forall yy \forall l(\textsc{PeoPair}^D(yy^l) \rightarrow yy^l \leq xx^l))^{\star^{\mathrm{Pair}}})$$

Intermediate distributive readings, on which the piano was lifted by combinations of pairs working in tandem, are given by various fitting covers.

The subject of (1) is syntactically plural. What happens when it's singular?

(4) The pair (of people) lifted the piano.

Let's assume that the subject denotes *aa*. Since there's only one pair, there's only one fitting cover, $\star^{\text{Pair}}$, whose single value is *aa* themselves. So, in this case, the truth-conditions of a distributive reading of the sentence

$$\textsc{lifted}^D(aa^{\star^{\text{Pair}}})$$

turn out to be equivalent to those of its collective reading

$$\textsc{lifted}(aa^{\star^{\text{Pair}}}).$$

The sentence admits a single interpretation, saying that the pair lifted the piano together.

3.2 Objection: stubbornly distributive predicates

We've just given our account of what's going on in cases where a predicate distributes down to collections of certain kinds but is prevented from distributing down to the individuals making those collections up. But there are also cases in which a predicate seems to be forced to distribute all the way down to the members of a collection. Consider the following pairs of sentences, borrowed from Magri (2012):

(5) The pairs (of people) are sick.

(6) The people are sick.

(7) The flocks of birds landed on the sea.

(8) The birds landed on the sea.

The sentences in each pair are equivalent: a pair of people is sick just in case each of them is, and so some pairs of people are sick just in case each of the individual people is sick; likewise, a flock of birds lands just in case each of them does, and so some flocks of birds land just in case each of the individual birds does. It might seem that our account doesn't accommodate the data, here. On our view, 'are sick' must be interpreted in (5) with respect to a cover which fits the noun phrase 'pairs of people'—that is, it must be interpreted with respect to a cover which divides the denotation of 'the pairs (of people)', not into individual people, but into *pairs* of people. If distributive predicates only distribute over the values of a cover, then 'are sick' will distribute only over those pairs, and not all the way down to the individual people. So, we seemingly have no way to explain the equivalence between (5) and (6). Likewise for (7) and (8).

In fact, though, our framework gives us the resources to explain what's going on in these cases. Following Schwarzschild (2011), Magri (2012, 9) suggests that predicates like 'are sick' and 'landed on the sea' are *stubbornly distributive*; they have no collective reading.[13] We can adopt this suggestion. Focusing first on 'are sick', what this amounts to is the following:

$$\textsc{Sick}(aa^i) \to \textsc{Sick}^D(aa^i)$$

$$\equiv: \textsc{Sick}(aa^i) \to \forall xx(xx^i \precsim aa^i \to \textsc{Sick}(xx^i))$$

13. Magri calls them 'inherently distributive', but we'll stick with Schwarzschild's terminology.

That is, '$\textsc{Sick}$' is true of aa with respect to i only if it's true of every value of $i(aa)$. What this means is that, if a cover has multiple 'levels', then '$\textsc{Sick}$' is forced to distribute down through each 'level' of the cover. Suppose we have the following cover:

$$abcd \begin{array}{l} \longrightarrow ab \begin{array}{l} \longrightarrow a \\ \longrightarrow b \end{array} \\ \longrightarrow cd \begin{array}{l} \longrightarrow c \\ \longrightarrow d \end{array} \end{array}$$

Now, if '$\textsc{Sick}$' is true of $abcd$ with respect to this cover, then it distributes down to ab and to cd:

$\textsc{Sick}(abcd^i)$
$\therefore \forall xx(xx^i \precsim abcd^i \to \textsc{Sick}(xx^i))$
$\therefore \textsc{Sick}(ab^i)\ \&\ \textsc{Sick}(cd^i)$

But, by the same token, it also distributes down to each of a, b, c, and d:

$\textsc{Sick}(ab^i)$
$\therefore \forall xx(xx^i \precsim ab^i \to \textsc{Sick}(xx^i))$
$\therefore \textsc{Sick}(a^i)\ \&\ \textsc{Sick}(b^i)$

$\textsc{Sick}(cd^i)$
$\therefore \forall xx(xx^i \precsim cd^i \to \textsc{Sick}(xx^i))$
$\therefore \textsc{Sick}(c^i)\ \&\ \textsc{Sick}(d^i)$

This explains the equivalence of (5) and (6). In each sentence, 'are sick' is interpreted with respect to a cover which fits the noun. In the case of (6), the cover must carve the denotation of 'the people' into individual people. This means that the cover must eventually bottom out in individual people, and so 'are sick' must distribute all the way down to them. In the case of (5), the cover must carve the denotation of 'the pairs (of people)' into pairs of people. But recall from Section 2.3 that some things count as a pair of people with respect to i just in case i carves them into two people. So, in this case, too, the cover must eventually bottom out in individual people, and 'are sick' must distribute all the way down to them.

Similar remarks apply to (7) and (8). If 'landed on the sea' is stubbornly distributive, we can capture this in our framework as follows:

$$\textsc{Landed}(aa^i) \to \textsc{Landed}^D(aa^i)$$
$$\equiv: \textsc{Landed}(aa^i) \to \forall xx(xx^i \precsim aa^i \to \textsc{Landed}(xx^i))$$

Now, in (8), 'landed on the sea' is interpreted with respect to a cover which carves the denotation of 'the birds' into individual birds. This means that the cover must eventually bottom out in individual birds, and so 'landed on the sea' must distribute all the way down to them. In the case of (7), 'landed on the sea' is interpreted with respect to a cover which carves the denotation of 'the flocks of birds' into flocks of birds. But some things count as a flock of birds with respect to i just in case i carves them into individual birds. So, in this case, too, the cover must eventually bottom out in individual birds, and 'landed on the sea' must distribute all the way down to them.

Our framework has the resources to explain why in some cases — e.g. (1) ‘The pairs (of people) lifted the piano’ — a predicate distributes only over collections of individuals and cannot distribute over the individuals making them up. A distributive predicate is forced to distribute over a cover which fits the head noun which, in these cases, means a cover which carves the denotation of the subject term into collections of the relevant kind. But our framework *also* has the resources to explain why in some cases — e.g. (5) ‘The pairs (of people) are sick’ and (7) ‘The flocks of birds landed on the sea’ — a predicate distributes, not merely over collections of individuals, but over the individuals making up those collections. These cases involve predicates which are *stubbornly distributive*, distributing down through each ‘level’ of a cover; and the fitting covers, in these cases, bottom out in individuals of the relevant kind.

3.3 Counting

We now give our semantics for counting. We start with the case of numerals combined with simple nouns. On our view, to say that some things are n Fs is to say that the basic cover for ‘F’ divides them into n Fs. For instance, the people s_1–s_4 are four persons because the basic cover for ‘PERSON’ divides them into four persons. They’re also two pairs of people because the basic cover for ‘PAIR’ divides them into two pairs.

Putting this in terms of $\mathcal{CPL}$, we specify the property of *being exactly n Fs* using our numerically specific quantifier phrases from Section 2.3 and the basic cover $\star^{\mathrm{F}}$.

Number ascription schema
‘aa is/are n Fs’ is true iff:
$\exists_n xx(xx^{\star^{\mathrm{F}}} \precsim aa^{\star^{\mathrm{F}}})$ & $\forall yy(yy^{\star^{\mathrm{F}}} \precsim aa^{\star^{\mathrm{F}}} \rightarrow F(yy^{\star^{\mathrm{F}}}))$

And we adopt the following semantics for numerals.[14]

Numeral schema
$[\![\mathrm{n}]\!] = \lambda i \lambda \phi \lambda xx[\exists_n yy(yy^{\star^{\phi}} \precsim xx^{\star^{\phi}})$ & $\phi(\star^{\phi})(xx)]$

Fudging slightly, the extension of a numeral ‘n’ is a function from a property ϕ to the property of being ϕ and being carved into exactly n things by $\star^{\phi}$.

In simple cases, the property ϕ is just the intension of a plural predicate, such as ‘people’.

(9) aa are four people.

The relevant intensions are:

$[\![\mathrm{four}]\!] = \lambda i \lambda \phi \lambda xx[\exists_4 yy(yy^{\star^{\phi}} \precsim xx^{\star^{\phi}})$ & $\phi(\star^{\phi})(xx)]$
$[\![\mathrm{people}]\!] = \lambda i \lambda xx[\textsc{Person}^D(xx^i)]$
$[\![\mathrm{aa}]\!] = \lambda i[aa]$

The extension of a numeral operates on an entity of type <c,<e,t>>, in this case the intension of the noun phrase ‘people’. So, we just need to index ‘four’ and ‘aa’ to some covers i and j. Moreover, recall that the plural ‘people’ inherits the basic cover of ‘person’, so $\star^{\mathrm{Per}^D} = \star^{\mathrm{Per}}$.

14. Notice that the final occurrence of ‘xx’ isn’t indexed to a cover. That’s because the indexing of the term which saturates ‘xx’ will be achieved compositionally, by the application of ‘ϕ’ (see below).

$$\llbracket aa \text{ are four people}\rrbracket^{i,j} = \llbracket\text{four}\rrbracket^{i}(\llbracket\text{people}\rrbracket)(\llbracket\text{aa}\rrbracket^{j})$$
$$= \lambda\phi\lambda xx[\exists_4 yy(yy^{\star^{\phi}} \precsim xx^{\star^{\phi}}) \ \&\ \phi(\star^{\phi})(xx)](\lambda k\lambda zz[\textsc{Person}^{D}(zz^{k})])(aa)$$
$$= \lambda xx[\exists_4 yy(yy^{\star^{\mathrm{Per}}} \precsim xx^{\star^{\mathrm{Per}}}) \ \&\ \lambda k\lambda zz[\textsc{Person}^{D}(zz^{k})](\star^{\mathrm{Per}})(xx)](aa) \qquad \text{By IFA}$$
$$= \lambda xx[\exists_4 yy(yy^{\star^{\mathrm{Per}}} \precsim xx^{\star^{\mathrm{Per}}}) \ \&\ \textsc{Person}^{D}(xx^{\star^{\mathrm{Per}}})](aa) \qquad \text{By FA} \times 2$$
$$= \exists_4 yy(yy^{\star^{\mathrm{Per}}} \precsim aa^{\star^{\mathrm{Per}}}) \ \&\ \textsc{Person}^{D}(aa^{\star^{\mathrm{Per}}}) \qquad \text{By FA}$$

As promised, (9) is true just in case $\star^{\mathrm{Per}}$ carves *aa* into exactly four people.

We derive the truth-conditions of a sentence about pairs in an exactly analogous way.

(10) *aa* are two pairs (of people).

The relevant intensions are:

$$\llbracket\text{two}\rrbracket = \lambda i\lambda\phi\lambda xx[\exists_2 yy(yy^{\star^{\phi}} \precsim xx^{\star^{\phi}}) \ \&\ \phi(\star^{\phi})(xx^{i})]$$
$$\llbracket\text{pairs (of people)}\rrbracket = \lambda i\lambda xx[\textsc{PairPeo}^{D}(xx^{i})]$$
$$\llbracket\text{aa}\rrbracket = \lambda i[aa]$$

Indexing 'two' and 'pairs' to some i and j, and skipping some steps:

$$\llbracket aa \text{ are two pairs}\rrbracket^{i,j} = \llbracket\text{two}\rrbracket^{i}(\llbracket\text{pairs}\rrbracket)(\llbracket\text{aa}\rrbracket^{j})$$
$$= \exists_2 yy(yy^{\star^{\mathrm{Pair}}} \precsim aa^{\star^{\mathrm{Pair}}}) \ \&\ \forall xx(xx^{\star^{\mathrm{Pair}}} \precsim aa^{\star^{\mathrm{Pair}}} \to \textsc{Pair}(xx^{\star^{\mathrm{Pair}}}))$$

As promised, (10) is true just in case $\star^{\textsc{Pair}}$ carves *aa* into exactly two pairs of people.

The same account applies to one of our initial examples, assuming an unpronounced copy of 'pairs' after the numeral:

(3) The pairs (of people) are two (pairs).

We get

$$\llbracket\text{The pairs are two pairs}\rrbracket^{i,j} = \llbracket\text{two}\rrbracket^{i}(\llbracket\text{pairs}\rrbracket)(\llbracket\text{the pairs}\rrbracket^{j})$$

and the rest of the derivation proceeds on a similar way.

So far, we've focused on cases where a numeral combines with a simple noun. In such cases, $\phi = \star^{D}$, so our stated truth-conditions are redundant. For instance, the truth of '$\exists_2\ yy\ (yy^{\star^{\mathrm{Pair}}} \precsim aa^{\star^{\mathrm{Pair}}})$' guarantees that of '$\forall xx(xx^{\star^{\mathrm{Pair}}} \precsim aa^{\star^{\mathrm{Pair}}} \to \textsc{Pair}(xx^{\star^{\mathrm{Pair}}}))$'. However, in cases where the numeral is combined with a complex noun phrase, the second condition has an important role to play.

Consider:

(11) *aa* are four students who kissed.

What's required for some things to be four students who kissed? On a plausible reading of the sentence, 'kissed' is interpreted distributively with respect to some cover i. For instance, supposing that two pairs of students kissed and that these two pairs just *are* the four students, (11) is true just in case 'kissed' is interpreted with respect to some i which carves the four students into those two pairs. The problem is that the whole noun phrase 'students who kissed' can't be interpreted with respect to *that* cover. If it were, then by our semantics for numerals, *aa* would be *two* students who kissed, not four.

What's needed is a way to ensure that the elements of a complex noun phrase can be interpreted with respect to different covers. This is where the second of our two conditions comes in. The extensions of ‘students’ and ‘kissed’ are, with respect to some i and j:

$$[\![\text{students}]\!]^i = \lambda xx[\textsc{Student}^D(xx^i)]$$
$$[\![\text{kissed}]\!]^j = \lambda xx[\textsc{Kissed}^D(xx^j)]$$

Combining these by Predicate Modification, we get:

$$[\![\text{students who kissed}]\!]^{i,j} = \lambda xx[\textsc{Student}^D(xx^i) \;\&\; \textsc{Kissed}^D(xx^j)]$$

By λ-abstracting on i, we get an intension of type <c,<e,t>>, something on which the extension of a numeral can operate:

$$[\![\text{students who kissed}]\!]^j = \lambda i\lambda xx[\textsc{Student}^D(xx^i) \;\&\; \textsc{Kissed}^D(xx^j)]$$

Notice that the interpretation of ‘$\textsc{Student}^D$’ depends on a cover i, while that of ‘$\textsc{Kissed}^D$’ depends on a potentially different cover j.

The complex noun phrase ‘students who kissed’ inherits the basic cover of ‘student’, ‘$\star^{\text{Stu}}$’, so we can derive the truth-conditions of (11) as follows.

$$\begin{aligned}
&[\![aa \text{ are four students who kissed}]\!]^{i,j,k} = [\![\text{two–pairs}]\!]^i([\![\text{students who kissed}]\!]^j)([\![\text{aa}]\!]^k)\\
&= \lambda\phi\lambda xx[\exists_4 yy(yy^{\star^{\phi}} \precsim xx^{\star^{\phi}}) \;\&\; \phi(\star^{\phi})(xx^i)](\lambda l\lambda zz[\textsc{Student}^D(zz^l) \;\&\; \textsc{Kissed}^D(zz^j)])(aa)\\
&= \lambda xx[\exists_4 yy(yy^{\star^{\text{Stu}}} \precsim xx^{\star^{\text{Stu}}}) \;\&\; \lambda l\lambda zz[\textsc{Student}^D(zz^l) \;\&\; \textsc{Kissed}^D(zz^j)](\star^{\phi})(xx^i)](aa)\\
&= \lambda xx[\exists_4 yy(yy^{\star^{\text{Stu}}} \precsim xx^{\star^{\text{Stu}}}) \;\&\; \textsc{Student}^D(xx^{\star^{\text{Stu}}}) \;\&\; \textsc{Kissed}^D(xx^j)](aa)\\
&= \exists_4 yy(yy^{\star^{\text{Stu}}} \precsim aa^{\star^{\text{Stu}}}) \;\&\; \textsc{Student}^D(aa^{\star^{\text{Stu}}}) \;\&\; \textsc{Kissed}^D(aa^j)
\end{aligned}$$

The sentence is true just in case (i) $\star^{\text{Stu}}(aa)$ has exactly four values, each of which is a student, and (ii) aa kissed with respect to j. So, even if $j(aa)$ has as its values two pairs of students, aa get counted as *four* students, rather than two. Here, the numeral combines with a complex noun phrase and ‘$\textsc{Student}^D(aa^{\star^{\text{Stu}}}) \;\&\; \textsc{Kissed}^D(aa^j)$’ adds an important restriction to ‘$\exists_4 yy(yy^{\star^{\text{Stu}}} \precsim aa^{\star^{\text{Stu}}})$’.

Our account has much in common with the one in Payton (2022). But there are some important differences. Payton is concerned to show that, for instance, four people can be identical to two pairs of people, even though the former are four, the latter are two, and *being four* and *being two* are *prima facie* incompatible properties.[15] To this end, he defends the following:

Payton’s number ascription schema
‘aa is/are n Fs’ is true iff:
(i) $n = 1$ and $\exists xx(Fxx \;\&\; xx = aa)$; or
(ii) $n \geq 2$ and $\exists xx_1 \ldots \exists xx_n(Fxx_1 \;\&\; \ldots \;\&\; Fxx_n \;\&\; xx_1 \neq xx_2 \;\&\; \ldots \;\&\; xx_{n-1} \neq xx_n \;\&\; xx_1 \ldots xx_n = aa)$

That is, aa are n Fs just in case there are n Fs to which aa, taken together, are identical. On our view, this is only a *necessary* condition on the truth of ‘aa are n Fs’: that sentence is true only if $\star^{\text{F}}$ carves aa into n subpluralities each of which is F; given how covers are defined, those Fs taken together must be identical to aa.

15. In that paper, Payton is also concerned to defend ‘composition as identity’, the thesis that a whole is identical to its parts, despite the *prima facie* impossibility of something’s being both one and many. We’re not concerned with this thesis here.

The account in Payton (2022) also makes no use of covers, or of any other mechanism which could connect the truth of 'the n Fs are G' to the way in which the Fs are carved up and counted. Thus, while it's easy to see how, on that account, 'the four people' and 'the two pairs' could be co-denoting (the same plurality could be identical *both* to some four people *and* to some two pairs), it's hard to see how it could explain why in (1) 'lifted the piano' distributes over pairs while in (2) it distributes over people. Our account patches this hole.

4 Alternatives

Our framework allows us to give a unified, cover-based solution to both aspects of the distribution problem. We close by comparing our solution to two alternatives.

4.1 Covers versus genericity

Magri (2012) aims to explain why certain predicates, whether collective or distributive, are *shareable* — that is, if they're had by a group then they're had by the members of that group, and *vice-versa* — while others aren't. We've investigated non-shareable distributive predicates in Section 3.1 and shareable ones in Section 3.2, in both cases making reference to Magri's own data. It's worth comparing our approach to his.

Magri's starting hypothesis is that non-shareable predicates are *individual-level predicates* (ILPs): roughly, predicates which stand for permanent, or at least relatively stable, properties. Shareable predicates, by contrast, are *stage-level predicates* (SLPs): roughly, predicates which stand for less stable, more transient properties. If every dog in a pack is hungry, or tired, or sick (these being SLPs), then we can say, 'The pack is hungry/tired/sick'. By contrast, even if every dog in a pack is tall, or fat, or intelligent (these being ILPs), we can't say, 'The pack is tall/fat/intelligent.' For Magri, the fact that these predicates are ILPs is what explains their non-shareability (2012, 16–17).

Following Chierchia (1995), Magri treats a sentence in which an ILP is attributed to some thing(s) as a kind of *generic*: where 'F' is an ILP, 'aa is/are F' means something like, 'In general, aa is/are F'.[16] More specifically, he treats ILPs as containing implicit variables for Kratzer's (1989) situations, and sentences in which an ILP is attributed to some thing(s) as containing a generic operator, 'GEN', which operates on properties of situations:

$$[\![\text{GEN}]\!] = \lambda\phi_{<\text{i, t}>}\lambda\psi_{<\text{i, t}>}[\text{Generally, if } \phi(s) \text{ then } \psi(s)]$$

What goes in for 'ψ' is always the result of saturating an ILP with the subject-term. What goes in for 'ϕ' is the (unpronounced) property *being a situation which contains the denotation of the subject-term*, or $\lambda s[\textsc{In}([\![\text{N}]\!], s)]$. In effect, '$aa$ are F' means, 'In general, if s contains $[\![\text{N}]\!]$ then $[\![\text{N}]\!]$ is F in s.

How does this account explain non-shareability? Modifying one of Magri's examples, let aa be some glasses, each of which is short, but which have been arranged in a

16. Throughout this discussion we continue to use plural terms and variables. Magri himself doesn't adopt the resources of plural logic. Rather, he adopts a *mereological* approach to plurals on which a plural term refers to the sum of some individuals — see e.g. Link (1983) and Landman (1989). Nothing hinges on this disagreement.

tall pile. In this case, 'the glasses' and 'the pile (of glasses)' refer to the same things, and yet (12) is true while (13) is false.

(12) The pile (of glasses) is tall.

(13) The glasses are tall.

Because we're considering a case in which (13) is false while (12) is true, (Magri, 2012, 18) posits a reading of 'is tall', $\lambda xx\lambda s[\textsc{Tall}_{\text{Pile}}(xx, s)]$, which is satisfied by the pile but not by the individual dishes — intuitively, $\lambda xx\lambda s[\textsc{Tall}_{\text{Pile}}(xx, s)]$ is true of aa in s just in case aa form a tall pile in s. In this sort of case, he thinks, (12) and (13) are analyzed, respectively, as follows:

$$\begin{aligned}&[\![\text{The pile (of glasses) is tall}]\!]\\ &= \lambda\phi\lambda\psi[\text{Generally, if } \phi(s_2) \text{ then } \psi(s_2)](\lambda s_1[\textsc{In}([\![\text{the pile (of glasses)}]\!], s_1)])(\lambda s_3[\textsc{Tall}_{\text{Pile}}(aa, s_3)]\\ &= \text{Generally, if } \textsc{In}([\![\text{the pile (of glasses)}]\!], s_2) \text{ then } \textsc{Tall}_{\text{Pile}}(aa, s_2)\end{aligned}$$

$$\begin{aligned}&[\![\text{The glasses are tall}]\!]\\ &= \lambda\phi\lambda\psi[\text{Generally, if } \phi(s_2) \text{ then } \psi(s_2)](\lambda s_1[\textsc{In}([\![\text{the glasses}]\!], s_1)])(\lambda s_3[\textsc{Tall}_{\text{Pile}}(aa, s_3)]\\ &= \text{Generally, if } \textsc{In}([\![\text{the glasses}]\!], s_2) \text{ then } \textsc{Tall}_{\text{Pile}}(aa, s_2)\end{aligned}$$

But now, crucially, while Magri thinks that 'the glasses' and 'the pile (of glasses)' are co-*referring*, he denies that they're co-*denoting*. The denotation of a term is an *individual concept* — or, since we're assuming plural reference in our discussion, a *plurality concept* — i.e. a function from situations to things. $[\![$The glasses$]\!]$ is constant, returning aa for any situation. By contrast, $[\![$the pile (of glasses)$]\!]$ returns aa only for situations in which aa form a pile (Magri, 2012, 17). Thus, (12) and (13) have different truth-conditions. (12) says that in general, if the glasses form a pile then they form a *tall* pile. This condition, Magri thinks, is satisfied in the imagined scenario. (13), by contrast, says that in general, if the glasses *exist* then they form a tall pile. And this condition is plausibly not satisfied.

There are a few problems with Magri's view. First, the starting hypothesis that all ILPs are non-shareable is questionable. Consider a few variants on Magri (2012, 16)'s examples:

(14) That group of philosophers is ugly.

(15) That bunch of apples is spoiled.

Magri thinks that (14) should strike us as odd — even if each philosopher is ugly, we can't say that *the group* is ugly. But compare the exclamation, 'What an ugly group of philosophers!', which suggests that the group as a whole shares the property of ugliness with its members. Similarly, Magri might suggest that even if each apple is spoiled, we can't say that *the bunch* is. But compare the old saying, 'One rotten apple spoils the bunch', which suggests that a bunch can share the property of being spoiled with its members.

Second, the truth-conditions Magri assigns to (12) and (13) are questionable. He says that (12) is true just in case it's true *in general* that if the glasses form a pile then they form a tall pile. But it's easy to imagine that (a) there are many ways of arranging the glasses into a pile which *isn't* tall while (b) the glasses are *in fact* arranged into a tall pile. (12) should be true in this case, but Magri counts it as false. Putting the point another way: whether (12) is true should depend on whether the pile which the glasses *actually* form is tall, not on whether most piles they *could* form are tall.

This point is perhaps easier to see if we substitute for 'The pile (of glasses)' its plural variant, 'The piles (of glasses)':

(16) The piles (of glasses) are tall.

If we have enough glasses to form multiple piles, it should be *very* easy to form multiple *short* piles. But then, Magri will count (16) false even when the glasses form tall piles. And again, this seems wrong; what matters for the truth or falsity of (16) is whether the glasses are *in fact* arranged in tall piles, not how easily they *could* be arranged in short piles.

Third, and relatedly, Magri says that there's a reading of (13) on which it's true just in case, *in general*, if the glasses exist then they form a tall pile. This reading is supposed to be induced in a case of the sort considered — one in which each individual glass is short, but they happen to form a tall pile. That strikes us as implausible: whether (13) is true should depend on whether each individual glass is tall; that is, the relevant reading of 'are tall' should be the straightforward distributive reading.[17]

Finally, Magri is able to distinguish the truth-conditions of (12) and (13) only because he's able to distinguish the individual concepts denoted by 'the pile' and 'the glasses', and he's able to do *this* only because there's a substantial condition that *aa* have to meet in order to constitute a pile; their mere existence isn't enough. But now recall, once again:

(1) The pairs (of people) lifted the piano.

(2) The people lifted the piano.

It seems to us that there's no extra, substantial condition that two people need to meet, in order to count as a pair: *any* two people are a pair of people, by virtue of being two. Likewise, any even-numbered plurality of people are pairs, by virtue of being even-numbered. But then, Magri can't distinguish the concepts denoted by 'the pairs' and 'the people': And if he can't distinguish these individual concepts, then he can't distinguish the truth-conditions of (1) and (2).

Our approach faces none of these problems. We have no need to hypothesize that shareability tracks with ILPs. And our view extends naturally to (12) and (13), repeated below for convenience:

(12) The pile (of glasses) is tall.

(13) The glasses are tall.

Let *aa* be the denotation of both 'the glasses' and 'the pile (of glasses)' in the circumstances Magri imagines — i.e. let them be some glasses, each of which is short, but which are in fact arranged into a tall pile. (13) requires that 'tall' be interpreted with respect to a cover that carves *aa* into all and only the glasses among them. If 'tall' is here being read distributively (as it plausibly is), it distributes down to each glass,

17. Magri makes use of the distributive reading of 'are tall' *only* to account for cases in which (13) is *true* due to the heights of the individual glasses (Magri, 2012, 18). So, he seems to be committed to thinking that that reading of the predicate is only available when (13) is true, and not when it's false, despite the fact that its availability would give the intuitively correct explanation for *why* (13) is false.

and (13) is true just in case each glass is tall. So, it comes out false, as desired. By contrast, in (12) 'tall' is read non-distributively, and interpreted with respect to a cover that carves *aa* into all and only the piles among them — that is, a cover the first 'level' of which is just *aa* themselves. Since we see no reason to think that, on this reading, its applicability should depend on the individual heights of the things making up a pile) we can again count it as true just in case *aa* form a tall pile. So, (12) is true.[18]

4.2 Covers versus higher-level plurals

Throughout this paper we've adopted a framework of plural logic. In recent years, philosophers interested in this kind of framework have wondered whether it can be extended to create *higher-level plural* logics. Just as we can refer to, and quantify over, pluralities of *individuals*, might we be able to refer to, and quantify over, pluralities of *pluralities*?[19] Oliver and Smiley (2016) argue that we can, and in particular that collective noun phrases are higher-level plural terms (2016, 305–307). 'The pairs (of people)', for instance, doesn't denote a plurality of individual people (at least, not in the same way that 'the people' does); rather, it denotes a plurality of pluralities. This approach suggests an alternative solution to the distribution problem.

Picturesquely, Oliver and Smiley divide pluralities into different levels. At level 0 are those pluralities which contain only one thing, i.e. the individuals. At level 1 are all those pluralities which can be 'built up' from two or more individuals. At level 2 are all those pluralities which can be 'built up' from one of these 'first-level' pluralities together with other first-level pluralities and/or individuals. And so on (p. 311, pp. 314–317). All these pluralities are ordered by a non-transitive *vertical inclusion* relation. So, for instance, if the two pairs are Alice and Beth, and Cathy and Diane, then (a) there is a first-level plurality which vertically includes each of the four women but neither of the two pairs and (b) there is a second-level plurality which vertically includes each of the two pairs but neither of the four women. Finally, suppose that distributive predicates can distribute across this relation. Then, in (1),

(1) The pairs (of people) lifted the piano.

'lifted the piano' distributes over the two pairs but not over the four people. By contrast, in (2)

(2) The people lifted the piano.

'the people' denotes the first-level plurality of Alice, Beth, Cathy, and Diane, and so 'lifted the piano' distributes over the four people but not over the two pairs.[20]

18. Magri (2012, 16) suggests that 'tall' is *stubbornly distributive*. If that's right, then our view would still count (12) as true just in case each glass in the pile is tall — for, the basic cover of 'pile of glasses' would *first* carve *aa* into the pile *aa* themselves and *then* into the individual glasses, and 'tall' would be forced to distribute down to each individual glass. (Compare our treatment of 'sick' in Section 3.2). But in fact, we think 'tall' is *not* stubbornly distributive. Witness the fact that (16), 'The piles (of dishes) are tall', doesn't entail (13) in the way that (5), 'The pairs (of people) are sick' entails (6), 'The people are sick'.

19. See e.g. Rayo (2006), Linnebo and Nicolas (2008), Oliver and Smiley (2016, ch. 15), Simons (2016), Florio and Linnebo (2021, ch. 9), and Grimau (2021a).

20. Oliver and Smiley don't explicitly discuss distributive predication in the context of their higher-level language, but this is a natural extension of their approach elsewhere in the book. The intended story may be more complex. In addition to vertical inclusion, which obtains only between pluralities at different levels, Oliver and Smiley (2016, 307–308) posit a *horizontal inclusion* relation which obtains between pluralities at the same level — e.g. between Alice and Beth on the one hand, and Alice, Beth, and Cathy on the other (where the latter is a first-level plurality 'built up' from the three women and vertically includes no pairs).

Similarly, although Oliver and Smiley don't give a detailed picture of the semantics of numerals, they do insist that pluralities, as well as individuals, can be counted (pp. 306–307). This suggests a picture on which, for instance, 'the four people' denotes a first-level plurality 'built up' from four individual people while 'the two pairs (of people)' denotes a second-level plurality 'built up' from two two-membered first-level pluralities of people. This would explain why in (3),

(3) The pairs (of people) are two.

each pair, and not each person, is counted 'as one'.

However, not everyone is convinced that higher-level reference is coherent.[21] Ben-Yami (2013, 82) expresses this skeptical view as follows:

> The difference between singular and plural referring expressions consists in the former referring to *a single individual*, and the latter referring to *more than a single individual*. If so, what could iterating the step from the singular to the plural mean?... [I]f it were expressed by iterating the difference between the two italicised phrases above — and why shouldn't it? — then a superplural expression should refer to *more than more than a single individual*. But what could that mean?

Even if these skeptical worries can be answered, problems remain for the suggested solution to the distribution problem. Oliver and Smiley defend higher-level reference by appeal to the idea that collective noun phrases (and perhaps other referring terms) are pseudo-singular. We've seen that, if that's right, then terms like 'the people' and 'the pairs' should, in the right contexts, be co-denoting. But it's not clear that Oliver and Smiley's higher-level plural logic secures this result. Recall their picture of pluralities as stratified into different levels and individuated by what they vertically include. If we take this picture seriously, then it seems that any first-level plurality of people must be distinct from any second-level plurality of pairs of people, even if the latter is in some sense 'ultimately built up' from those same individuals. The first-level plurality *Alice, Beth, Cathy, and Diane* vertically includes the four people but no pairs, while the second-level plurality *Alice and Beth, and Cathy and Diane* vertically includes the two pairs but no people. These pluralities are therefore distinct, just as the sets {Alice, Beth, Cathy, Diane} and {{Alice, Beth}, {Cathy, Diane}} are.[22] So it can't be that, in the case described above, 'the people' and 'the pairs' are co-denoting.

Of course, typical proponents of higher-level plural resources would deny that the picture *should* be taken seriously — although Simons (2016, personal communication) is an exception. The terms 'Alice, Beth, Cathy, and Diane' and 'Alice and Beth, and Cathy and Diane' aren't meant to denote different pluralities, they're meant to denote the same plurality in different ways. And likewise, of course, for 'the people' and 'the pairs'. However, at least one of us is skeptical that this attempt to replace a metaphysical commitment to distinct pluralities with an ideological commitment

---

They take this distinction to effectively pull apart two components of the more familiar idea of *inclusion* we've borrowed from ordinary plural logic. Since, when developing their first-level plural language, they treat distributive predicates as distributing across the latter relation (pp. 114–115), they may want to allow such predicates to distribute across *both* vertical *and* horizontal inclusion in their higher-level language.

21. See especially Simons (1982, 187-195), Uzquiano (2004, 438-440), McKay (2006, 46-53,137-39), and Ben-Yami (2013).

22. Oliver and Smiley (2016, 314–317) suggest that a higher-level plural logic, which includes a non-transitive vertical inclusion relation and a resulting stratification of pluralities into levels, could effectively mimic (Cantorian) set theory.

to different modes of reference can succeed (Payton, nd). In any case, our approach requires neither the metaphysical commitment to distinct pluralities nor the ideological commitment to higher-level plural reference.[23]

5 Conclusion

The distribution problem poses a serious challenge to the view that collective nouns are pseudo-singular. Cover-based semantics provides a partial solution, but the traditional, thoroughly pragmatic approach to cover-selection leaves us unable to explain all that needs to be explained. We've shown how cover-selection can be semantically constrained and used this hypothesis to provide a unified solution to the distribution problem. This approach is more satisfying than the traditional cover-based one, and also avoids the drawbacks of alternative solutions.

Acknowledgments. A previous version of this paper was presented at Bilkent University. Thanks to the members of that audience, and to two anonymous referees for *Synthese*, for their helpful comments. David Nicolas's work was supported by the grant ANR-17-EURE-0017. Jonathan Payton's work was supported by a Young Scientists Award (Genç Bilim İnsanları Ödülü) from the Science Academy of Türkiye (Bilim Akademisi).

References

Ben-Yami, H. 2013. Higher-level plurals versus articulated reference, and an elaboration of *Salva Veritate*. *Dialectica 67*(1): 81–102. https://doi.org/10.1111/1746-8361.12013 .

Black, M. 1971. The elusiveness of sets. *Review of Metaphysics 24*(4): 614–636 .

Champollion, L. 2017. *Parts of a Whole: Distributivity as a Bridge Between Aspect and Measurement*. Oxford: Oxford University Press.

Chierchia, G. 1995. Individual-level predicates as inherent generics, In *The Generic Book*, eds. Carlson, G.N. and F.J. Pelletier, 125–175. Chicago: University of Chicago Press.

Florio, S. and Ø. Linnebo. 2021. *The Many and the One*. New York: Oxford University Press.

Gillon, B.S. 1987. The readings of plural noun phrases in English. *Linguistics and Philosophy 10*(2): 199–219. https://doi.org/10.1007/bf00584318 .

Gillon, B.S. 1990. Plural noun phrases and their readings: Reply to Lasersohn. *Linguistics and Philosophy 13*(4): 477–485 .

23. For criticism of Oliver and Smiley's treatment of collective nouns from a very different direction, see Snyder and Shapiro (2021); for their reply, see Oliver and Smiley (2022).

Gillon, B.S. 1992. Towards a common semantics for English count and mass nouns. *Linguistics and Philosophy* 15: 597–639. https://doi.org/10.1007/bf00628112 .

Grimau, B. 2021a. In defense of higher-level plural logic: Drawing conclusions from natural language. *Synthese 198*(6): 5253–5280. https://doi.org/10.1007/s11229-019-02399-z .

Grimau, B. 2021b. Structured plurality reconsidered. *Journal of Semantics 38*(1): 145–193. https://doi.org/10.1093/jos/ffaa012 .

Heim, I. and A. Kratzer. 1998. *Semantics in Generative Grammar*. Malden: Blackwell Publishing, Inc.

Horden, J. and D. López de Sa. 2021. Groups as pluralities. *Synthese 198*(11): 10237–10271. https://doi.org/10.1007/s11229-020-02715-y .

Kennedy, C. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30: 1–45. https://doi.org/10.1007/s10988-006-9008-0 .

Korman, D.Z. 2015. *Objects: Nothing out of the Ordinary*. Oxford: Oxford University Press.

Kratzer, A. 1989. An investigation of the lumps of thought. *Linguistics and Philosophy 12*(5): 607–653. https://doi.org/10.1007/BF00627775 .

Landman, F. 1989. Groups I. *Linguistics and Philosophy 12*(5): 559–605. https://doi.org/10.1007/bf00627774 .

Landman, F. 2020. *Iceberg Semantics for Mass Nouns and Count Nouns. A New Framework for Boolean Semantics*. Springer Cham.

Link, G. 1983. The logical analysis of plurals and mass terms, In *Meaning, Use, and Interpretation of Language*, eds. Bäuerle, R., C. Schwarze, and A. von Stechow, 303–323. Walter de Gruyter. https://doi.org/https://doi.org/10.1515/9783110852820.302.

Linnebo, Ø. and D. Nicolas. 2008. Superplurals in English. *Analysis 68*(3): 186–197. https://doi.org/10.1093/analys/68.3.186 .

Magri, G. 2012. Collective nouns without groups. In *IATL 27: Proceedings of the 27th annual meting of the Israeli Association for Theoretical Linguistics*, Haifa, Israel, pp. 1–20.

McKay, T.J. 2006. *Plural Predication*. Oxford: Oxford University Press.

Nicolas, D. and J.D. Payton. n.d. Superplurals analyzed away .

Oliver, A. and T. Smiley. 2016. *Plural Logic: Second Edition, Revised and Enlarged*. Oxford: Oxford University Press.

Oliver, A. and T. Smiley. 2022. Snyder and Shapiro's critique of pseudo-singularity. *Thought 11*(4): 226–231. https://doi.org/10.5840/tht202451629 .

Payton, J.D. 2021. Composition as identity, now with all the pluralities you could want. *Synthese 199*(3-4): 8047–8068. https://doi.org/10.1007/s11229-021-03152-1 .

Payton, J.D. 2022. Counting composites. *Australasian Journal of Philosophy 100*(4): 695–710. https://doi.org/10.1080/00048402.2021.1938617 .

Payton, J.D. n.d. From singular to plural...and beyond? .

Rayo, A. 2002. Word and objects. *Noûs 36*(3): 436–464. https://doi.org/10.1111/1468-0068.00379 .

Rayo, A. 2006. Beyond plurals, In *Absolute Generality*, eds. Rayo, A. and G. Uzquiano, 220–254. Oxford: Oxford University Press.

Ritchie, K. 2017. Plural and collective noun phrases, In *The Routledge Handbook of Collective Intentionality*, eds. Jankovic, M. and K. Ludwig, 464–475. London: Routledge.

Schwarzschild, R. 1996. *Pluralities*. Dordrecht: Kluwer Academic Press.

Schwarzschild, R. 2011. Stubborn distributivity, multiparticipant nouns and the count/mass distinction, In *Proceedings of NELS 39*, eds. Lima, S., K. Mullin, and B. Smith, 661–678. North Eastern Linguistic Society.

Sider, T. 2014. Consequences of collapse, In *Composition as Identity*, eds. Cotnoir, A.J. and D.L.M. Baxter, 211–221. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199669615.003.0011.

Simons, P. 1982. Number and manifolds, In *Parts and Moments: Studies in Logic and Formal Ontology*, ed. Smith, B., 160–198. Philosophia Verlag.

Simons, P. 2016. The ontology and logic of higher-order multitudes, In *Unity and Plurality: Logic, Philosophy, and Linguistics*, eds. Carrara, M., A. Arapinis, and F. Moltmann. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198716327.003.0004.

Snyder, E. and S. Shapiro. 2021. Group nouns and pseudo-singularity. *Thought 10*(1): 66–77. https://doi.org/10.1002/tht3.476 .

Uzquiano, G. 2004. Plurals and simples. *The Monist 87*(3): 429–451. https://doi.org/10.5840/monist200487324 .

Yi, B.U. 2005. The logic and meaning of plurals. Part I. *Journal of Philosophical Logic 34*(5–6): 459–506. https://doi.org/10.1007/s10992-005-0560-9 .

Yi, B.U. 2006. The logic and meaning of plurals. Part II. *Journal of Philosophical Logic 35*(3): 239–288. https://doi.org/10.1007/s10992-005-9015-6 .