Restrictions on subkind coercion in object mass nouns.[1]

Peter Sutton — *Heinrich Heine University, Düsseldorf*
Hana Filip — *Heinrich Heine University, Düsseldorf*

Abstract. Rothstein (2015) raises the following puzzle. Why should it be the case that mass nouns like *rice* receive a coerced taxonomic plural reading when directly modified by a numerical expression, but object mass nouns like *furniture*, *jewellery* and *crockery* do not? For example, *three rices* can mean THREE KINDS OF RICE, but *three furnitures* cannot mean THREE KINDS OF FURNITURE. We attempt to solve this puzzle by providing an analysis of kind readings for concrete count and mass nouns generally, which is based on the analysis of their non-kind predicate readings in Sutton and Filip (2016a). The key property driving our analysis is the extensional overlap of subkinds at each level of categorisation. Object mass nouns have extensionally overlapping subkinds relative to a level of categorisation, while other mass nouns do not. We also differentiate between count and mass nouns in terms of counting contexts. Artefact denoting count nouns such as *vehicle* have a felicitous taxonomic plural, because, on our account, count nouns are linked to specific counting contexts, which force the resolution of potential overlap between objects in their denotation. Artefact denoting mass nouns such as *furniture* do not have a felicitous taxonomic plural, because mass nouns are saturated with the null counting context which leaves any overlap in the noun's subkind structure unresolved.

Keywords: kinds, subkinds, taxonomic plurals, count/mass distinction, context sensitivity.

1. Introduction

Concrete mass nouns can be generally coerced into count noun interpretations, PORTION or SUBKIND, modulo context (as is well known since at least Pelletier (1975)). Different classes of nouns differ with respect to the ease with which they can be coerced into a count interpretation. For example, *water* in (1a)-(1b) is easier to coerce into a count noun interpretation than *rice* in (2a)-(2b), while count interpretations of *mud* are possible only in highly specialised contexts such as in (3a)-(3b).

(1) a. Three waters, please!
e.g. three [GLASSES/BOTTLES OF] water. (portion)
b. I ordered three waters for the party: still, sparkling, and fruit-flavoured for the kids.
i.e. three [KINDS OF] water (subkind)

(2) a. We ordered the main courses with two plain rice, one egg fried rice and a nan, more than enough for the four of us.[2]

[1]This research was funded by the German Research Foundation (DFG) CRC991, project C09. We would like to thank the participants of *Sinn und Bedeutung 21*, the *Universität Köln Linguistischer Arbeitskreis Köln (LAK)*, and the *6th Annual Bar-Ilan Slavic & Semantics Workshop* for helpful feedback. In particular, we would like to thank Susan Rothstein and Fred Landman for their input and inspiration. Thanks, too, to Natalja Beckman (project C03 in CRC991) for help with sourcing German examples.

[2]`http://www.derbytelegraph.co.uk/speciality-dishes-star-turn-littleover-s-red/`

e.g. two [PORTIONS OF] plain rice. (portion)

b. *Context: three kinds of rice, Calmati, Texmati, Kasmati*
These three rices have basmati's viscosity and cooking style, but smaller individual grains.[3]
i.e. three [KINDS OF] rice (subkind)

(3) *Context: yield points of different mud samples before contamination*
The three muds experienced particles dispersion at the same temperature with different yield points.[4]

a. The three [SAMPLES OF] mud... (portion)

b. The three [KINDS OF] mud.... (subkind)

Object mass nouns (which have also been called, *inter alia*, 'fake' or 'neat' mass nouns) include, in English, *furniture*, *footwear*, *cutlery*, *crockery*, *equipment*. These nouns are, like prototypical count nouns such as *chair* and *cat*, considered to be 'naturally atomic' in that "what counts as one entity is not determined by context but by the naturally atomic structure of the stuff. What counts as one *P* is part of our knowledge of what a *P* is" (Rothstein, 2010). What is remarkable and complicates the analysis of this subset of mass nouns derives from the observation that objects that count as one *P* in the denotation of a given object mass *P* stand in a network of kind-subkind relations. For example, *furniture* is a SUPERORDINATE CATEGORY term in that its constituent members comprise terms for BASIC LEVEL categories: *table, chair, bed*, etc. Basic level categories, in turn, subsume members labeled by SUBORDINATE LEVEL categories: for *chair*, we have *kitchen chair, dentist chair*, for instance. There is also a level of subkinds of superordinate categories: for FURNITURE, we have *bedroom furniture*, which is a kind of *furniture*.

1.1. The Puzzle: Rothstein (2015)

Object mass nouns strongly resist coercion in numerical count constructions, in which either specific ordinary individuals of the same (sub)kind) or different (sub)kinds are counted. For example, in (4a), we see a resistance to the grammatical counting of individual items of furniture (that realise basic-level kinds). In (4b), we see resistance to counting basic level kinds, and in (4c), we see resistance to counting superordinate-level kinds.

(4) a. #I ordered three furnitures from Ikea: one table and two chairs.
b. #I ordered two furnitures from Ikea: chairs and tables.
c. #I ordered two furnitures from Ikea: bedroom and living room furniture.

This is puzzling because there are clearly identifiable and conceptually accessible 'atomic' entities in the denotation of object mass nouns (e.g., single items of furniture such as single tables or chairs), and yet they cannot be accessed via coercion. Subordinate categories labeled

`story-20536589-detail/story.html` ACCESSED: 10/10/2016.

[3] *The Ultimate Rice Cooker Cookbook*, Harvard Common Press, 2003. p. 23

[4] From: Adekomaya, Olufemi A. 2013. Experimental analysis of the effect of magnesium saltwater influx on the behaviour of drilling fluids. Journal of Petroleum Exploration and Production Technology 3. 61–67.

by object mass nouns have identifiable sub-/superkinds, and yet members at these levels cannot be accessed for grammatical counting in coerced environments, even when the relevant sub-/superkinds are easily retrievable from the context.

The above data are well-known, but rarely directly addressed, with a few notable exceptions. Among the recent ones is Sutton and Filip (2016a) dealing with the restrictions on grammatical counting of object mass nouns that concerns the cardinality of ordinary objects that realise basic-level kinds (4a). Rothstein (2015), as a side observation, raises the question about the lack of pluralisation with a subkind interpretation (illustrated by our example (4b) above): If mass nouns like *water* (1b) can pluralise with a coerced taxonomic subkind interpretation, why do object mass nouns not have such taxonomic plurals? However, the answer to this question lies beyond the scope of her main agenda. A recent attempt at answering this question can be found in Grimm and Levin (2016), who focus on the failure of artefact denoting nouns to form well-formed taxonomies (see Section 5). Building on some ideas Sutton and Filip (2016a), in this paper, we pose our main question as follows: Why do object mass nouns resist mass-to-count coercion for counting taxonomic subkinds?

1.2. Superordinacy

The (apparent) SUPERORDINATE nature of object mass nouns like *furniture* cannot be what underlies the data. There are SUPERORDINATE COUNT nouns like *vehicle, weapon*, which are grammatically countable, and which have natural plurals with a subkind interpretation, as in (5a). Some of these have mass counterparts that cannot be coerced into countable subkind readings (5b).

(5) a. The brief for the government-backed project is to produce four vehicles ranging in size from the Ford Fiesta to the Vauxhall Cavalier.[5]
b. #The brief for the government-backed project is to produce four transports...

It is worth noting that the restriction on a coerced subkind interpretation of object mass nouns is less strict in other languages. For instance, in German, there are subkind readings for indefinite NPs formed with object mass nouns. Take *Gebäck* ('pastry' or 'baked good'), for instance, as in (6a). However, felicity is significantly worsened for direct numerical attachment, as in (6b). Such cross-linguistic differences have not yet been noticed, to our knowledge, but they lie outside of the scope of this paper and we plan to investigate them in our further research.

(6) a. Ein Gebäck, das in der Osterzeit auf keinem Kaffeetisch fehlen darf,
A pastry REL in the Easter.time on no.DAT coffee.table lack.INF may.3SG
ist Mamoule.
be.3SG ma'amoul
'A type of pastry that is a must on any coffee table over Easter is ma'amoul.'[6]

[5] `http://www.independent.co.uk/arts-entertainment/motoringfamily-planning-1577828.html` (15.03.2016).

[6] Obtained from the DWDS corpus `https://www.dwds.de`

b. #Ich habe zwei Gebäck gekauft: Schweineohren und Kekse.
I have.1SG two pastry buy.PAST pig.ear.PL and cookie.PL
Int: 'I bought two kinds of baked goods: palmier pastries and cookies.'

1.3. Outline

Object mass noun concepts like FURNITURE are associated with sub-/superkind structures that are inconsistent with a well-formed taxonomy, as has been observed (see e.g., Wierzbicka, 1985; Wisniewski et al., 1996). Related to this, we propose that they have sub-/superkind structures that overlap in their extensions, and such overlap cannot be contextually coerced into disjointness. To count kinds, overlap must be resolved; if overlap cannot be resolved, counting of 'kind units' goes wrong.

We implement this idea by relying on ideas independently introduced in three accounts of the mass/count distinction, which will be briefly summarized in Section 2: Rothstein (2010), Landman (2011), and Sutton and Filip (2016a). In Section 3, we give a semantics for kind readings of count and mass nouns by generalising restrictions on counting of particular individuals, proposed in Sutton and Filip (2016a), to restrictions on counting of kinds. In Section 4, we provide a semantics for explicit kind-extracting expressions such as *kind of* and *type of* and we show how our account from Section 3 can be used to derive the restrictions on subkind coercion. In Section 5, we briefly discuss Grimm and Levin (2016), the one other account of these data we are aware of, in the light of our analysis.

2. Background

2.1. Rothstein (2010): Count nouns are indexed to counting contexts

There is an agreement that prototypical mass nouns (*water, air*) are not naturally atomic (i.e., do not have stably discrete, non-overlapping objects in their denotation across all contexts), and are divisible (proper parts of *P*s are also *P*s). Building on Zucchi and White (1996, 2001) (and B. Partee, p.c.), Rothstein (2010) focuses on puzzling nouns, the denotations of which are divisible and not naturally atomic, but which are nonetheless lexicalized as count nouns. For example, take *fence*. What counts as a single fence is determined in context, so by Rothstein's definition, *fence* is not naturally atomic. Furthermore, nouns such as *fence* denote entities that are divisible in that a long fence could itself be divided up into smaller fence units, which also felicitously fall under the denotation of *fence*.

Rothstein's innovation is to introduce the notion of a counting context. Whereas mass nouns are of type $\langle e,t\rangle$, and denote sets of entities, count nouns are of type $\langle e\times k,t\rangle$ and denote *semantic atoms*, that is, atomic entities indexed to counting contexts. Grammatical counting is a context dependent operation. We count, in a particular context k, instances of the noun denotation which in that context are considered atomic instances of that noun denotation.

Count predicates are derived from number neutral 'root' predicates via a $COUNT_k$ operation. For example, suppose the root predicate *FENCE* denotes $\{f_1, f_2, f_3, f_4, f_1 \sqcup f_2, ..., f_1 \sqcup f_2 \sqcup f_3 \sqcup f_4\}$

(the set of atoms closed under mereological sum, $\sqcup$) and that there are two contexts k_1 and k_2:

$$\begin{aligned}
FENCE_{\text{count}} &= COUNT_k(FENCE) = \{\langle d,k\rangle : d \in FENCE \cap k\} \\
k_1 &= \{f_1, f_2, f_3, f_4, g_1, g_2, \ldots\} \\
k_2 &= \{f_1 \sqcup f_2 \sqcup f_3 \sqcup f_4, g_1, g_2, \ldots\} \\
COUNT_{k_1}(FENCE) &= \{\langle f_1,k_1\rangle, \langle f_2,k_1\rangle, \langle \{f_3,k_1\rangle, \langle f_4,k_1\rangle\} && \Rightarrow \textit{Four fences} \\
COUNT_{k_2}(FENCE) &= \langle f_1 \sqcup f_2 \sqcup f_3 \sqcup f_4, k_2\rangle && \Rightarrow \textit{One fence}
\end{aligned}$$

In default cases, applying $COUNT_k$ to a predicate at a context results in a disjoint set. In our analysis of kind readings for concrete nouns, we will, inspired by Rothstein (2010), also appeal to contexts as devices for yielding disjoint sets, and hence countable sets.

2.2. Landman (2011, 2016) on neat mass nouns

The central concept in Landman (2011, 2016) are sets that generate N denotations under mereological sum, $\sqcup$. These are referred to as 'generator sets' in Landman (2011) and base-sets in Landman (2016) (henceforth we mostly use terminology from Landman (2016)). Landman analyses noun denotations in terms of pairs of sets $\langle \mathbf{body}, \mathbf{base}\rangle$ (an "i-set"). Bodies are the sets that determine truth conditions for predicates. Bases generate bodies under sum. The count/mass status of a noun is determined by overlap or disjointness in the base: For i-set X, "X is count iff $\mathbf{base}(X)$ is disjoint, otherwise X is mass. (Landman, 2016: p. 8).

Of particular interest in the context of this paper, are what Landman calls *neat mass* nouns. Neat mass nouns are, approximately, those nouns otherwise referred to as object mass nouns (*furniture, kitchenware*). Neat mass nouns are mass as per the above definition, but are also neat (as opposed to mess nouns like (*water, meat, salt*): "X is neat iff $\mathbf{min}(\mathbf{base}(X))$ is disjoint and $\mathbf{min}(\mathbf{base}(X))$ generates $\mathbf{base}(X)$ under $\sqcup$, otherwise X is mess." (Landman, 2016: p. 9).

Entities in the base sets for neat mass nouns represent that which intuitively counts as 'one'. For example, the base-set for *kitchenware* would include, *inter alia*, single pestles, mortars, pans, and lids, but also pestle and mortar sums, and pan and lid sums. With such a denotation, *kitchenware* would be neat since e.g. $\mathbf{min}(\mathbf{base}(KITCHENWARE)) = \{pestle, mortar, pan, lid\}$, a disjoint set. However, *kitchenware* would be mass since e.g. $\mathbf{base}(KITCHENWARE) = \{pestle, mortar, pan, lid, pestle \sqcup mortar, pan \sqcup lid\}$, an overlapping set.

The reason Landman emphasizes overlap in base sets is that overlap makes counting go wrong. In Landman's (2011) terminology, variants, V of a set X are maximally disjoint subsets X and $V \subseteq X$ such that $\sqcup X \in^* V$. For the above example for *kitchenware*, possible variants would be:

$V_1 =$	$\{mortar, pestle, pan, lid\}$	(4 items)	$V_2 =$	$\{pestle \sqcup mortar, pan, lid\}$	(3 items)
$V_3 =$	$\{mortar, pestle, pan \sqcup lid\}$	(3 items)	$V_4 =$	$\{pestle \sqcup mortar, pan \sqcup lid\}$	(2 items)

The existence of different variants leads to multiple different answers to the question *how many?*, hence counting goes wrong (even if we can count the items in each variant, of course). The connection between Rothstein's (2010) default counting contexts and variants should be relatively transparent (however see Sutton and Filip (2016b) for an in-depth discussion). In our account, we will make use of a similar notion, albeit in the context of kind and subkinds.

2.3. Sutton and Filip (2016a): Mass Ns are saturated will the null counting context

Our analysis of object mass nouns as predicates of entities (Sutton and Filip, 2016a) is, in part, a synthesis of Landman's variants and Rothstein's counting contexts. Counting contexts are indices on interpretations of predicates. We assume a domain of counting contexts $\mathcal{C} = \{c_0, c_1, \ldots, c_n\}$ such that $c_1, ..., c_n$ are default counting contexts in the sense of Rothstein (2010), roughly Landman's variants, and c_0 is the null counting context. We define the NULL COUNTING CONTEXT c_0 to model 'context' in the sense of Landman's "can all count as one simultaneously in the same context" (Landman, 2011: pp. 34-5) (which makes counting go wrong). The interpretation of a predicate at the null counting context c_0 is the union of the interpretations of the predicate at all counting contexts (i.e. variants). Examples for specific, disjoint counting contexts and the null, overlapping, counting context are given in Figure 1.

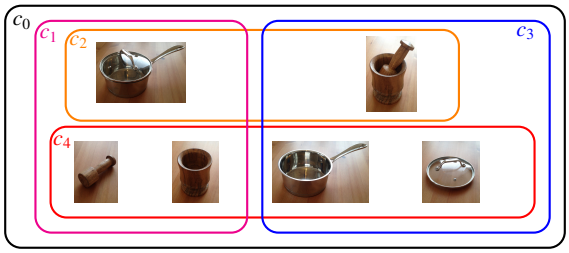

Figure 1: Null counting context c_0 and disjoint counting contexts c_1-c_4

We propose that all concrete nouns contain a context variable in their lexical entries, which means that they are of type $\langle c, \langle e,t \rangle \rangle$. However, the lexical entries of mass nouns are saturated with the null counting context, which β-reduces them to type $\langle e,t \rangle$. For example, the lexical entry for the mass noun *kitchenware* is given in (7). Although uttered in c_i, the entry is saturated with the null counting context c_0.

$$[[\text{kitchenware}]]^{c_i} = \lambda x \langle \text{K_WARE}(x), \mathbf{IND}(\text{K_WARE})(c_0)(x) \rangle \quad (7)$$

The null counting context allows for overlap in noun denotations, as c_0 in Figure 1 shows. This makes them grammatically uncountable.

The lexical entries of count nouns are NOT saturated with the null counting context. Instead, the counting context argument is filled by the specific context of use. This has the effect that, when used in a given utterance, their context argument saturated with a particular counting context. So, count nouns also β-reduce to type $\langle e,t \rangle$ in a particular context. For example, the lexical entry for the German count noun *Küchengerät* ('kitchenware', lit. kitchen.device) is given in (7). The expression is uttered and evaluated at c_i.

$$[[\text{Küchengerät}]]^{c_i} = \lambda x \langle \text{K_WARE}(x), \mathbf{IND}(\text{K_WARE})(c_i)(x) \rangle \quad (8)$$

Specific counting contexts only determine discrete, non-overlapping objects, as we see in c_1, c_2, c_3, c_4, in Figure 1. If there is any overlap present in the noun's denotation, specific counting contexts will resolve it.

A major strength of this account is that it allows us to easily account for crosslinguistic variation in count/mass lexicalization patterns (as shown in (7) and (8)). The only difference between

a mass noun like the English *kitchenware* and a count noun like the German *Küchengerät* ('kitchenware', lit. kitchen.device), is that the IND-set of individuated entities is evaluated at the null counting context c_0 in the first instance and at the counting context of utterance c_i in the second. This results in an overlapping IND-set for *kitchenware* and a disjoint IND-set for *Küchengerät*. This explains why *kitchenware* is mass, but *Küchengerät* is count.

Another advance made in Sutton and Filip (2016a) lies in the explanation of why mass-to-count coercion is blocked for object mass nouns, even if they have 'natural atoms' (ordinary individuals like chairs, tables, in the denotation of *furniture*) in their denotation. That is, we explain why the mass-to-count coercion operation cannot access 'natural atoms' that are available in their denotation.

Our analysis hinges on contrasting the implicit classifier concepts needed in cases of mass-to-count coercion (*three* [PORTIONS OF] *water*) with the role of explicit unit-extracting expressions in e.g. *piece/item of furniture*. We analyse unit-extracting classifier expressions such as *piece of*, *item of* as functions that forcibly insert the counting context of utterance into the interpretation of the whole unit-extracting phrase. In other words, such expressions shift the interpretations of nouns from being indexed at the null context (c_0) with overlapping IND-sets to being indexed to disjoint counting contexts (c_1, c_2, c_3, c_4). The effect on *kitchenware*, for example, is that *item of kitchenware* is interpreted with the same entry as the right hand side of (8), a pair with a disjoint, and so countable, IND-set. Based on these two assumptions, the reason why *#three kitchenwares* cannot be coerced to mean 'three ITEMS OF kitchenware' is that IMPLICITLY provided unit-extracting classifier concepts cannot perform the 'heavy handed' semantic operation of rewriting the null-context as the context of utterance. We must, we argue, evaluate at the null counting context provided by the lexical entry of the object noun, so overlap is not resolved. This blocks grammatical counting for object mass nouns in terms of a cardinality of a particular set of object instances. Counting goes wrong, even if the units for counting are conceptually accessible.

3. A semantics for kind readings of count and mass nouns

In this section, we argue that there are parallels between counting particular individuals and counting kinds (or subkinds). From this parallelism, we are justified in applying much of the basic architecture for analysing predicate readings of count and mass nouns, developed in Sutton and Filip (2016a) to the analysis of kind readings of count and mass nouns.

One parallel between counting particular individuals and counting kinds is that, in both cases, overlap makes counting go wrong. As we saw in Section 2, Landman's (2011; 2016) idea was that counting means non-overlap (or overlap made irrelevant). This main property behind mass/count lexicalization was adopted and built on in Sutton and Filip (2016a). We generalise this from particular individuals to kinds. First, we will argue that object mass nouns have an extensionally overlapping subkind structure, while other mass nouns do not. Then we point out that this tracks the felicity of countable subkinds in coercion environments. Finally, we provide some evidence that extensional overlap on the level of kinds prohibits felicitous counting.

By way of a working example, compare the English mass nouns *furniture* and *rice*. We do

not claim that there is, definitively, a single way to analyse the kind-subkind structure of these nouns, however, two possible (partial) representations are given in Figures 2 and 3. Subkinds of **furniture** include categories such as **bedroom furniture**, **living room furniture**, and **dining room furniture**. In turn, subkinds of these categories include basic level category kinds such as **beds**, **chairs**, and **cabinets**. As Figure 2 shows, when we look to the extensional level, such as the extension of the basic level kind **chairs**, there are entities that are chairs that count as **bedroom**, **living room**, and **dining room** furniture. But this means that the subkinds **bedroom**, **living room**, and **dining room** furniture overlap extensionally. This makes sense, given that, for example, chairs labelled as 'bedroom chairs' (say in department store catalogue) can look and function just like chairs labelled 'kitchen chairs' or 'dining room chairs'.

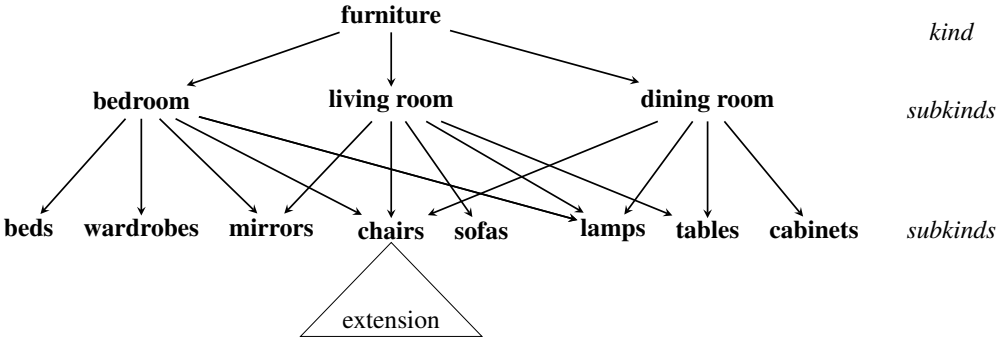

Figure 2: Kind-subkind structure of *furniture*

In contrast, subkinds of **rice** plausibly include categories such as **long grain** and **short grain**. In turn, subkinds of these categories include species kinds such as **jasmine** and **arborio**. Notice, however, that relative to one horizontal level in Figure 3, the subkinds do not overlap.[7] It will be important, as part of our fuller account to say more about the significance of such horizontal levels. We give a more precise characterisation of this in Section 3.1.

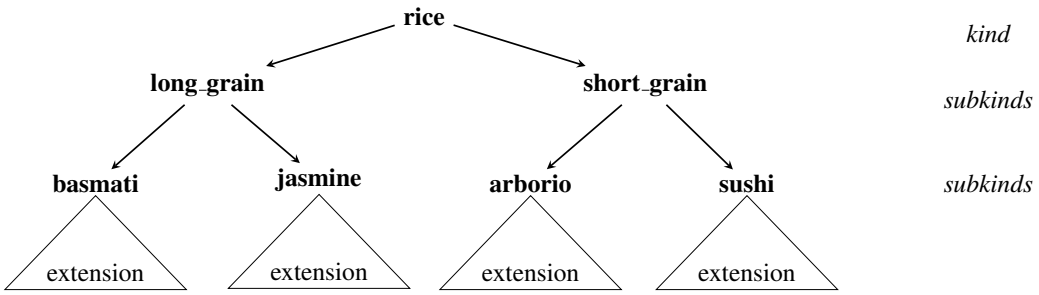

Figure 3: Kind-subkind structure of *rice*

[7]It is possible that, as with most natural language expressions, the borderline between categories may be vague. We do not rule out, therefore, that there may be, for example, some rice that is a borderline case between long grain and short grain rice. However, this is distinct from the overlapping case for furniture in which, for example, some chair could be a clear case of both a dining room chair and a bedroom chair. When we use the term 'overlapping' with respect to kinds, we therefore mean 'clearly overlapping' as is the case with the denotation of *furniture*, as opposed to possibly vaguely overlapping, as could be argued for most natural language expressions.

We hypothesise that there are two sources for the kind of overlap we see for *furniture*-like nouns but not for *rice*-like nouns. First, nouns like *rice* form a well-formed biological taxonomy, and, as such form disjoint subkinds. Nouns such as *furniture* do not denote well formed taxonomies.[8] Second, nouns such as *furniture* denote highly heterogeneous collections of artefacts. Members of collective artefact categories like **furniture**, **footwear**, **jewellery** are related through a network of 'family resemblance' style overlapping and criss-crossing relations. Furthermore, their 'vertical' kind-subkind relations are inconsistent with a well-formed taxonomy (inheritance of properties, transitivity, (see e.g., Wierzbicka, 1985; Wisniewski et al., 1996)).

These two properties (ill-formed taxonomies and heterogeneity), together, seem to make the probability of extensional overlap between kinds very high. Furthermore, we have seen from the distinction between *vehicle* and *transport* in (5a) and (5b), that the inaccessibility of countable subkind readings is restricted to a subclass of mass nouns. These factors, taken together, yield the following prediction. If a noun is mass and denotes a collection of homogenous artefacts, then there will be no felicitous mass-to-subkind reading when the noun is directly modified by a numerical or is pluralised.

In summary, we have diagnosed some properties that we think give rise to the kind of extensional overlap between kinds that we have described above. We gave some reason to associate overlap with a failure of counting entities in our discussion of (Landman, 2011, 2016): overlap makes counting go wrong because it gives rise to multiple answers to the question 'how many?'. We end this section with some evidence that non-overlap matters for counting kinds.

In English, one can explicitly access subkinds for mass nouns by using kind classifier expressions such as *kind (of)*, *type (of)*, and *sort (of)*. This is useful, because we can see that in cases where there is nonetheless overlap between explicitly referred to subkinds, counting of subkinds is only possible if any overlap is resolved. For example, (9) would be felicitous in cases where two sets of chairs were bought (one for the kitchen, one for the dining room), even if either set of chairs could have been used in the other location.

(9) I ordered two kinds of furniture: kitchen furniture and dining room furniture.

That is to say, if we ignore or remove overlap and treat the sets as disjoint, then counting is possible. However, if we make the removal of overlap impossible, we get infelicity in numerical constructions. For example, the truth-conditions of (9) exclude a case where one ordered one set of chairs that could function both as kitchen chairs and dining room chairs.

3.1. Formal Account

The main idea to be formally elaborated upon here is that counting subkinds is sensitive to KIND COUNTING CONTEXTS that remove overlap between subkinds in their constituent members. We introduce kind counting contexts as an index in the formalism, however, their precise mechanisms

[8] We will address an alternative account proposed by Grimm and Levin (2016) which focuses on the failure of taxonomic properties for artefact denoting nouns in Section 5.1.

are likely to be dependent on pragmatic speaker-hearer decisions based on intensional criteria. From the set of kind counting contexts, we define a NULL KIND COUNTING CONTEXT. The null counting context does not resolve overlap. Akin to the treatment of mass nouns in Sutton and Filip (2016a), the lexical entries of mass nouns are saturated with the null kind counting context. Count nouns will take the kind counting context of utterance.

We restrict our semantics to kind (or subkind) readings, modelled as Boolean semilatticies closed under sum. Models M are a tuple $M = \langle \mathcal{D}, \mathcal{C}, \mathcal{L} \rangle$:

$\mathcal{D}$ = $\{\mathbb{D}_e, \mathbb{D}_k, \mathbb{D}_t\}$
$\mathbb{D}_e$ = the domain of entities, a semilattice structure of atomic entities closed under $\sqcup_e$.
$\mathbb{D}_k$ = the domain of (sub)kinds, k_i, a semilattice structure of atomic subkinds closed under $\sqcup_k$.
$\mathbb{D}_t$ = $\{0, 1\}$
$\mathcal{C}$ = set of counting contexts c_i
$\mathcal{L}$ = set of levels of categorisation l_i

Other formal terminology is as standard except we mark relations between entities with a subscript e and relations between kinds with a subscript k.

$\sqcup_e$	mereological sum for entities	$\sqsubseteq_k$	part of (kinds and subkinds)
$\sqcup_k$	mereological sum for (sub)kinds	$\sqsubset_k$	proper part of (kinds and subkinds)
$\sqsubseteq_e$	part of (entities)	$a : k$	a is of (sub)kind k
$\sqsubset_e$	proper part of (entities)	$k \circ k'$	k overlaps with k'

We do not commit to a particular analysis of kinds, we assume a basic intuitive taxonomic 'kind of' relation. Notably, the mereology of kinds mirrors the mereology of entities. For example, the definition of the part of relation for kinds is *mutatis mutandis* identical to the definition of 'part of' in standard mereology (10). However, subkinds also imply relations between entities of that kind (like subtypes in type theory) as defined in (11). Because we wish to capture overlap between parts of entities as well as whole entities, we adopt Landman's (2016) use of Boolean part sets defined in (12). Extensional overlap between kinds is defined in (13).[9]

$$\forall k, k' [k \sqsubseteq_k k' \leftrightarrow k \sqcup_k k' = k'] \tag{10}$$

$$\forall k, k' [k \sqsubseteq_k k' \leftrightarrow \forall x_{\langle e \rangle} [x : k \rightarrow x : k']] \tag{11}$$

$$(x] = \{y \mid y \sqsubseteq_e x\} \tag{12}$$

$$\forall k, k' [k \circ k' \leftrightarrow \exists x, y, z_{\langle e \rangle} [y : k \wedge z : k' \wedge x \in (y] \wedge x \in (z]]] \tag{13}$$

In words, two kinds extensionally overlap if an entity or a proper part some entity that is in one is an entity or a proper part of an entity in the other.

Levels of categorisation: When it comes to counting subkinds, the requisite resolution of overlap with respect to a counting context amounts to the restriction of not simultaneously counting entities at DIFFERENT LEVELS of categorisation. For example, it is highly marked to say *I bought two kinds of furniture: tables and living room*. We semi-formally define levels of categorisation in the following way.

[9] We presume this is relative to a world but we suppress such details here.

k, k' are admissible on the same level of categorisation, iff

(i) $\neg(k \sqsubset k')$ and $\neg(k' \sqsubset k)$

(ii) k and k' form part of a natural comparison class of subkinds.

We assume a pre-theoretic notion of *natural comparison class* and give only the examples in Table 1 to demonstrate it. Examples of levels of categorisation for the kinds **rice** and **furniture** are also given in Table 1.

Table 1: Examples of levels of categorisation

Kind	level	subkinds forming a natural comparison class
rice	*colour*	**brown**, **white**
	grain length	**long_grain**, **short_grain**
	origin	**USA**, **India**
furniture	*furnishing space*	**bedroom**, **office**
	item type	**chairs**, **sofas**

Counting contexts: Kind counting contexts, c_i, are mappings from sets of subkinds (at some level) to maximally disjoint sets of subkinds (at that level). For example, there are different EXTENSIONAL RESOLUTIONS of the sub-kinds **dining_room** and **bedroom** furniture. We will call these e.g. $\textbf{dining_room}_1, ..., \textbf{dining_room}_n, \textbf{bedroom}_1, ..., \textbf{bedroom}_n$. Extensional resolutions differ, for example, with respect to whether some particular chair counts as dining room furniture or bedroom furniture. In this sense, they are different from *variants* since they have no condition of maximal disjointness or generation of the extension of the full extension under sum. Some of the resolutions from one subkind may overlap with some resolutions from the other, but some of the resolutions from one subkind may be disjoint with some resolutions from the other. Kind counting contexts resolve the overlap inherent in the subkind structure of the kind by selecting non-overlapping extensional resolutions for each subkind in a set of subkinds. Effectively, counting contexts force a choice for specific items in the domain: "Is this dining room furniture or bedroom furniture in this context?"

At the *null counting context* (c_0), subkinds may be overlapping in their constituent members and extensional overlap makes counting of subkinds go wrong. The *null counting context* (c_0) is defined in (14), relative to a counting level, in terms of the union of counting contexts $c_{i\geq 1} \in \mathcal{C}$ (all maximally disjoint subsets of X).

$$X_{c_0,l_i} = \bigcup X_{c_i,l_i} \text{ computed from all } c_{i\geq 1} \in \mathcal{C} \tag{14}$$

IND_k function: Kind predicates are interpreted relative to levels ($l_i \in \mathcal{L}$) of categorisation and counting contexts ($c_i \in \mathcal{C}$). The intuitive notion of 'counting as one' is here recast as a function IND_k. We assume that IND_k is intensional and also dependent on, *inter alia*, the purposes and goals of the communicative agents. For simplicity, we set these complications aside. IND_k accesses the subkind structure of a kind predicate (a set of subkinds $\langle k,t\rangle$) and introduces a counting context argument c, and a counting level argument l. At a context and a level the subkinds in the set count as 'one'. IND_k is of type $\langle k, \langle l, \langle c, \langle k,t\rangle\rangle\rangle\rangle$. Given a kind, a level and a context, it outputs a set of (sub-)kinds that can be counted if they do not overlap extensionally.

$$\text{IND}_k = \lambda k_1 \lambda l \lambda c \lambda k_2 [k_2(l)(c) \wedge k_2 \sqsubseteq k_1] \tag{15}$$

Table 2: Examples of subkinds for *furniture* at two different levels of categorisation, indexed for different extensional resolutions.

Level	Subkinds
l_1	$\mathbf{dining_room}_1, ..., \mathbf{dining_room}_n, \mathbf{bedroom}_1, ..., \mathbf{bedroom}_n, ...$
l_2	$\mathbf{chairs}_1, ..., \mathbf{chairs}_n, \mathbf{tables}_1, ..., \mathbf{tables}_n, ...$

Table 3: Some extensional resolutions for subkinds of *furniture*.

dining_room			**bedroom**		
$\mathbf{dining_room}_1$	$\mathbf{dining_room}_2$	$\mathbf{dining_room}_3$	$\mathbf{bedroom}_1$	$\mathbf{bedroom}_2$	$\mathbf{bedroom}_3$
|	|	|	|	|	|
extension	extension	extension	extension	extension	extension
↓	↓	↓	↓	↓	↓
$\{a,b\}$	$\{a,b,c\}$	$\{a,b,c,d\}$	$\{e,f\}$	$\{d,e,f\}$	$\{c,d,e,f\}$

Examples of IND_k, counting contexts and levels: In order to solidify how these different parts of our formal apparatus operate, we provide a toy example of how the subkinds structure of *furniture* can be broken down into different levels of categorisation and made disjoint at different counting contexts. Table 2 shows some of the subkinds for *furniture* at two different levels indexed for different extensional resolutions. Table 3 demonstrates how, relative to level l_1, the extensions of the different extensional resolutions can overlap. In this example, there are two elements, c and d, which are in the extension of more than one subkind.

There are different ways to resolve this overlap by applying kind counting contexts in order to ensure that the extensional resolutions of subkinds are non-overlapping. Three of the possible counting contexts are given in Table 4. At each counting context, one extensional resolution is chosen from each subkind (at a particular level) such that the subkinds are disjoint.

Finally, we give an example of how the IND_k function provides the argument structure for levels and counting contexts. The formulas below will form parts of lexical entries of kind readings of nouns. We give details of this in Section 3.2.

$$\text{IND}_k(\mathbf{furniture}) = \lambda l \lambda c \lambda k[k(l)(c) \wedge k \sqsubseteq \mathbf{furniture}] = \{\mathbf{dining_r}_1, ..., \mathbf{dining_r}_n, \mathbf{bedroom}_1, ..., \mathbf{bedroom}_n, \mathbf{chairs}_1 ..., \mathbf{chairs}_n, \mathbf{tables}_1 ..., \mathbf{tables}_n, \}$$

$$\text{IND}_k(\mathbf{furniture})(l_1) = \lambda c \lambda k[k(l_1)(c) \wedge k \sqsubseteq \mathbf{furniture}] = \{\mathbf{dining_r}_1, ..., \mathbf{dining_r}_n, \mathbf{bedroom}_1, ..., \mathbf{bedroom}_n\}$$

$$\text{IND}_k(\mathbf{furniture})(l_1)(c_1) = \lambda k[k(l_1)(c_1) \wedge k \sqsubseteq \mathbf{furniture}] = \{\mathbf{dining_r}_1, \mathbf{bedroom}_3\}$$
(two non-overlapping subkinds relative to level l_1 & context c_1)

Table 4: Some possible counting contexts for *furniture* at l_1.

Context	Subkind Resolution
c_1	$\{\textbf{dining_room}_1, \textbf{bedroom}_3\}$
c_2	$\{\textbf{dining_room}_2, \textbf{bedroom}_2\}$
c_3	$\{\textbf{dining_room}_3, \textbf{bedroom}_1\}$

3.2. Lexical entries for kind denoting expressions

We wish to remain neutral with respect to the relationship between kind readings and predicate readings for nouns in English. However, a reasonable working hypothesis is that there is a kind shifting operation from predicates to kinds (generally available for mass nouns, for count nouns, licensed by the bare plural noun morphology or definite article plus singular noun constructions).

Kind readings of nouns are interpreted as pairs $\langle \textbf{kind}, \textbf{counting_base} \rangle$. The **kind** is the kind for a noun (the result of a kind shifting operation). The **counting_base** is the result of the application of the IND_k function to a kind ($\text{IND}_k(\mathbf{n})$). Recall that the IND function for kinds applies to a kind **n** and yields a function from counting levels l to a function from counting contexts c, to a set of subkinds k each of which count as 'one' with respect to the kind **n**. Counting of kinds is licensed only when the counting base does not overlap extensionally. Count nouns are interpreted at the level and counting context of utterance. Mass nouns are interpreted at the level of utterance, but come saturated with the null counting context (c_0). This means that any overlap in the IND-set at a level remains and so can block the counting of subkinds.

Rice: the entry for the mass noun *rice* is given in (16). The first in the pair is the kind **rice**. The second in the pair is the set of subkinds of rice at the level l_j and the null counting context c_0.

$$[\![\text{ rice }]\!]^{c_i, l_j} = \langle \textbf{rice}, \text{IND}_k(\textbf{rice})(l_j)(c_0) \rangle \tag{16}$$

As Figure 4 helps to show, even at the null counting context, relative to a level, the subkinds of *rice* are extensionally disjoint. This means that they are available for counting. We will give an account of mass-to-subkind coercion in Section 4, however, in principle, this means that directly attaching a numerical expression to *rice* can yield a felicitous subkind reading as in (2b).

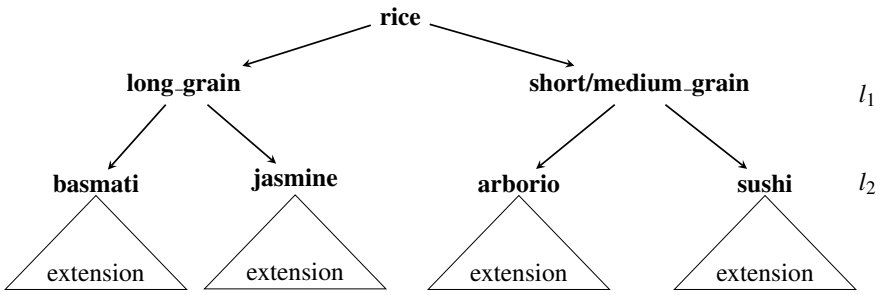

Figure 4: Subkinds at each level have *Disjoint Extensions*

Notice, also, that the subkind structure for *rice* preserves *property inheritance*: the subkind has the same properties as the superkind plus one or more additional properties. For example, Basmati rice is a type

of long-grain rice and has the properties of long-grain rice; long-grain rice is a type of rice and has the properties of rice.

Furniture: The lexical entry for the mass noun *furniture* is given in (17) and (18). The only difference between them is how the argument for level of categorisation has been filled.

$$[\![\text{ furniture }]\!]^{c_i,l_1} = \langle \textbf{furniture}, \text{IND}_k(\textbf{furniture})(l_1)(c_0) \rangle \quad (17)$$

$$[\![\text{ furniture }]\!]^{c_i,l_2} = \langle \textbf{furniture}, \text{IND}_k(\textbf{furniture})(l_2)(c_0) \rangle \quad (18)$$

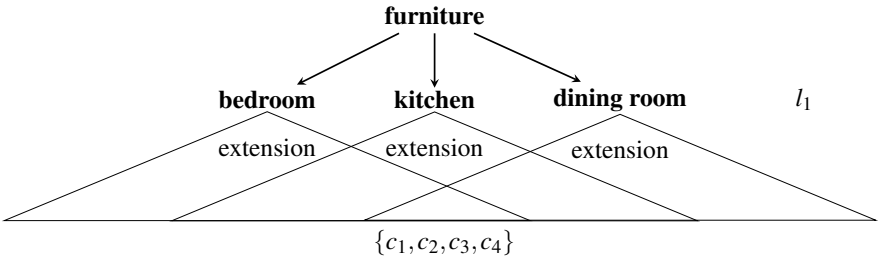

Figure 5: Subkinds at l_1 have *Overlapping Extensions*

Like *rice*, *furniture* is a mass noun and so the entry is saturated with the null counting context. Unlike *rice*, however, the subkinds for *furniture* are not extensionally disjoint at the null counting context at either level as Figures 5 and 6 help to show. At level l_1, for example, some chairs ($\{c_1, c_2, c_3, c_4\}$) count as bedroom, kitchen and dining room furniture. There are, therefore, different answers to the question *how many kinds?* which makes counting of subkinds go wrong. At level l_2, for example, the basic level category **vanities** extensionally overlaps with the basic level categories **chairs**, **tables** and **mirrors**. There are, therefore, different answers to the question *how many kinds?* which makes counting of subkinds go wrong.

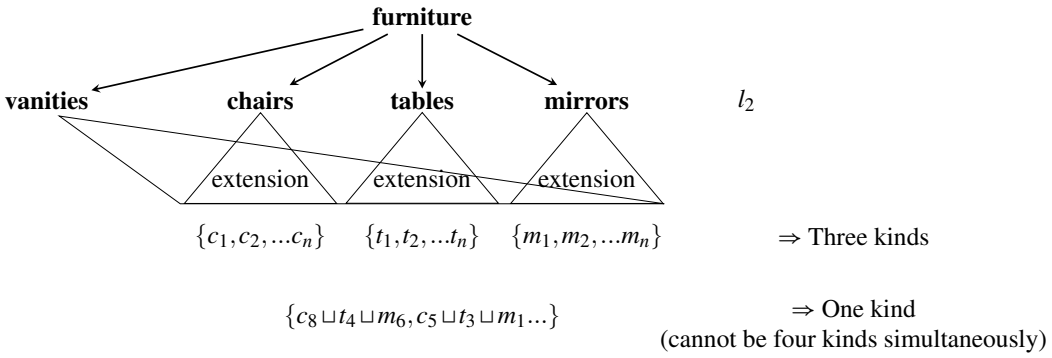

Figure 6: Subkinds at l_2 have *Overlapping Extensions*

Vehicles: The entry for the plural count noun *vehicles* is given in (19). The kind-subkind structure of this collective artefact denoting noun also has overlapping subkinds (just like *furniture*). However, as a count noun, *vehicles* is NOT saturated with the null counting context. Therefore, it is interpreted relative to the

counting context of evaluation.

$$[\![\text{ vehicles }]\!]^{c_i,l_2} = \langle \textbf{vehicle}, \text{IND}_k(\textbf{vehicle})(l_2)(c_i) \rangle \tag{19}$$

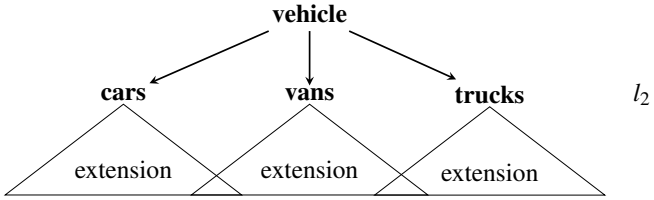

Figure 7: Subkinds at l_2 have *Overlapping Extensions*

That is to say, that although the structure of *vehicles* looks just like the structure of *furniture* in terms of overlapping subkinds, but, because the counting context of evaluation is not the null counting context, extensional overlap between subkinds at each level of categorisation is resolved. Specific counting contexts force the resolution of overlap and so, there is, in any one context, only one answer to the question '*how many kinds?*'. This makes the counting of subkinds possible.

4. Kind extracting expressions and restrictions on subkind coercion

First, in Section 4.1, we give an outline for the semantics of kind extracting classifier expressions such as *kind of*, *type of* and *sort of*. The outcome of this will be a compositional analysis of complex noun expressions such as *kind of furniture*. Our main focus will be on how such expressions compose with lexically simple mass nouns to yield something countable. Second, in Section 4.2 we will provide an analysis of mass-to-subkind cases of coercion. We will show how, combined with or account from Section 3.1, we can explain why coercion is possible for nouns such as *rice* (see example (2b)), but not for collective artefact nouns such *furniture* (see examples (4b) and (4c)).

4.1. Kind extracting expressions

Our account of the kind-subkind structure denoted by concrete nouns allows us to give a very straightforward analysis of kind extracting expressions such as *type of*, and *kind of*. On the assumption that English count nouns are standardly predicate denoting and English mass nouns are standardly kind denoting (Chierchia, 1998), first kind extracting expressions license a type shift to a kind reading for singular count nouns such as *vehicle* in *kind of vehicle* (details of this shifting operation suppressed below). However, more importantly, second, kind extracting expressions force a resolution of any extensional overlap between subkinds at some level of categorisation. Formally, we can represent this as function that introduces a further context argument into the counting base (we use the standard functions π_i, π_2 such that for an expression $\langle X, Y \rangle_{\langle a \times b \rangle}, \pi_1(\langle X, Y \rangle) = X$ and $\pi_2(\langle X, Y \rangle) = Y$):

$$[\![\text{ n}_{\text{kind reading}}]\!]^{c_i,l_i} = \langle \mathbf{n}, \text{IND}_k(\mathbf{n})(l_i)(c_i) \rangle_{\langle k \times \langle k,t \rangle \rangle} \tag{20}$$

$$[\![\text{ kind of }]\!] = \lambda \mathcal{K}_{\langle k \times \langle k,t \rangle \rangle} \lambda c \langle \pi_1(\mathcal{K}), \pi_2(\mathcal{K})(c) \rangle \tag{21}$$

Since kind counting contexts are mappings from sets of subkinds (at some level) to maximally disjoint sets of subkinds (at that level), counting contexts can be stacked. Applying a specific counting context to an overlapping set of subkinds evaluated at the null counting context, returns a disjoint set of subkinds.

For kinds denoted by mass nouns, therefore, the kind extracting expression operation, in effect, replaces the null counting context, c_0, with the counting context of utterance for the kind-extracting expression. This allows the counting of subkinds. For instance, for *kind of furniture* in (22):

$$\begin{aligned}
[\![\text{ kind of furniture }]\!]^{c_i,l_j} &= [\![\text{ kind of }]\!]^{c_i}([\![\text{ furniture }]\!]^{c_i,l_j}) \\
&= [\![\text{ kind of }]\!]([\![\text{ furniture }]\!]^{c_i,l_j})(c_i) \\
&= \lambda\mathcal{K}.\lambda c.\langle\pi_1(\mathcal{K}),\pi_2(\mathcal{K})(c)\rangle(\langle\textbf{furn},\text{IND}_k(\textbf{furn})(l_j)(c_0)\rangle)(c_i) \\
&= \lambda c.(\langle\textbf{furn},\text{IND}_k(\textbf{furn})(l_j)(c_0)(c)\rangle)(c_i) \\
&= \langle\textbf{furn},\text{IND}_k(\textbf{furn})(l_j)(c_0)(c_i)\rangle \\
&= \langle\textbf{furn},\text{IND}_k(\textbf{furn})(l_j)(c_i)\rangle
\end{aligned} \tag{22}$$

The context of evaluation for the kind extracting expression applies to the counting base and enforces the extensional resolution of overlap among subkinds in that context. This prevents multiple, and possibly inconsistent, category assignments to ordinary individuals, e.g., one chair *simultaneously* realising two different subkinds, **kitchen_furniture** and **dining_room_furniture**, for instance, *in the same context at the same time*; of course, one and the same chair can be categorised as a kitchen chair in one context, and a dining room chair in another context.

4.2. Restrictions on subkind coercion

Subkind coercion, we assume, is one possible strategy for resolving a type-mismatch such as the one created by applying a numerical expression directly to a mass noun. One way of attempting to resolve this mismatch is to access the subkind structure in the kind reading of the mass noun's lexical entry and attempt to enumerate these subkinds. However, this strategy can only be successful if the kind interpretation of the noun has a disjoint subkind structure. If it does not, the grammatical counting operation is not defined. For example, for an expression $\mathcal{K}_{\langle k\times\langle k,t\rangle\rangle}$, the type for a kind-interpretation of a noun, applying a numerical expression such as *three*, will have the result shown in (23).

$$[\![\text{ three }]\!](\mathcal{K}) = \begin{cases} \langle\pi_1(\mathcal{K}),|\pi_2(\mathcal{K})| = 3\rangle & \text{presupposing } \pi_2(\mathcal{K}) \text{ is disjoint} \\ \bot & \text{otherwise} \end{cases} \tag{23}$$

This will give different results for nouns such as *rice* compared to nouns such as *furniture*. Nouns like *rice* have an extensionally disjoint subkind structure at each level of categorisation. Therefore, if one coerces *rice* as a result of direct numerical modification, the felicitous reading will be to enumerate subkinds. This is the reading for sentences such as (2b). Nouns like *furniture* do not have an extensionally disjoint subkind structure at each level of categorisation. Therefore, if one coerces *furniture*, as a result of direct numerical modification, there will not be a felicitous reading that enumerates subkinds. This explains the infelicity of sentences such as (4b) and (4c).

Furthermore, this will give different results for mass nouns such as *transport* compared to count nouns such as *vehicle*. On the assumption that, in English, the pluralisation of count nouns licenses a kind shifting operation (Chierchia, 1998), counting of subkinds is felicitous because, although **vehicle**, like **transport**, has an overlapping subkind structure, the lexical entry for *vehicle* is not, as a count noun entry, saturated with the null counting context whereas the lexical entry for *transport*, as a mass noun entry is saturated with the null counting context. This means that in every context of utterance, the overlap between subkinds for **vehicle** is resolved by the counting context of use. On the other hand, the saturation of the lexicon with the null counting context for *transport* means that extensional overlap between subkinds is not resolved. This explains the felicity of examples such as (5a) and the infelicity of examples such as (5b).

5. Conclusions and comparisons with other accounts

5.1. An alternative account: Grimm and Levin (2016)

As far as we know, there is only one other account (Grimm and Levin, 2016) that tries to accommodate the data we have considered here. Grimm and Levin emphasise that collective artefact (*furniture*-like) nouns and their purported subordinate terms do not stand in a well formed taxonomic relation *viz.* they do not participate in the kind-subkind relations necessary to form a well-formed taxonomy. For example, in well formed taxonomies, subkinds inherit all properties from superkinds. Take the kind **apple**, the subkind **green apple** and the subsubkind **granny smith**. Green apples have all the properties of apples, and Granny Smiths have all the properties of green apples. The inheritance relation is transitive. Collective artefacts do not preserve property inheritance. The reason for this, on Grimm and Levin's (2016) analysis, derives from their claim that denotations of collective artefact nouns include (potential) associated events. For example, part of the lexical semantics of *furniture* is that items in the denotation of *furniture* potentially participate in furnishing events, or at least has the potential to serve this function. This is meant to contrast with expressions such *chair*, the denotata of which participate in sitting-on events. The result is that "$[\![$ chair $]\!] \not\subseteq [\![$ furniture $]\!]$, even though chairs (in stereotypical worlds) always satisfy the associated event of furniture." and"*Furniture*-like nouns and their constituent entities have different associated events, and therefore do not participate in the sub-/super-kind relation necessary to form an artifactual taxonomy." (Grimm and Levin, 2016). To take another example, *furniture* is for furnishing spaces and mirrors are a subkind of furniture, however, mirrors are not all for furnishing spaces. Compact make-up mirrors and telescope mirrors, for instance, do not share this function.

We have made no mention, in our account, of associated events, but we do not see any reason why there should not be an event-related component in the analysis of kinds denoted by artefact nouns. It strikes us as plausible that, if events are a part of the lexical semantics for artefact denoting nouns, then reference to event types may well underspecify extensions in a way that will give rise to extensional overlap in subkinds. In this sense, our account and Grimm and Levin's may well be compatible. Further work must be done to establish the veracity of this possibility, however.

5.2. Conclusions

The crucial difference between *furniture*- and *rice*-like nouns, on our account, is that subkinds of *furniture*-like nouns extensionally overlap with respect to their constituent subordinate entities, i.e., 'lower' subkinds and ordinary individuals. *Furniture*-like nouns do not lack subkinds entirely (*pace* Grimm and Levin (2016)), but their taxonomies are defective, and so inconsistent with well-formed taxonomies. This is not the case for prototypical mass nouns like *water* or granular nouns like *rice*, which have taxonomic subkinds whose members we view as conceptually not overlapping.

The crucial difference between *furniture*- and *vehicle*-like nouns, on our account, is that, as mass nouns, *furniture*-like nouns are indexed to the null counting context (c_0) whereas count nouns like *vehicle* are indexed to the counting context of utterance. This means that when interpreted with kind readings, despite the fact that both kinds of nouns denote extensionally overlapping subkind structures, the subkinds of *vehicle*-like nouns are countable because the counting context of utterance forces a choice which resolves this overlap. In context, we are forced to decide if some vehicles are trucks or vans. The result of this choice determines whether there are two kinds of vehicles or just one.

The indexing of mass nouns to the null counting context explains why a simple 'kind of' coercion is insufficient for object mass nouns: they have overlapping taxonomic (sub)kinds. The coercion operation consists of accessing, and attempting to enumerate, the subkind structure in the lexical entry. Counting subkinds is only then possible if the subkinds at some level are extensionally disjoint. This contrasts with the explicit use of a kind extracting expression such as *kind of* which also provides a counting context of utterance that resolves extensional overlap.

References

Chierchia, G. (1998). Reference to kinds across language. *Natural Language Semantics 6*(4), 339–405.

Grimm, S. and B. Levin (2016). Artifacts: Reference, countability, and categorization. Presentation at Heinrich-Heine-University Düsseldorf, Workshop on Countability.

Landman, F. (2011). Count nouns – mass nouns – neat nouns – mess nouns. *The Baltic International Yearbook of Cognition 6*, 1–67.

Landman, F. (2016). Iceberg semantics for count nouns and mass nouns: The evidence from portions. *The Baltic International Yearbook of Cognition 11*, 1–48.

Pelletier, F. J. (1975). Non-singular reference: Some preliminaries. *Philosophica 5*(4), 451–465.

Rothstein, S. (2010). Counting and the mass/count distinction. *Journal of Semantics 27*(3), 343–397.

Rothstein, S. (2015). Object mass nouns from a crosslinguistic perspective. Handout for the SFB 991 Colloquium.

Sutton, P. R. and H. Filip (2016a). Counting in context: Count/mass variation and restrictions on coercion in collective artifact nouns. *Semantics and Linguistic Theory (SALT) 26*, 350–370.

Sutton, P. R. and H. Filip (2016b). Mass/count variation, a mereological, two-dimensional semantics. *The Baltic International Yearbook of Cognition Logic and Communication 11*, (To appear).

Wierzbicka, A. (1985). *Lexicography and Conceptual Analysis*. Karoma.

Wisniewski, E. J., M. Imai, and L. Casey (1996). On the equivalence of superordinate concepts. *Cognition 60*, 269–298.

Zucchi, S. and M. White (1996). Twigs, sequences and the temporal constitution of predicates. In T. Galloway and J. Spence (Eds.), *Semantics and Linguistic Theory (SALT) 6*, pp. 223–270. Ithaca, NY: Linguistic Society of America.

Zucchi, S. and M. White (2001). Twigs, sequences and the temporal constitution of predicates. *Linguistics and Philosophy 24*(2), 223–270.