Baltic International Yearbook of Cognition, Logic and Communication

Volume 6 *FORMAL SEMANTICS AND PRAGMATICS. DISCOURSE, CONTEXT AND MODELS* Article 12

2011

Count Nouns - Mass Nouns, Neat Nouns - Mess Nouns

Fred Landman
Linguistics Department

Recommended Citation
Landman, Fred (2011) "Count Nouns - Mass Nouns, Neat Nouns - Mess Nouns," *Baltic International Yearbook of Cognition, Logic and Communication*: Vol. 6. https://doi.org/10.4148/biyclc.v6i0.1579

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2011 Volume 6: *Formal Semantics and Pragmatics: Discourse, Context, and Models*

pages 1-67 DOI: 10.4148/biyclc.v6i0.1579

FRED LANDMAN
Linguistics Department
Tel Aviv University

COUNT NOUNS – MASS NOUNS
NEAT NOUNS – MESS NOUNS

ABSTRACT: In this paper I propose and formalize a theory of the mass-count distinction in which the denotations of count nouns are built from non-overlapping generators, while the denotations of mass nouns are built from overlapping generators. Counting is counting of generators, and it will follow that counting is only correct on count denotations.

I will show that the theory allows two kinds of mass nouns: *mess* mass nouns with denotations built from overlapping minimal generators, and *neat* mass nouns with denotations built from overlapping generators, where the overlap is not located in the minimal generators. Prototypical mass nouns like *meat* and *mud* are of the first kind. I will argue that mass nouns like *furniture* and *kitchenware* are of the second type.

I will discuss several phenomena—all involving one way or the other explicitly or implicitly individual classifiers like *stuks* in Dutch—that show that both distinctions mass/count and mess/neat are linguistically robust. I will show in particular that nouns like *kitchenware* pattern in various ways like count nouns, and not like mess mass nouns, and that these ways naturally involve the neat structure of their denotation. I will also show that they are real mass nouns: they can involve measures in the way mess mass nouns can and count nouns cannot.

I will discuss grinding interpretations of count nouns, here rebaptized *fission* interpretations, and argue that these interpretations differ in crucial ways from the interpretations of lexical mass nouns. The paper will end with a foundational problem raised by fission interpretations, and in the course of this, atomless interpretation domains will re-enter the scene through the back door.

1. ...OR NOT TO COUNT

Count nouns, like *boy*, can be counted, mass nouns, like *salt*, cannot:

(1) a. ✓one *boy*/✓two *boys*/✓three *boys*,...
b. #one *salt*/# two *salt*/ # three *salt*,...

The standard assumption about count nouns is that the denotation of a count noun like *boys* is a structure of singularities and pluralities, where the singularities are the semantic building blocks of the structure, and we count pluralities in terms of these semantic building blocks.

Why can't we similarly count the denotations of mass nouns like *meat* and *salt*? Is it something about the building blocks of mass noun denotations, and if so, what?

In this section I discuss some answers to this question.

1.1. *We could count mass nouns if we wanted to, but we don't want to.*

Let us make the most minimal assumption: there is a lexical feature [+C] that count nouns and numerical expressions have and mass nouns don't, and felicitous combination of the numerical with the noun requires unification of this feature. This theory predicts that count nouns can be counted and mass nouns cannot.

Since the feature is not semantically interpreted, on this minimal account, the lexicon *chooses* to assign the feature [+C] to some nouns and not to others, and mass nouns cannot be counted, *simply because they are mass*.

Jef Pelletier (in many papers from Pelletier 1975 to his presentation at the Riga conference, see also Pelletier & Schubert 1989/2002) has defended an account along these lines. For Pelletier the feature [+C] is not semantically interpreted *because it shouldn't be*: mass nouns and count nouns have in essence the *same* denotations. The arbitrariness of the choice between mass nouns and count nouns is illustrated by the fact, for instance, that *spaghetti* is mass and *noodle* is count, by the existence of minimal pairs like *shoes/footware*, and by the free shiftability of nouns between mass and count uses.

While the facts about arbitrariness need to be acknowledged, it is also well-known that the arbitrariness is not absolute: languages that have the mass-count distinction tend to agree on what nouns are prototypically mass and what are prototypically count, and one can ask the same question for these classes: why don't we count prototypical mass nouns, and do we count prototypical count nouns?

Pelletier's shiftability argument aims to show that we are better off assuming that mass nouns and count nouns have the same denotation, except for a bit of contextual restriction. According to Pelletier, nouns shift *freely* between mass and count uses: mass nouns are *packaged* as count nouns, as in (2a), while count nouns are ground into mass nouns, as in (2b):

(2) a. We'd like three *waters*, please.
b. After the failed repair attempt, there was *watch* all over the table.

Pelletier assumes that the easiest account of these facts is the assumption that there is no semantic difference between the mass noun and the count noun: when the feature [+C] is assigned, as in (2a), you can count objects in the denotation of *water*, and they will be counted like objects are counted (this one, and this one, and this one), while in (2b) counting is lexically disabled, even though what there is on the table is *conceptually* countable.

I will argue later in this paper, following Rothstein (2009a), that, while grinding is indeed an operation available in the grammar, an operation that maps count noun interpretations onto mass interpretations, the output interpretations of grinding differ *semantically* in crucial ways from normal mass interpretations. I will argue that the ground interpretation of a count noun cannot be regarded as simply the same interpretation with the count feature removed and maybe some contextual restriction: grinding is a real semantic operation that

maps the input (count) meaning onto a *different* (mass) meaning.

I will argue here that the same is true for packaging. I have argued myself in Landman (1991) that packaging on noun phrases is a contextually available operation. Look at (3):

(3) Lord Peter, we have examined ***both*** *the coffee in the cup and the coffee in the pot*, and ***neither*** have strychnine in them.

Both and neither require a sum of two things. Since *coffee* is a mass noun, the coffee in the cup and the coffee in the pot sum up in the mass domain to the mass sum of coffee, which is not a sum of two things, because it is mass. We get the correct reading of (3) by packaging the coffee in the cup and packaging the coffee in the pot as two count atoms, and letting the conjunctive noun phrase denote the sum of the two packages. In this contextual shift, there are no constraints on packaging: you just treat the contextually relevant mass entities as packages.

Now look at (2a). Suppose the waiter of *Chez Jef* comes back with a tray on which stand: a scotch glass of carbonated water, a 2 liter bottle of distilled water, and a test tube of water from the canal. We wouldn't find this an appropriate reaction to our request in (2a).

Why not? On the minimal account, we have added the [+C] feature and made objects in the denotation of water available to be selected as count packages. If I say *excuse me, we asked for three waters*, the waiter can answer: *well, that's what you got*.

The inappropriateness of the waiter's reaction is naturally accounted for, if we assume that packaging of nouns is a *grammatical* construction in which an implicit classifier is added:

$$[_{NOUN}\ \textit{water}] \Rightarrow [_{NOUN[+C]}\ [_{CLASSIFIER}\ \mathrm{e}]\ [_{NOUN}\ \textit{water}]]$$

On this view, the null classifier has a contextually provided classifier meaning. Thus, a natural interpretation in context would be that $[_{CLASSIFIER}\ \mathrm{e}]$ is interpreted as *glasses of*, or (in Israel) *half liter bottles of*. Since there is no relevant contextual classifier meaning that comprises the three things that the waiter brings, his reaction is inappropriate.

On this view, packaging of nouns is a real semantic operation that maps the input meaning of the mass noun onto a *different* output meaning of the count noun.

I claim, then, that linguistic evidence suggests that grinding and packaging are not inverse operations that switch freely and without much more semantics than a bit of contextual restriction between count and mass uses of nouns. The fact that these operations have a real semantics suggests, if anything, that the meanings of the count nouns and the mass nouns are semantically different. So the question remains: why can't we count mass?

1.2. We could count mass nouns if we wanted to, but we choose not to.

We are now concerned with theories that *do* distinguish the meanings of mass nouns and count nouns semantically. We said that pluralities are counted in terms of their semantic building blocks. Following Link (1983), we can define what counts as semantic building blocks in terms of the notion of a plurality structure: we choose a particular part-of relation, *plural-part-of*, in terms of which we define counting.

Link (1983) creates a sortal distinction between count nouns and mass nouns: count nouns have a denotation in a plurality structure on which a counting operation is defined, mass nouns have a denotation in a structure which is disjoint from the count structure, a mass structure on which a counting operation is not defined.

Krifka (1989) does not make this sortal distinction, but defines different partial orders on the same domain. The interpretation domain is ordered by a partial order. For nouns that are [+C], the noun intension X and natural unit function NU determine in world w a set: $\mathrm{A}_{X,w} = \lambda x.\mathrm{NU}_{X,w}(x){=}1$, the set of objects that naturally function as units that count as 1 X in w.

A second partial order of *plural part of* is defined on a superset of $\mathrm{A}_{X,w}$ in which the elements of $\mathrm{A}_{X,w}$ are the building blocks (the minimal elements). As Krifka argues in a footnote, crucially his plural-part of order cannot be simply lifted from the general order on the domain, because the elements in $\mathrm{A}_{X,w}$ may overlap in terms of that order.

The analysis in Rothstein (2010) is similar both to Link's and to Krifka's. With Link, she assumes a typal distinction between the denotations of mass nouns and count nouns. With Krifka, she derives the count noun denotations from a counting function which assigns the value 1. Her analysis differs from Krifka's in that she assumes that the minimal elements need not be conceptually natural objects, but can be contextually selected, and conceptually rather arbitrary.

These theories distinguish mass denotations from count denotations and assume that the latter, but not the former, are interpreted relative to part-of structures for which counting is defined. If we don't go any further, we get what I would call a *we-choose-not-to* answer to the question of why we can't count mass nouns: we have two types of structures, one that comes equipped with a counting operation and one that does not, and mass nouns are interpreted in the structure that is not equipped with the counting operation. In other words, count nouns have the possibility of counting built into their meaning, while mass nouns do not.

In that case, the answer to why we can't count mass nouns is simple: because we decided not to build the possibility of counting into their meaning. Thus, we interpret the feature [±C] semantically, but rather minimally: we don't count mass nouns, because we have equipped count nouns but not mass nouns with a counter: we can choose between equipped and non-equipped interpretations of nouns, and that's all there is to it (an answer along this line is suggested in Sybesma (2009)).

Interestingly enough, neither Link, nor Krifka, nor Rothstein seem to regard this as a sufficiently insightful analysis. While separating mass noun and count noun denotations typally, Link (1983) (and Landman 1991) make further assumptions about mass structures which makes it *structurally* impossible to equip the mass domain with a counting operation (see below).

Krifka (1989) assumes that the objects that count as one in a noun denotation are selected to count as one by a *natural unit* function, and suggests that mass nouns cannot be counted for conceptual reasons: mass nouns do not come conceptually with natural units of counting, while count nouns do.

Rothstein (2010) argues against this: many count nouns are contextually count and their semantic units are not necessarily 'natural units' at all. She assumes that the count domain selects in context the objects of count 1 to be objects that are in that context mutually disjoint. (following an earlier incarnation of the present paper).

The present paper is not about count nouns: I am sympathetic to all three approaches to count nouns. However, as far as mass nouns are concerned, I share the discomfort: I find the *we-choose-no-to*-account uninsightful and practically circular, and I think we can do a bit better.

1.3. We can't count mass nouns because they have no semantic building blocks or instable semantic building blocks.

The semantic building blocks of a count noun denotation are generally assumed to be the minimal elements in that denotation. Hence counting is counting of minimal elements. A very common assumption in the earlier literature is that mass noun denotations differ from count noun denotation in that mass noun denotations are not built from minimal elements, or don't have minimal elements at all (see e.g. ter Meulen 1980, Bunt 1985, Link 1983, Landman 1991). A representative example is given by the following (almost) quote:

> "What are the minimal parts of water? Chemistry tells us that they are the water molecules. But water molecules can be counted, while water cannot be counted. This shows that natural language semantics does not incorporate the insights of chemistry in its models: in our semantic domains, the water molecules are not the minimal parts of water. In fact, the real semantic question is: is there any evidence, semantic evidence, to assume that mass entities like water are built from minimal parts at all, either from minimal parts that are water, or from minimal parts that aren't water? If there is no such semantic evidence, it is theoretically better to assume that the semantic system does not impose a requirement of minimal parts.
> Since there is no semantic evidence for minimal parts, we should assume non-atomic structures for the mass domain. That has the added bonus that we can nicely explain why we cannot count mass entities, because counting is counting of atoms." (paraphrase of Landman 1991, pp 312-313)

Chierchia 1998 challenges this view by pointing at mass nouns like *furniture* (and others discussed in Pelletier & Schubert 1989/2002): *furniture* consists of pieces of *furniture*, and just as parts of pieces of *furniture* are not necessarily themselves pieces of *furniture*, in the same way parts of pieces of *furniture* do not necessarily themselves count as *furniture*. (4a) seems equivalent to (4b), neither are made true by (4c):

(4) a. I moved the furniture around.
 b. I moved the pieces of furniture around.
 c. I switched the top drawer and the middle drawer in the dresser.

If so, it is very unattractive to assume that the denotation of *pieces of furniture* contains minimal elements, namely the individual pieces of furniture, but the denotation of *furniture* does not. The first problem, then, concerns what we could call non-prototypical mass nouns: mass nouns with naturally minimal parts.

A second problem concerns prototypical mass nouns. It is what I call the problem of *homeopathic semantics*. Look at (5):

(5) There is *salt* on the viewing plate of the microscope, one molecule's worth.
 [–C]

The observation is that the mass noun *salt* in (5) is felicitous, though intuitively, what is on the viewing plate doesn't have any parts that are themselves salt. If we assume that semantically the denotation of *salt* is divisible, and *salt* has no minimal elements, then we are forced to *invent* here an infinite structure of non-existent salt parts that are themselves in the denotation of *salt*.

I call this *homeopathic semantics*: to postulate arcane semantic structures solely to avoid counting: we "dilute" the *salt* so far that not a single molecule remains, yet semantically we continue to divide it into parts that semantically count as *salt*.

But such an approach is implausible. The real observation is that divisibility is plausible at a macro level, because at a macro level we can unproblematically divide, say, *water* into two parts that both have the right characteristics to count as water. But at a micro level, this is no longer plausible, and at some level you reach a point where what you have doesn't divide anymore into two parts that can both count as water.

So what is in the microscope is *salt*, but cannot be split into parts that are themselves *salt*, hence, what is in the microscope is a pretty good candidate for a minimal *salt* part.

The micro level doesn't have to be this small. Look at the following picture: It's a piece of wallpaper, of the kind that I would call in

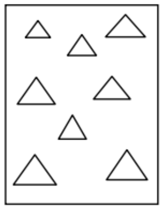

Dutch *driehoekjesbehang, triangle-patterned wallpaper*, wallpaper with little triangles.

Now, there is a sense in which any part of a piece of triangle-patterned wallpaper can be called triangle-patterned wallpaper, even if it doesn't have the pattern on it (i.e. a piece that was cut out of a role of triangle-patterned wallpaper). But there is another sense, and that is the one I am interested in here, in which in order for a piece of wallpaper to count as *triangle-patterned wallpaper*, it must contain the pattern, i.e. a triangle. In this sense, if I cut a circle out of the triangle patterned wallpaper above as in A, I wouldn't call the piece I have cut out *driehoekjesbehang*, but if I cut it as in B, I would: On

A: NO

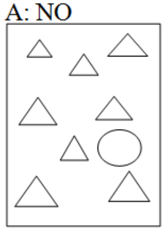

B: YES

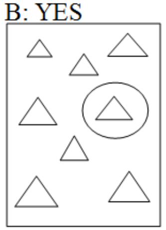

this interpretation, the circular piece in B can no longer be cut into two parts, each of which counts itself as *driehoekjesbehang*. And this piece can be part of a *partition* of the piece of wallpaper into parts that all count as *driehoekjesbehang*, but cannot themselves be split into two parts that both count as *driehoekjesbehang*. And these parts are good candidates for (contextually provided) minimal parts. Note that

B: YES

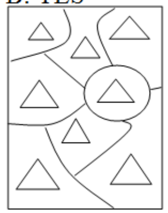

the division into *two* parts that *both* count as *driehoekjesbehang* is important here, because we are concerned with *count*. If we take one piece of *driehoekjesbehang* we can often divide that into two pieces, *one* of which continues to count as *driehoekjesbehang*, while the other does not. I call this *shaving*: the piece that still counts as *diehoekjesbehang* after you have cut off a snippet counts as *the same piece* of *driehoekjesbehang*, just as a clean-shaven version of me counts as the same person. This means that I assume that the division into minimal parts is a division *up to intensional identity* (the relation that makes the shaved version and the unshaved version of me count as *one*). I will refrain from developing this part of the theory, but I do assume that an appropriate intensional identity relation has to be added.

Examples like this can be multiplied for mass nouns denoting patterned materials. They show that the idea that elements in the denotation of a mass noun can always be split into parts that are also in the denotation of the mass noun is unwarranted.

In many respects the theory developed in Chierchia (2010) moves back from the proposals in Chierchia (1998) towards the homeopathic theory. Chierchia (2010) assumes that you cannot count mass nouns because the minimal elements in the denotation of the mass noun can be, what he calls, *instable*, where an element in a denotation N is instable if it is vague whether the element should count as one or as two.

Obviously if you don't know of your minimal elements whether they count as one or as two or as many, you can't count them.

Since Chierchia (1998) argues that this situation is not what goes on in the denotation of nouns like furniture, in the recent paper, Chierchia pooh-poohs the importance of mass nouns like furniture, moving to a position that, after all, these aren't really 'real' mass nouns.

I find this move disappointing, and will argue later in this paper that *furniture* nouns really are 'real' mass nouns (though 'neat' ones). Neither am I charmed by the analysis in terms of instable minimal elements, because I think that, on closer view, this is just a variant of the homeopathic theory: while the theory doesn't assume that you can continue to divide elements in a mass noun denotation infinitely, i.e. smaller and smaller, it seems to assume that you *can* divide them *indefinitely*, smaller and smaller, approaching but never surely reaching the 'true' non-vague minimal elements.

The problem is that this theory is also homeopathic. The cases of *salt in the microscope* and *triangle-patterned wallpaper* are as problematic for Chierchia's later theory as they are for the 'no minimal elements' theory: what there is in the microscope is not an instable element in Chierchia's sense. In fact, put *two* molecules in the microscope: in that case you can partition the salt into two parts, neither of which can be partitioned into salt. Also in this case, there is *salt* in the microscope, but not instable salt in Chierchia's sense: the two-molecule structure can only be bi-partitioned into parts that count as salt that cannot themselves be bi-partitioned into salt. The instability that Chierchia requires for mass noun denotations ('don't know whether it's one or two') is absent.

The same is true in the case of *driehoekjesbehang*: none of the parts in the partition indicated for the triangle-patterned wallpaper are instable in Chierchia's sense, because they can only count as one. Yet, we cannot count *triangle-patterned wallpaper*.

I think that the homeopathic account, whether in the classical form or in the form of Chierchia (2010), is untenable.

1.4. We can't count mass nouns because they have vague building blocks.

This is suggested by Chierchia (1998). Chierchia (2010) can be regarded as his way of making this suggestion precise in such a way that

the vagueness involved distinguishes mass nouns from count nouns: mass nouns have instable minimal elements, while count nouns have stable minimal elements. I am not denying that Chierchia's notion of instability may be a useful notion. I do not think though that it can be used to distinguish mass nouns from count nouns, and I do not think that it can be used to explain why mass nouns cannot be counted.

As Chierchia (2010) realizes very well, other notions of vagueness discussed in the vagueness literature are patently not notions that tell mass nouns apart from count nouns, and hence cannot be used to distinguish the two.

-Cardinal vagueness

Look at the examples in (6):

(6) a. How many *quarks* are there in the water in the sea?
[+C]
b. #There is more than two *water* in the sea.
[-C]

We don't know how many quarks there are in the water, and the number may even be truly undetermined (because of quantum mechanics). But that doesn't prevent *quark* from being count, and (6a) from being felicitous. We don't know how many minimal parts of water there are in the sea, but arguably, whatever the number, it's more than two: if we divide the water naturally into two parts that are water, normally these will divide themselves into parts that are water. So the statement in (2b) should be true in a natural context; but, of course, that doesn't make it felicitous.

-Borderline vagueness Maybe we can assume that the denotation of mass nouns like *salt* is generated from building blocks that are not salt, nor non-salt, but *borderline* salt. The problem is that in the nominal domain, borderline vagueness is typically found with *classifier nouns, count nouns that include a quantitative size dimension* in their meaning, like *grain* and *heap*:

-you have to have the right size to be a grain, and the right size to be a heap, and what is the right size is vague, precisely because the meanings of these count nouns involves a *quantitative dimension*, and this is what brings in problems like the Sorites paradox. But prototypical mass nouns like *water* and *salt* are not vague in this sense.

-Higher-order vagueness
On accounts of higher order vagueness, it is not the set of atoms which is vague, but the whole part-of structure itself. Such an account needs to be formalized, of course.

Chierchia's (2010) account can be understood as an analysis in this spirit, and so can the analysis that I will present here (though I don't think of it in terms of vagueness myself).

So I have nothing as such against the idea that we cannot count mass noun denotations because they involve higher order vagueness, and count nouns do not, since it may well be possible to reformulate my account in those terms.

1.5. We cannot count mass noun denotations because we cannot pull the semantic building blocks out of the mass noun denotation.

In the analysis of Chierchia (1998), count nouns have access to the set of building blocks, while mass nouns do not. The account can be illustrated with the Dutch triple in (7):

(7)

meubel	*meubels*	*meubilair*
piece of furniture	pieces of furniture	furniture
singular count	plural count	mass

For the count noun, Chierchia follows Hoeksema's (1983) account, in which the singular noun denotes a set of atoms, the plural noun denotes the closure of that set under sum *minus the set of atoms*. The mass noun, for Chierchia, denotes the union of the two, i.e. just the closure of the singular noun under sum:

$$*P = \{y\colon \exists X \subseteq P\colon y = \sqcup X\}$$

(8)

meubel ⟶	MEUBEL	a set of atoms
meubels ⟶	∗MEUBEL – MEUBEL	
meubilair ⟶	∗MEUBEL	

This semantic choice was always unfortunate, in the light of the well known problems that Hoeksema's analysis of plurals faces (as discussed in Lasersohn (1988), Rothstein (1992), summarized in Landman (2000)). But the main idea of Chierchia's analysis can be maintained without having to rely on Hoeksema's account. The essence of the analysis

is that the interpretations of the singular and plural count noun are derived from a lexically provided set of atoms in terms of which we count, and the interpretation of the mass noun is not. The access to this set of atoms is preserved in the count noun interpretations, and the interpretations of numericals refer to it.

We can easily deal with this, if we make our semantic representations a bit richer. Let the noun meaning be a pair, of which the first element is the standard interpretation, and the second element the Chierchia set, which is lexically provided, and semantically accessed by the interpretations of the numericals. Thus:

(9) *meubel* ⟶ <MEUBEL, MEUBEL> with MEUBEL a set of atoms
meubels ⟶ <∗MEUBEL, MEUBEL>
meubilair ⟶ <∗MEUBEL, – >

On this account, the plural and the mass noun have the same interpretation, but only the first allows access to the set of atoms MEUBEL, the second element of the pair.

As we will see below, my own proposal for the semantics of nouns like *furniture* is very close to this, except that in the theory to be developed, the second element of the pair plays a different and more fundamental role.

On Chierchia's theory, the difference between the plural *meubels* and the mass noun *meubilair* lies only in the lexical access to the set of atoms, there is no difference in structure. Since the theory is a general theory of mass nouns, mass nouns like *salt* and *mud* are assumed to have the same kinds of denotations:

(10) *salt* ⟶ <∗MIN-SALT, – > where MIN-SALT is a set of atoms

Chierchia assumes that the set MIN-SALT is vague, unlike the set of minimal elements of count nouns. As we have seen, it isn't clear what notion of vagueness would be at stake here (note that Chierchia's later theory is not a modification of the theory under discussion here, but a rejection of it.)

As I see it, the main problem with the theory of Chierchia (1998) is that the mass and the count denotations are so close that one seriously wonders why languages have the mass-count distinction at all.

We encode lexical access to the set of atoms only in count nouns, and hence mass nouns cannot be counted. But why don't they shift in context, when the set of atoms is made salient: The language easily allows me to package water into macro packages, but not into minimal water parts:

(11) a. I would like two coffees, two cognacs and two *waters*, please.
b. #There are far more than a billion *waters* in this cup of water.

It seems that, if there isn't any deeper reason why we cannot count mass nouns than Chierchia gives, the packaging in (11b) ought to be the most obvious one available. But it isn't. And is it just an imperfection of mass nouns that the set of atoms is not available? Why do languages bother distinguishing mass nouns and count nouns?

A set of atoms is sitting at the bottom of the mass noun denotation and at the bottom of the count noun denotation. The theory postulates that it can be pulled out in the second case, but not in the first case, and this is why you can't count. The problem is that it is not particularly difficult to semantically or contextually pull a set of atoms out of an atomic structure.... a child can do it. And there, of course, is the problem: the child doesn't do it.

2. VARIANTS

2.1. Counting and non-overlap.

All theories of count nouns that define counting in terms of a partial order take care to distinguish the order relative to which counting takes place from the partial order these things stand in in the mass domain. In all these theories, there is a counting function that will count a plurality in count denotation X in terms of its parts in X that count as one.

For prototypical count nouns, these parts that count as one will not overlap in the mass sense either, they will have no part in common: i.e. prototypically the denotations of *boy*, *soccer ball* and *planet* are sets of the elements that are mutually disjoint. But mass overlap is, of course, not impossible: my two hands and my ten fingers are objects in the count domain, and there are predicates like body part that may include them all. But the way counting works is that, if we can count entities simultaneously as one, the partial order in terms of which the counting is done, starts from the elements that count as one, and *ignores their potential mass overlap*, i.e. the count domain treats these entities *as if* they do not overlap.

As we said, prototypically count nouns have minimal elements that do not overlap in the first place. As Rothstein (2010) argues, nouns that include overlapping entities normally restrict their denotation to eliminate the overlap. Thus, we may count a fence structure put up by four farmers as one fence, or as four fences, but not normally as five.

There are situations where the overlap is not eliminated. Recently, I ordered a set of Krifka-Rothstein Outfits For All Occasions (cf. Krifka (2009)):

1. The pants and the shirt (for informal meetings)
2. The pants and the shirt and the tie (for informal meetings with Europeans)
3. The pants and the shift and the jacket (for formal meetings)
4. The pants and the shirt and the tie and the jacket (for formal meetings with Europeans)
5. The pants and the shirt and the tie and the jacket and the vest (in case I get invited to dine at High Table)

And, I have a fitting *kipah*, *yarmulka*, in case any of these occasions involves a religious ceremony, which makes all together 10 outfits (in fact, there are more combinations, but I don't have occasions for them).

As Krifka (2009) argues, these outfits are intensional entities, in that they do not all simultaneously exist in one and the same situation. And in fact, this is shown in the following counting situation:

(12) *Customs officer:* What's in the suitcase?
Me: My Krifka-Rothstein outfits.
Customs officer: How many outfits?
Me: Ten.
Customs officer: I am sorry Sir, custom regulations are that you can only bring five outfits into the country.
Me: Ok, leave out the *kipah*.

(12) is, of course, unnatural: the counting has at most joke status. Real counting is what we find in (13):

(13) There will be no religious ceremonies this trip, so only five outfits are relevant.

For the arithmetic to come out correctly here, we must count each outfit as one, and *ignore the mass overlap*. We do this either by packaging—treating each sum of clothing as an atomic individual in its own right (following Link 1984)—or by defining a new count part-of relation on the sums of clothing, which too treats the mass overlap as irrelevant (following Krifka 1989). And we *must* do this, because, as Krifka 1989 stresses, counting is an additive measure in that 1+1=2 only holds for 1's that do not overlap, and in the count domain 1+1 is indeed 2.

So, we all agree, then, that count means non-overlap, or overlap made irrelevant. If so, maybe the problem with counting in the mass domain is *overlap*, or *overlap not made irrelevant*. This is the underlying idea of the present analysis.

2.2. Variants

All the proposals discussed so far can be seen as being formulated one way or other in terms of *underspecification*:
-Mass is mass because it isn't specified as count.
-Mass is mass because it isn't equipped with a counting function
-Mass is mass because looking down in a mass denotation, you don't see any building blocks.
-Mass is mass because looking down in a mass denotation, you don't see the building blocks clearly.
-Mass is mass because you see the building blocks all right, but cannot pull them out.
My proposal is formulated in terms of *overspecification*: I propose that when you look down in a mass denotation you see *too many building blocks*. And hence, when you count building blocks in a mass denotation, *you will count them wrong*. We will take our inspiration from the following example. The picture shows a body of water, and sentence (14), with mass noun *salt*, is felicitous and true:

(14) There is *salt*[–C] in the water, two molecules worth.

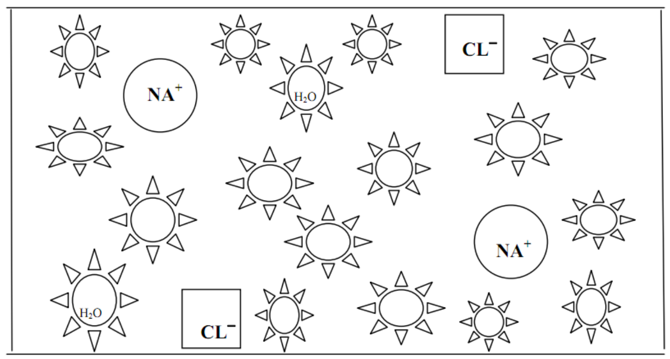

Now, there is two molecules worth of salt in the water. But *which* two molecules?
$SALT_1+SALT_2$ or $SALT_3+SALT_4$?

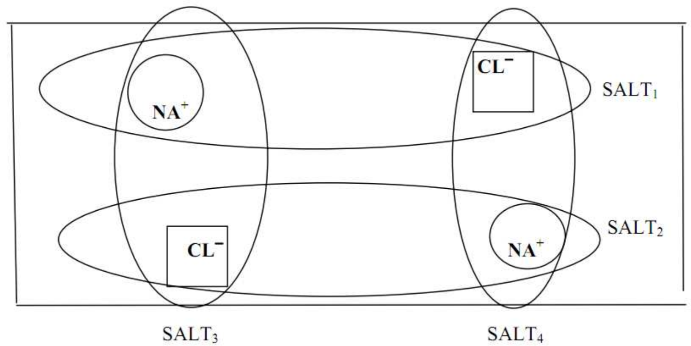

On the perspective on which we count, we have two *variants* of salt each with two non-overlapping building blocks (in the example, the molecules): $SALT_1+SALT_2$ *versus* $SALT_3+SALT_4$. For counting we *choose* one of these variants, and we count relative to it.
I am proposing here that for mass noun denotations we do not make the choice between these variants: as far as the mass denotation of *salt* is concerned, it is equally appropriate to regard the salt as being built from $SALT_1+SALT_2$ as it is to regard it as being built from

$SALT_3+SALT_4$, and in fact, we don't make the choice and *regard the salt as built, simultaneously if you want, from both variants*. Thus, the mass perspective *merges* all variants into one part-of structure, so to say scrambles them and gives (in the example) four *overlapping* building blocks.

We assume that counting is counting of semantic building blocks. If you insist on counting the building blocks in the denotation of the mass noun *salt*, you will count overlapping building blocks (four, in the example), *and you are guaranteed to count wrong!*
This is the proposal of the present paper:

> *The denotations of mass nouns cannot be counted, because counting goes wrong!*

On this proposal, the reason you cannot count prototypical mass noun denotations is not 'vertical': it's not that when you look down you see nothing, or nothing very well. The reason is 'horizontal': when you look around you at the other building blocks, you see a multitude of overlapping building blocks coming from different variants.

In general, we get variants by dividing objects into parts in different ways, without making a choice between these different ways of division. We took the case of salt dissolved in water as our model. But many other cases come to mind. The unit structure of a crystal like *diamond* forms a lattice structure. But the structure is part of a larger lattice structure and there is more than one way of partitioning the crystal into its crystal units: Division A is not more 'real' than division B: What about, say, *gold*, which in its metal state is neatly built up from gold atoms? Where are the variants? If you insist, I will maintain that each *gold* atom in your ring is built from 79 nucleons and 79 electrons, but for each gold atom, one of its electrons wanders freely through your ring. Now, *with which electron* does each gold atom form a gold atom?
Chemistry, I think, doesn't care, since chemistry doesn't really count gold-atoms, it *measures* how much gold there is.

I will not go down further on the path of speculating how we get variants given various chemical substances. I take the model as an *inspiration* for the semantics of mass nouns, rather than as a straitjacket to fit chemistry into. A better picture of the semantics of prototypical mass nouns is the following.

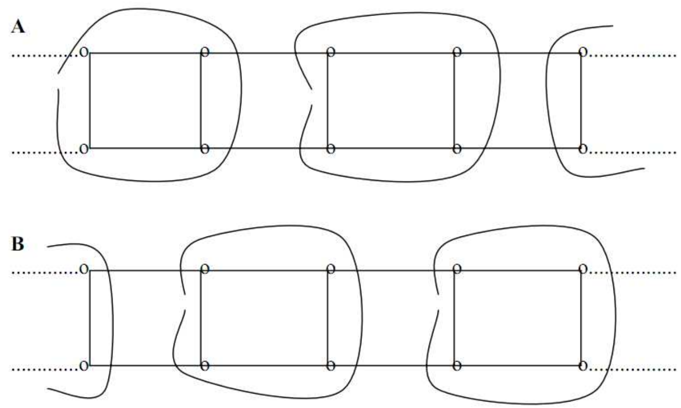

Take a big juicy slab of meat. With Chierchia (1998), I think that we can think of this as being built from minimal parts. Not natural meat-parts, but minimal parts that are appropriately minimal in a context. For instance, they are the pieces as small as a skilled butcher, or our special finegrained meat-cutting machine can cut them. Suppose the meat cutting machine consists of two sharp knife-lattices that cut the meat from left to right, and then from front to back, snap-snap. This will cut the meat into very many minimal meat pieces.

But if I move the knife-lattices slightly, the front-back knife to the left, the left-right knife to the front, and cut snap-snap, I get a different partition into minimal meat pieces. And of course, there are many ways of moving the knives. All these partitions cut into pieces which, in context, can count as minimal meat pieces. None of these partitions has a privileged status, and none of these partitions provides *its* minimal pieces with the privileged status of being the 'real' minimal pieces. On my view, all of these pieces count equally as minimal meat pieces in the context given, and the meat is built from all of them.

Similarly for the case of *driehoekjesbehang* discussed above. We gave one partition into minimal pieces of *driehoekjesbehang* above, but, of course, there are many other such partitions. Since it is a partition, each partition will consist of one square and some space. Since there is enough space that needs to be divided up, many partitions exist, and hence *driehoekjesbehang* is built from minimal pieces of *driehoekjesbe-*

hang, many of which overlap.

In fact, I think this case is instructive as a model even for mass nouns like *water*. When we think of partitioning a body of water, up to water molecules, we may be inclined to regard the structure as consisting of non-spatio-temporally realized Mickey Mouse molecules: And

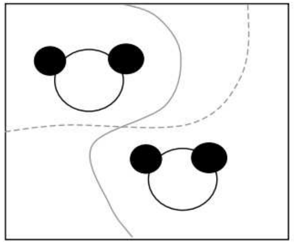

the two minimal elements are two molecules. But *the space between the molecules* is part of the body of water and shouldn't be ignored. Which means that here too we can argue that a minimal element in the denotation of the mass noun *water* will be something that consists of some essential structure (a Mickey Mouse) and some space. And again, there are many ways of dividing the space, and hence, many ways of partitioning the water into minimal mass-parts.

In sum: I propose that mass noun denotations are built from overlapping building blocks coming from a multiplicity of simultaneous variants, different ways of dividing the stuff into minimal parts. Count noun denotations, on the other hand, are built from building blocks that are, or are made, non-overlapping, denotations that form a single variant.

3. REGULAR SETS

We want to build a theory of mass and count noun denotations that generalizes the standard Boolean semantics for count nouns to include mass nouns, based on the idea that mass noun denotations are built from simultaneous variants. In this, we want to stay as close to the Boolean semantics as we can.

We start by spelling out a list of standard notions.

We assume that the domain in which mass nouns and count nouns are interpreted forms a complete atomic Boolean algebra **BOOL** = <BOOL, $\sqsubseteq, \neg, \sqcup, \sqcap$, 0,1>.

Let X ⊆ BOOL.
$^*X = \{b \in BOOL: \exists Y \subseteq X{:}b = \sqcup Y\}$

Let x,y ∈ BOOL–{0}
x and y are *disjoint* iff $x \sqcap y=0$
x and y *overlap* iff $x \sqcap y \neq 0$

Let X ⊆ BOOL–{0} X is *disjoint* iff ∀x,y ∈ X: x and y are disjoint
X *overlaps* iff X is not disjoint.

Two (non-zero) elements overlap if they have a non-zero part in common ($x \sqcap y$), otherwise they are disjoint. A set is disjoint if any two elements in it are disjoint.

X is maximally disjoint in Y iff X is disjoint and X ⊆ Y and
for every Z ⊆ Y: if Z is disjoint and Z ⊇ X then X=Z

X is *maximally disjoint* in Y if X is a disjoint subset of Y and adding any more elements of Y to X makes X overlap.

x is a *minimal element* of X iff x ∈ X–{0} and
for every y ∈ X–{0}: if $y \sqsubseteq x$ then y = x
min(X) is the set of minimal elements of X.

A *generating set for* X is a set **gen**(X) ⊆ X–{0} such that:
∀x ∈ X: ∃Y ⊆ **gen**(X): x = ⊔Y

Generating here means *generating under complete sum.*
If **gen**(X) is a generating set of X, then every element of X is generated as the (complete) sum of elements in **gen**(X).

The following facts are important for our purposes:
-If 0 ∈ X, 0 is generated by any set **gen**(X), since generation is under complete sum, ∅ is a subset of every set, and ⊔∅ = 0.
-If **gen**(X) is a generating set of X, then **min**(X) ⊆ **gen**(X). This is because generation is under sum, and **gen**(X) ⊆ X. Minimal elements

in X can only be generated under $\sqcup$ from **gen**(X) by being already in **gen**(X).
-But sets can have more than one set of generators. If X is itself a Boolean algebra, **min**(X) is a generating set for X, and hence so is any set Y such that **min**(X) $\subseteq$ Y $\subseteq$ X, including X itself.

So far all notions introduced are completely standard. We now use these to introduce the notions we are after.

A *generated set* is a pair $X = $ <X,**gen**(X)>, with **gen**(X) a generating set for X.

In the theory to be developed, the denotations of lexical nouns are going to be generated sets.

Standard notions are lifted to generated sets in the obvious way, For instance:

Let $X = $ <X,**gen**(X)>be a generated set
X is *bounded* iff X is bounded
X is bounded iff 0, $\sqcup$X $\in$ X

We now define the notion of a variant.

Let $X = $ <X,**gen**(X)>be a bounded generated set.

V is a *variant for X* iff
1. V is a maximally disjoint subset of **gen**(X)
2. *V is a subset of X such that $\sqcup$X $\in$ *V

A variant for X is a maximally disjoint subset of **gen**(X) whose closure generates elements of X, including the top element $\sqcup$X.
The closure $*$V of variant V for X is a Boolean algebra with $\sqcup$X as maximal element and V as the atoms. This means then that for each variant V for X, $\sqcup$X is generated as $\sqcup$V.

X is *generated by variants* iff
1. For every x $\in$ X there is some variant V for X such that x $\in$ *V
2. Every disjoint subset of **gen**(X) is part of some variant for X.

A generated set X is generated by variants if every element of X is generated as the sum of atoms in some Boolean algebra $*$V, with V a variant for X.
The second condition says that every disjoint subset can be extended to a variant. This condition guarantees that if Y is a disjoint subset of **gen**(X), then $\sqcup$Y $\in$ X.
Namely, by the second condition, Y is part of some variant V. This means that $\sqcup$Y $\in$ $*$V, and hence, by the definition of variant, $\sqcup$Y $\in$ X.

We now define the notions of Boolean parts of b, X-parts of b, minimal X-parts of b, X-generators of b:

Let b $\in$ BOOL:

(b] = {y $\in$ BOOL: y $\sqsubseteq$x}	the *Boolean part set* of b (also called the ideal generated by b.)
Let b $\in$ X.	
$\mathbf{ps}_X$(b) = (b] $\cap$ X	the *X-part set of* b is the intersection of the Boolean part set of b with X
$\mathbf{min}_X$(b) = **min**($\mathbf{ps}_X$(b))	the set of minimal X-parts of b

Let X be a generated set, b $\in$ X.

$\mathbf{gen}_X$(b) = $\mathbf{ps}_X$(b) $\cap$ **gen**(X)	the set of generators of b in X is the set of X-parts of b that are in the set of generators for X.
$\mathbf{ps}_X$(b) = <$\mathbf{ps}_X$(b), $\mathbf{gen}_X$(b)>	the generated X-part set of b is the pair consisting of the X-part set of b and the set of generators of b in X.

With this, we define a notion that is a bit stronger than the notion of a set generated by variants: a set closed under variants:

X is *closed under variants* iff for every b $\in$ X: $\mathbf{ps}_X$(b) is generated by variants.

So, if X is closed under variants it is not just X (= $\mathbf{ps}_X$($\sqcup$X)) that is generated by variants, but the generated X-part set of every element in X is as well.

Next we introduce the notion of Boolean relative complement:

Let x,z $\in$ BOOL and x $\sqsubseteq$ z

$\neg_z x = \sqcup\{y \in (z]\colon x \sqcap y = 0\}$ The relative complement of x in z is the sum of all the Boolean parts of z that do not overlap x.

Let X $\subseteq$ BOOL
X is *relatively complemented* iff for every x and z in X:
if x $\sqsubseteq$ z then $\neg_z x \in$ X.

This means that for every b $\in$ X, $\mathbf{ps}_X$(b) is closed under relative complement.
With these notions we define the notion of a *regular* set:

Let X be a bounded generated set.

X is *regular* iff *X* is *closed under variants* and *X* is *relatively complemented*.

We impose the following interpretation constraint on lexical nouns:

Constraint on lexical nouns:
Mass nouns and count nouns denote regular sets.

(More precisely, *plural* count nouns denote regular sets. Singular count nouns denote sets <V,V> such that <*V,V> is a regular set.)

Fact 1: If B is a complete atomic Boolean algebra with set of atoms ATOM_B, then B = <B,ATOM_B> is a regular set.

Fact 2: Let S = NA $\cup$ CL be a disjoint subset of BOOL.
Let SALT = $\{b \in S\colon |\mathbf{min}_{*S}(b) \cup \text{NA}| = |\mathbf{min}_{*S}(b) \cup \text{CL}|\}$
(SALT is the set of those sums of NA and CL elements, that are built from as many NA-elements as CL-elements.)
Then *SALT* = <SALT,**min**(SALT)> is a regular set.

Regular sets are meant to be generalizations of Boolean algebras that stay as close to Boolean algebras as is possible.

The idea is that the denotation of a count noun is generated from a *single variant*, a set of non-overlapping elements. The mass noun denotation is a simultaneous multiplicity of such variants, each a Boolean algebra which represents a different way of partitioning the *same* stuff (i.e. with the same supremum). These Boolean algebras are scrambled together into a regular set, collecting the variants together in one set of generators. This means that the set of generators is going to contain mutually overlapping elements, since the variants represent different partitions of the same stuff.

The guiding intuition about the set of generators, **gen**(X), of regular set *X* is that it is the set of *semantic building blocks*. And these are the things that we would want to count as *one*.

4. THE BOOLEAN INTUITIONS

Regular sets generalize Boolean algebras. Regular sets are not always Boolean algebras. The question is: if we move away from Boolean algebras, aren't we giving up on Boolean properties that motivated the Boolean approach to the semantics of count nouns in the first place? I will make a few remarks here.

1. Cumulativity.
If noun denotations are Boolean algebras, then they are cumulative, closed under sum. The validity of cumulativity for mass nouns and plural nouns has been a motivating principle for the Boolean approach:

-If x and y are salt. then x $\sqcup$ y is salt.
-If x and y are horses, then x $\sqcup$y are horses.

By moving to regular sets, it may seem that we are giving up on cumulativity, since cumulativity is not valid for regular sets in general, because, unlike Boolean algebras, regular sets are not necessarily closed under sum. Counterexamples can be found in the set SALT defined above:

Let Na $\in$ NA and Cl_1, Cl_2 $\in$ CL.

Then Na $\sqcup$ Cl_1 $\in$ SALT and Na $\sqcup$ Cl_2 $\in$ SALT (since the amount of Na and Cl is the same).

However, $(Na \sqcup Cl_1) \sqcup (Na \sqcup Cl_2) = Na \sqcup Cl_1 \sqcup Cl_2 \notin SALT$, since the amount of Na and Cl is not the same.

The observation is: cumulativity is not valid for salt with overlapping building blocks. To which I add: *and it shouldn't be!*
Regular sets *do* satisfy the form of cumulativity that *is* intuitively valid (cf. Krifka 1989):

-If x and y are salt *and x and y are disjoint* then $x \sqcup y$ is salt.

Namely: if $x, y \in SALT$, and $x \sqcap y = 0$, then x is generated by a disjoint subset V_x of **gen**(SALT) and y is generated by a disjoint subset Vy of gen(SALT), and since x and y are disjoint, $V_x \cup V_y$ is disjoint.
Since SALT is a regular set, $V_x \cup V_y$ is part of a variant for SALT, and hence
$\sqcup(V_x \cup V_y) \in SALT$, which is $x \sqcup y$.

2. Remainder.
Noun denotations are closed under remainder:

-Take some, but not all of the salt away, there is something left, and what is left is salt.

This principle, of course, stays valid for lexical nouns, because what is left is the relative complement, and regular sets are closed under relative complement.

What about complex noun phrases? Lønning (1987) assumes that both nouns and adjectives denote Boolean part-of sets. Since $(x] \cap (y] = (x \sqcap y]$, intersecting a noun with an adjective automatically gives you a set which is itself a Boolean part set.

I am actually not following Lønning here even for count nouns: I assume that a count noun is generated by a disjoint set, but I am not requiring this set to be a set of atoms in **BOOL** (in this respect I am following Krifka (1989), rather than Link (1983) or Landman (1989)).

But what about intersective adjectives? Shouldn't they be Boolean? The answer is that we need to look at the semantics in each particular case, and if we do so, we can, I claim, get the right semantics with regular sets.

For instance, look at the locative modifier *in the shaker* and the noun phrase *salt in the shaker*. I think, with Lønning, that this noun phrase ought to denote a regular set.
But note that the *semantics* of locatives will tell us that if something is in the shaker, its parts are in the shaker. And this means that it is not difficult to make sure compositionally that *salt in the shaker* denotes $\mathbf{ps}_{SALT}(b)$ for some $b \in SALT$: the salt-parts that are in the shaker (i.e. salt-parts that are part of the sum of salt-parts that are in the shaker). Since *SALT* is a regular set, $\mathbf{ps}_{SALT}(b)$ is also a regular set, so the Lønning intuition is satisfied in this case.

But we don't *want* the semantics to work like this in all cases. In the count domain, numericals like *at least three* intersect with the noun interpretation, but their interpretation is not Boolean, and the intersection is not either, and shouldn't be.

Now think about adjectives in the mass domain, like *yellow*. If we assume that *yellow* is a property that mass entities only acquire *in some bulk*, then *yellow* is like *at least three* in the count domain, and we shouldn't expect the noun phrase to denote itself a regular set.

Let us assume that the salt is yellow: $\sqcup(SALT) \in SALT \cap YELLOW$.
Let us assume that *yellow* comes in bulk and that the single salt molecule (NaCl) is not yellow:
$(NaCL) \in SALT - YELLOW$. We have a lot of salt, though, and the color comes in bulk, so intuitively taking that one molecule away leaves us with yellow salt:
$\sqcup(SALT) - (NaCL) \in SALT \cap YELLOW$.

But this means that the denotation of *yellow salt* (on the bulk-interpretation) is not a regular set, since it is not closed under relative complement. And, I think, this is the way it should be: the case is completely parallel to that of *at least three boys* in the count domain, the denotation of which is also not closed under relative complement.

The problem, then, is with Lønning's identification of intersective adjectives with Boolean adjectives, not with the generalization from Boolean denotations to regular sets. We can do the semantics on regular sets just as well as we did on Boolean sets (and, in some cases, we can do better).

5. COUNTING GENERATORS

A regular *noun intension* maps every world onto a regular set. Lexical mass nouns and (plural) count nouns have regular intensions. For noun intension N and world w, we write N_w for the regular set which is the extension of N at w, where $N_w = $ <N_w, **gen**(Nw)>.

We define a function **COUNT** which maps every regular noun intension N and world w onto a relation between the elements of N_w and natural numbers in $\mathcal{N}$: **formulae** Count specifies different ways of counting the elements of N_w, for noun intension N:
-The generators of N_w count as 1. -The count of b $\in N_w$ is the arithmetic sum of the count of its generating parts in N_w. -The count of b $\in N_w$ is also the arithmetic sum of the count of its generating parts *per generating variant*.

To this we add a *correctness criterion for counting*:

Correctness criterion: COUNT is *correct* on a regular noun intension N iff
for every world w: N_w is a *function* from N_w into $\mathcal{N}$

The idea is:

Count nouns have intensions on which **COUNT** is correct.
Mass nouns have intensions on which **COUNT** is incorrect.

The intensional definition takes care of borderline cases of mass denotations with 1 or 0 elements. It is obviously hard to distinguish an empty mass denotation from an empty count denotation, and one may ask why (15a) is felicitous, but (15b) is much less felicitous. I use Dutch examples because in English the facts are muddled, because of the use of *zero* as *no*:

(15) a. Vandaag waren er nul studenten in de klas.
Today were there null students in the class
b. #Vandaag was er nul brood in huis.
Today was there null bread at home

Similarly, we discussed above cases where the noun is mass, but its denotation consists of a single element. (16a) is felicitous, but (16b) is not:

(16) a. There is salt in the microscope, one molecule of salt.
b. #There is one salt in the microscope.

On the account of mass nouns developed here, what we can do with two elements, we cannot do with one: two Na-ions and two Cl-ions make two partitions of salt molecules, similarly two gold atoms and two electrons. But one gold atom and two electrons give two combinations, but not two combinations that partition into variants, since the two elements overlap. In other words, such a set is not closed under relative complement. (This problem was raised by Dafna Rothstein Landman at the first presentation of this talk at the Palmyr conference, and by Manfred Krifka at the Riga conference.)

On the analysis given, **COUNT** is incorrect for *salt* despite null and singleton denotations: to be correct for a regular intension, **COUNT** must be correct for the regular denotation of *salt* in each world, which, of course, it isn't (Krifka (1989) also defines count for noun intensions, albeit for different reasons).

We look at two prototypical examples.

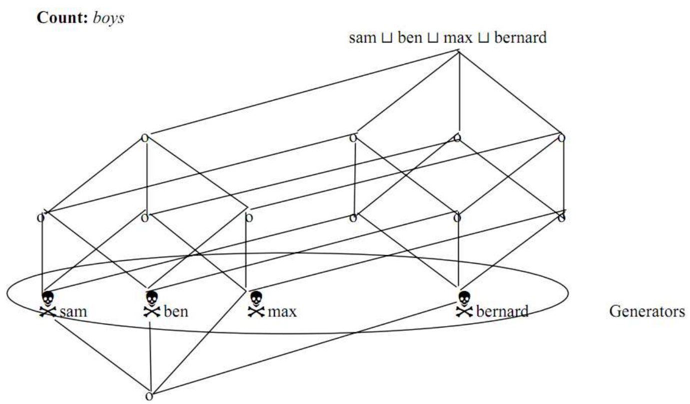

We assume that the count noun *boys* has an intension *BOY* which at each world determines a regular set $BOY_w = $ <∗BOYw,BOYw>, where BOYw is a *disjoint* subse5t of **BOOL.**
In the example the set of generators is: {sam, ben, max, bernard}.

$$\begin{aligned}
&\textbf{COUNT}_{BOY_w}(\text{sam} \sqcup \text{ben} \sqcup \text{max} \sqcup \text{bernard})\\
&= \Sigma\{\textbf{COUNT}_{BOY_w}(\text{a}):\ \text{a} \in \textbf{gen}_{BOY_w}(\text{sam} \sqcup \text{ben} \sqcup \text{max} \sqcup \text{bernard})\}\\
&= \textbf{COUNT}_{BOY_w}(\text{sam}) + \textbf{COUNT}_{BOY_w}(\text{ben}) + \textbf{COUNT}_{BOY_w}(\text{max})\\
&\quad + \textbf{COUNT}_{BOY_w}(\text{bernard})\\
&= 1 + 1 + 1 + 1 = 4
\end{aligned}$$

For every world w, **gen**(BOY_w) forms a single variant for BOY_w. Hence we do not need to check condition 4 of **COUNT** independently.

This means that for every world w, **COUNT** is indeed a *function* on BOY_w, and **COUNT** is correct on *BOY*.

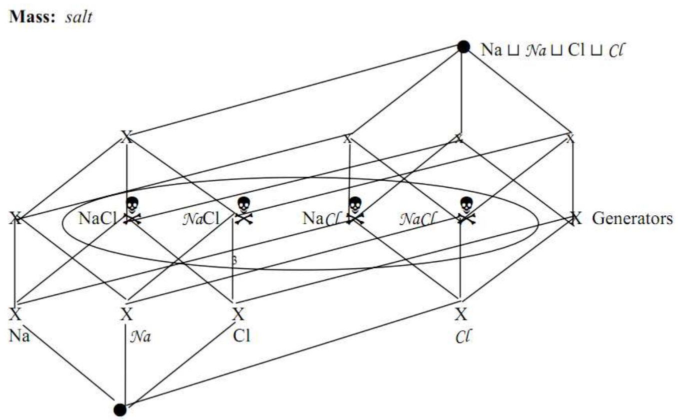

The mass noun *salt* is mapped onto an intension *SALT* which maps each world w onto a regular set $SALT_w = <\text{SALT}_w, \textbf{gen}(\text{SALT}_w)>$. We take the above structure as an example.

By the definition we gave above, SALT_w is the set of elements that contain equal amounts of Na and Cl:

$$\text{SALT}_w = \{0, \text{Na} \sqcup \text{Cl}, \mathcal{Na} \sqcup \mathcal{Cl}, \mathcal{Na} \sqcup \text{Cl}, \text{Na} \sqcup \mathcal{Cl}, \text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}\}$$

We assume in this example that the set of generators equals the set of minimal elements:

$$\textbf{gen}(\text{SALT}_w) = \{\text{Na} \sqcup \text{Cl}, \mathcal{Na} \sqcup \mathcal{Cl}, \mathcal{Na} \sqcup \text{Cl}, \text{Na} \sqcup \mathcal{Cl}\}$$

This set is built from two variants:

$$\begin{aligned}
&\text{V1} = \{\text{Na} \sqcup \text{Cl}, \mathcal{Na} \sqcup \mathcal{Cl}\} &&(\text{Na} \sqcup \text{Cl}) \sqcap (\mathcal{Na} \sqcup \mathcal{Cl}) = 0\\
&\text{V2} = \{\mathcal{Na} \sqcup \text{Cl}, \text{Na} \sqcup \mathcal{Cl}\} &&(\mathcal{Na} \sqcup \text{Cl}) \sqcap (\text{Na} \sqcup \mathcal{Cl}) = 0
\end{aligned}$$

gen(SALT_w) itself is not disjoint:

$$(\text{Na} \sqcup \text{Cl}) \sqcap (\text{Na} \sqcup \mathcal{Cl}) \neq 0 \text{ and } (\mathcal{Na} \sqcup \mathcal{Cl}) \sqcap (\text{Na} \sqcup \mathcal{Cl}) \neq 0$$

Hence the salt is built from building blocks that overlap.
Now we count:

$$\begin{aligned}
&\textbf{COUNT}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}) =\\
&\quad = \Sigma\{\textbf{COUNT}_{SALT_w}(\text{a}):\ \text{a} \in \textbf{gen}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl})\}\\
&\quad = \textbf{COUNT}_{SALT_w}(\text{Na}) + \textbf{COUNT}_{SALT_w}(\mathcal{Na}) + \textbf{COUNT}_{SALT_w}(\text{Cl})\\
&\qquad + \textbf{COUNT}_{SALT_w}(\mathcal{Cl})\\
&\quad = 1 + 1 + 1 + 1 = 4
\end{aligned}$$

$$\begin{aligned}
&\textbf{COUNT}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}) =\\
&\Sigma\{\textbf{COUNT}_{SALT_w}(\text{y}):\ \text{y} \in \text{V}_1\} = 2
\end{aligned}$$

$$\begin{aligned}
&\textbf{COUNT}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}) =\\
&\Sigma\{\textbf{COUNT}_{SALT_w}(\text{y}):\ \text{y} \in \text{V}_2\} = 2
\end{aligned}$$

Hence, all in all:
$\textbf{COUNT}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}) = 4$ and $\textbf{COUNT}_{SALT_w}(\text{Na} \sqcup \mathcal{Na} \sqcup \text{Cl} \sqcup \mathcal{Cl}) = 2$

So, **COUNT** is not a function on *SALT*w, and **COUNT** is incorrect on *SALT*.

6. COUNT AND MASS – NEAT AND MESS

We now take up the motivating idea concerning overlap and define count nouns as nouns whose intension at every world specifies a regular set built from a set of non-overlapping generators, while mass nouns are nouns whose intension at every world specifies a regular set built from a set of non-overlapping generators (if the denotations are big enough to allow this):

Let *X* be a function from worlds to regular sets.

X is [+C], **count**, iff for every w: **gen**(X_w) is disjoint
i.e. the generators of X_w do not overlap
X is [–C], **mass**, iff for every w: if $|X_w|>1$ then **gen**(X_w) is not disjoint,
i.e. the generators of X_w overlap.

While defined for intensions, we will freely use these features for the structures X_w themselves, and call these structures count and mass.

This definition in principle allows nouns that are neither mass nor count. I don't assume that lexical nouns can be specified that way. While I could have defined [–C] as the complement of [+C], the definition given presents mass and count more strongly as different semantic perspectives: always overlapping generators *versus* always disjoint generators.

Interestingly enough, the theory of regular sets allows a *second kind of mass structure*, which is mass, but in several ways closer to count. For this, we introduce the following opposition:

X is [+N], **neat**, iff for every w: **min**(X_w) is disjoint
i.e. the minimal elements of X_w do not overlap
X is [–N], **mess**, iff for every w: if $|X_w|>1$ then **min**(X_w) is not disjoint,
i.e. the minimal elements of X_w overlap.

On this definition, neat nouns are nouns whose intension at every world specifies a regular set whose set of minimal elements is non-overlapping, while mess nouns are nouns whose intension at every world specifies a regular set whose set of minimal elements is overlapping (again, if the denotations are big enough to allow this).

By definition, **count** entails **neat**: [–N] ⇒ [–C]
Equivalently, **mess** entails **mass**: [+C] ⇒ [+N]

The mass structure for *salt* given in the previous section is a structure that is **mess mass** [–C, –N]. Its set of generators overlap ([–C]), and since the set of generators is the set of minimal elements, its set of minimal elements overlap ([–N]).

But the theory allows structures that are **neat mass**: [–C, +N]. These are structures in which the set of generators overlaps, but the set of minimal elements does not.

I propose that these structures are precisely suited for mass nouns like *furniture* and *kitchenware*:

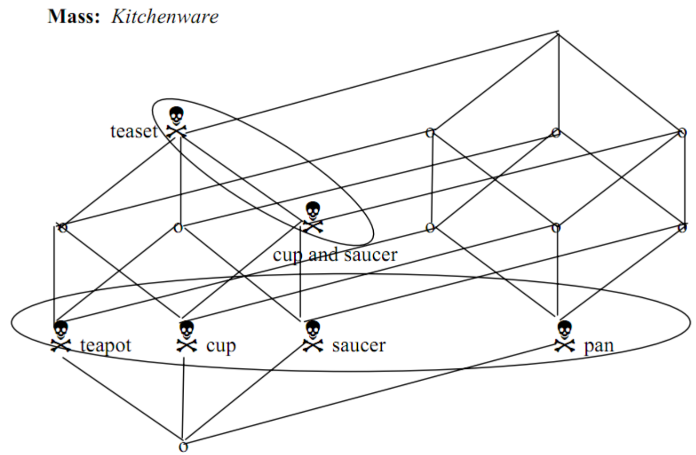

In this structure, the set of generators includes more than just the minimal elements. The building blocks are what we intuitively want to count as one. Thus, in this structure singularities and pluralities are counted as one simultaneously, without making sure that they do not overlap.

The difference with **count** is that for count nouns a plurality of *boys* does not itself count as *one boy*. But a plurality of *kitchenware*, like the *cup and saucer*, can count itself as *kitchenware*, and can also count as one. For instance, it counts as one on an inventory listing where everything that is sold as one item has its own price.

Rothstein (2010) discusses count nouns like *line, highway, mirror*: objects in the denotation of these nouns typically have objects as parts that themselves can be in the denotation of these nouns: a line divides into lines, a highway into highways, a mirror breaks into mirrors.

But before the mirror breaks, we do not, in a normal context, count the mirror and its parts that would count as mirrors when broken as *more than one*: only the maximal mirror counts. Thus the mirrors that we *do* count don't overlap, or we *make* them not overlap by packaging.

Neat mass denotations are different: the teapot, the cup, the saucer, the cup and saucer all count as kitchenware and can all count as one simultaneously in the same context.

Neat mass nouns differ from *mess mass* nouns like *salt* and *meat*, in that the minimal building blocks of neat nouns are non-overlapping: the minimal building blocks of *meubilair*, furniture, are the *meubels*, the pieces of furniture. The generating set of *furniture* overlaps, but *the overlap is only vertical*: a sum and its parts count as one simultaneously.

In other words: the denotations of *neat nouns* are sets in which the distinction between *singular individuals* and *plural individuals* is not properly articulated.

In context, the denotation of *furniture* may be equated with its set of generators. Then we would get the following denotations:

(17) *meubel* $\longrightarrow$ $<$MEUBEL, MEUBEL$>$ with MEUBEL a disjoint set
meubels $\longrightarrow$ $<*$MEUBEL, MEUBEL$>$
meubilair $\longrightarrow$ $<*$MEUBEL, $*$MEUBEL$>$

We see that this is in fact very close to what I proposed as a reasonable version of the theory of Chierchia 1998, except, of course, that I add here an interpretation to these pairs, that fits them naturally into the theory of mass nouns I am developing here. Hence, the reason why you cannot count neat nouns is not the reason that Chierchia gives. The reason is that counting goes wrong.

In the above example, **gen**($KITCHENWARE_w$) =
{ *the teapot, the cup, the saucer, the pan, the cup and saucer, the teaset* }

We calculate the count for *the teaset*:

- $\mathbf{COUNT}_{KITCHENWARE_w}$(*the teaset*) = 1
because *the teaset* $\in$ **gen**($KITCHENWARE_w$)
Generators count as 1.
- $\mathbf{COUNT}_{KITCHENWARE_w}$(*the teaset*) = 5
because *the teaset* = *the teaset* $\sqcup$ *the teapot* $\sqcup$ *the cup* $\sqcup$ *the saucer*, and *the teaset, the teapot, the cup, the saucer* $\in$ **gen**($KITCHENWARE_w$)
The count of the teaset is the arithmetic sum of the count of its generator parts.

- $\mathbf{COUNT}_{KITCHENWARE_w}$(*the teaset*) = 2
because *the teaset* = *the teapot* $\sqcup$ *the cup and saucer*, and
the teapot, the cup and saucer $\in$ **gen**($KITCHENWARE_w$)
Here we count relative to the variant: {*the teapot, the cup and saucer*}.
- $\mathbf{COUNT}_{KITCHENWARE_w}$(*the teaset*) = 3
because *the teapot* $\sqcup$ *the cup* $\sqcup$ *the saucer*, and
the teapot, the cup, the saucer $\in$ **gen**($KITCHENWARE_w$)
Here we count relative to the variant: {*the teapot, the cup, the saucer*}.

Clearly, then, **COUNT** is incorrect, and *kitchenware* is mass. It is neat mass, because the minimal elements are neatly disjoint.

7. INDIVIDUATED SETS AND THE TWO FEATURE SYSTEM

Rothstein (2010) assumes that the mass nouns *furniture* and *kitchenware* are like the count nouns *boys* and *peas* in that their sets of minimal elements are *individuated* or *naturally atomic*. The following is an attempt at (partially) formalizing this notion.

> Let X be a regular noun intension and D be a set of naturalistic properties, like properties of Form, *square, round*,..; properties of Size, *big, small*,..; properties of Weight, *heavy, light*,..; properties of Color, *red, green*,.. etc...
>
> A subset of D is a *dimension set* for X, D_X, if D_X consists of properties of which it is natural (in every salient world w) for the generators of X_w, the elements of **gen**(X_w), to have them.

By this we mean that the generators of X_w are the kind of things that we distinguish in terms of whether they are *big* or *small*, *red* or *green*, etc...

> X is *individuated by dimension set* D_X if each property in D_X is a bipartition on **gen**(X_w), and the properties in D_X jointly determine the partition into singletons: $\{\{x\}: x \in \mathbf{gen}(X_w)\}$ (for every salient world w, and non-trivial regular set X_w).

The idea is that D_X consists of natural properties, and enough of them, to tell the generators apart. Individuation is not counting: you can individuate the generators of a noun denotation in w with natural properties, partition them into finegrained natural units down to the level of singletons, without ending up with non-overlapping objects.

But counting is itself individuation: we assume that generators that are made non-overlapping in context (i.e. count) are *ipse facto* individuated.

We give this the following form:

The *extensional dimension set* E_{Xw} is:

$E_{Xw} = \{\ \lambda x \in \mathbf{gen}(X_w)\colon\ \forall y \in \mathbf{gen}(X_w)-\{x\}\colon\ x \sqcap y = 0\ \}$

The set consisting of the property that a generator has if it is disjoint from all other generators.

-Noun intension X is [+I], *individuated*, iff there is a *salient* dimension D_X which for every world w individuates X_w (if X_w is non-trivial).

-We assume that E_X, the extensional dimension set, is always salient.

We let X be [–I], *non-individuated*, if X is not individuated.

On this formalization, X_W is individuated by E_X iff X is count, and hence **count** entails **individuated**: [+C] ⇒ [+I]
We have now three features: [±C], [±N] and [±I].
I will adopt a *Two Feature System* in which the structural notion **neat** (no overlapping minimal elements) and the more intensional notion **individuated** are taken to coincide:

Strong Mess Mass assumption: [+N] ⇔ [+I]

This is a constraint on noun intensions: we restrict the noun intensions available for the interpretation of natural language to those that satisfy the equivalence [+N] ⇔ [+I]. This makes no difference for count nouns, which are by definition neat and extensionally individuated, but the constraint says that when a mass noun has a neat denotation, its generators are interpreted as individuated by a salient (intensional) dimension set, and it says that when a mass noun has a mess interpretation, there is no natural salient dimension set individuating its generators.

The Two Feature System gives the following set of features, which we assume to be lexically specified on nouns in English:

[+C, +N]	= [+C]	count:	*boys, peas*
[–C,–N]	= [–C, –N]	mess mass:	*meat, cheese*
[–C,+N]	= [–C, +N]	neat mass:	*furniture, kitchenware*

The theory makes the following natural distinctions:

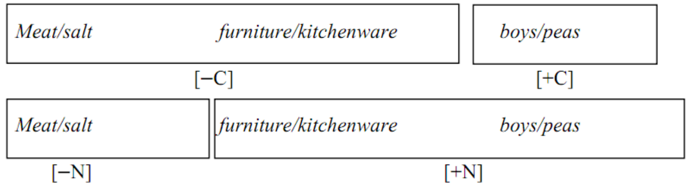

And the hypothesis is that these contrasts are semantically robust, meaning that natural languages will cluster properties around these two boundaries, both within one language and cross-linguistically.

For the feature [±C] this is, of course, well attested in the literature. See for example the following table:

	[–C]		[+C]	
1. Plural:	salt	#salts	boy	√boys
	furniture	#furnitures		
2. Numericals	#one salt	#two salt	√one boy	√two boys
	#one furniture	#two furniture		
3. Quantifiers:	#every meat		√every boy	
	#many meat	#many furniture		√many boys
	√much meat	√much furniture	#much boy	#much boys

8. THE ROBUSTNESS OF THE FEATURE [±N]

In this section I discuss four phenomena which show that the feature [±] is semantically robust.

8.1. *The classifier stuks in Dutch.*

Dutch has a classifier *stuks* with a meaning similar to the English *head* (as in *head of cattle*) but with a much wider use. Doetjes (1997) observes that *stuks* applies to count nouns and to neat mass nouns, but not to mess mass nouns, and that the noun phrase [*stuks* NOUN] is count:

(18) COUNT

a. Hoeveel *hemden* neem je mee op vakantie? Drie
How-many *shirts* take you with on vacation Three
stuks.
items

b. Hoeveel *croquetten* heb je gegeten? Zes *stuks*.
How-many *meat rolls* have you eaten? Six *items*

(19) NEAT MASS

a. Hoeveel *meubilair* heb je besteld? Drie stuks
How-much *furniture* have you ordered? Three *items*

b. Hoeveel *keukenwaar* heb je aangekruist in the
How-much *kitchenware* have you checked in the
catalogus? Acht *stuks*
catalogue? Eight *items*

c. Hoeveel *vee* heb je gekocht? Drie *stuks*, twee
How-much *cattle* have you bought Three *items* two
schapen en een koe.
sheep and a cow

(20) MESS MASS

a. Hoeveel *kaas* heb je gekocht? #Drie stuks.
How-much *cheese* have you bought? #Three items

b. Hoeveel *vlees* heb je gegeten? #Drie *stuks*.
How-much *meat* have you eaten? #Three *items*

The classifiers can also occur prenominally, but this is most natural in list contexts:

(21) *Checking a sorted order list from an online Department store:*
U heeft drie stuks meubilair, zes stuks keukenwaar, twaalf
You have three items furniture, six items kitchenware, twelve
stuks fijne vleeswaren, en zes stuks sport artikelen aangekruist.
items cold cuts and six items sports-products checked.

The singular *stuk* appear in the expression *per stuk/per item*:

(22) a. De *bloemen* worden niet *per stuk* verkocht, maar
The *flowers* are not *per item* sold, but
per tien stuks.
per ten items

b. De *keukenwaar* wordt niet *per stuk* verkocht, maar
The *kitchenware* is not *per item* sold, but
per tien stuks
per ten items

c. #Het *vlees* wordt niet *per stuk* verkocht, maar
The *meat* is not *per item* sold, but
per tien stuks.
per ten items

A caveat: as one can easily find out by searching the internet, workers in the catering branch do not, in their internet exchanges, distinguish very carefully between the classifier *stuks (items)* and the plural noun *stukken (pieces)*. This means that they produce data which contradicts the data in (20):

(23) a. Een bitter garnituur bestaat uit zes stuks worst, zes stuks kaas en zes stuks bitterballen.
A meat roll dish consists of six items sausage, six items cheese and six items meat rolls

This may be a linguistic innovation or sloppiness. Both for Doetjes, for me and my informants, (23a) is ungrammatical; its content should be expressed as (23b):

(23) b. Een bitter garnituur bestaat uit zes *stukjes* worst, zes *stukjes* kaas en zes stuks bitterballen.

A meat roll dish consists of six pieces sausage, six pieces cheese and six items meat rolls

I will ignore the internet innovation here.

The classifier *stuks* takes neat nouns denotations as input and turns them into count noun denotations. How does it do this?

That depends on the semantics of the input noun. All neat nouns are individuated, but some are more individuated than others. *Vee* is a mass noun in Dutch, while *cattle* is a plural noun in English. There is no doubt, however, either in Dutch or in English, *which* elements count as the most elementary building blocks of *vee*: the heads of cattle. In this, *vee* is like prototypical count nouns.

[*Vee* means *domesticated farm animals, live-stock*, typically cows, sheep, goats, but also chickens (*pluimvee/feathered live-stock*). However, out of the blue, *vee* means *cattle* (and that's the only thing my pocket dictionary Dutch-English gives). Below, I will, for ease, translate *vee* as *cattle*, except where I explicitly mean *live-stock*.]

Rothstein (2010) uses the term *inherently atomic* for prototypical count nouns, to distinguish them from count nouns like *fence*. For inherently atomic neat mass nouns, the interpretation of *stuks* is the following:

Let X be an inherently atomic neat noun intension. For every world w:
$\mathbf{stuks}(X_w) = <*(\mathbf{min}(X\mathrm{w})), \mathbf{min}(X\mathrm{w}))>$

Stuks vee has the same denotation as the count noun phrase *domesticated farm animals*.

Neat nouns like *kitchenware* are less inherently atomic, in the sense that we saw above: in context, it is not automatically obvious whether something is meant to count as three or as one. In this case, we can assume that different choices are possible:

For regular set X, let V_X be the set of all variants in X.

Let X be a neat noun intension which is not inherently atomic.
In context k, let $\mathbf{stuks}_k$ be a function which maps X and world w onto a set

$\mathbf{stuks}_{k,X,w} \in \mathrm{V}_{Xw}$, a variant for X_w.
We define, for context k and world w:
$\mathbf{stuks}\mathrm{k}(X_w) = <*\mathbf{stuks}_{k,X,w} , \mathbf{stuks}_{k,X,w}>$

In context k, we choose a variant of the generators of *kitchenware*, say, *the pan, the teapot*, and *the cup and saucer*, and let *stuks keukenwaar* denote the closure under sum of that set.

For inventory list counting contexts, we may want to count *all* the generators, not just a variant. In that case, we have to make the generators disjoint:
Let $\uparrow (\mathrm{X}) = \{\uparrow(\mathrm{x}): \mathrm{x} \in \mathrm{X}\}$, with $\uparrow$ the packaging operation defined in section 9 below.
Let X be a neat noun intension which is not inherently atomic. Let i be an inventory context.

$\mathbf{stuks}_i(X_w) = <* \uparrow(\mathbf{gen}(X)), \uparrow(\mathbf{gen}(X))>$

This will package the generators that are pluralities as count atoms.

In all these cases, the resulting noun phrase is a count noun phrase, and COUNT is *correct* on its denotation.

8.2. Counting in Chinese.

The Dutch classifier *stuks* is very similar to Chinese individual classifiers, like the general individual classifier *ge*.

We follow Chierchia (1998) and Li (1983) in assuming that semantically all lexical nouns in Chinese are [–C]. We assume that numericals in Chinese, like numericals in English, require noun denotations as input that are [+C]. It follows from this, that numericals cannot combine with lexical nouns in Chinese:

(24) *#Liǎng níu*   *#Liǎng ròu*
two cow   two meat

We assume that the Chinese nouns that correspond to English prototypical count nouns are neat nouns, [+N], in Chinese. And we assume that *ge* is much like *stuks* in Dutch, in that it maps [+N] nouns onto noun phrases that are count.

The difference with Dutch, then, is that the class of [–C,+N] nouns in Chinese is much larger, and, by necessity, the classifier construction is fully productive. For most of these nouns, the first interpretation strategy given for *stuks* above—mapping the neat noun interpretation on the closure of its *set of minimal elements*—will be appropriate for Chinese *ge* as well.

Since *ge* requires a neat noun as input, we find the contrast in (25):

(25) a. *ròu* [–N] *níu* [+N]
meat cow

b. *#Liăng ge ròu* ✓*Liăng ge níu*
two CL [meat$_{[-N]}$] two [CL cow$_{[+N]}$]

8.3. Distributive adjectives.

Rothstein (2009b) and Rothstein (2010) discuss distributive adjectives, adjectives that have distributive interpretations and resist collective interpretations (I see no reason to call them 'stubbornly distributive' as Schwarzschild does: in my usage 'not distributive' means 'not necessarily distributive' (since any predicate can be made to distribute)).

Distributive: *Small, big, large, round, square,. . .*
Not distributive: *noisy, successful,. . .*

(26) a. The boys are *noisy/successful*
b. The boys are *small/big*

(26a) can mean that the individual boys are noisy/successful, but also that the boys are noisy/successful as a group. (26b) only means that the individual boys are small/big, not that the boys are small/big as a group.

Schwarzschild and Rothstein (independently) observe that distributive adjectives modify neat mass nouns like *furniture* and *kitchenware* in the same way as they modify count nouns, while mess mass nouns pattern differently:

(27) a. The *furniture* is big.
b. The *big furniture* is exhibited on the third floor.

(28) a. The *meat* is big.
b. The *big meat* is in the other fridge.

For *furniture* in (27) we find exactly what we found for count nouns: (27a) expresses that the furniture generators, the pieces of furniture, are big. The big furniture in (27b) consists of the pieces of furniture that are individually big, like the sofa's and the pianola's.

This reading is absent for mess mass nouns like *meat* in (28): (28a) does not mean that the meat-generators are big; (28b) does not mean that all big meat-generators are in the other fridge. (28a) means that the meat comes in the form of a big hunk, while (28b) means that the meat that is packaged in big hunks is in the other fridge. *Big* clearly does not distribute to the meat-generators: in the contexts of our meat-cutting machine, the generators are all small. The *small-big* distinction is not a salient individuating dimension on meat generators, and in fact, applying distributive adjectives to mess mass nouns is, out of context, often not fully felicitous. The distributive adjectives are typically the ones that are naturally used as part of *individuating dimension sets*.

The question is: how do distributive adjectives distribute in the neat mass domain? In the count domain, the distributive operator is standardly regarded as a null version of the adverbial *each*. Now that we discover that there is distributivity in the (neat) mass domain, questions are raised about the connection between distributivity and *each*, because, of course, *each* cannot apply to mass nouns (*#each furniture*).

I propose that the connection with *each* is not given up. I assume that the distributive operator on neat noun denotations semantically shifts the neat noun denotation X to count interpretation $\mathbf{stuks}_k(X)$, and distributes on the latter as usual:

Let X be a *neat noun* intension, for every world w and context k:

$$\mathbf{big} = \lambda X_w.\ <*(\mathbf{gen}(\mathbf{stuks}_k(X_w)) \cap \mathbf{big}),\ \mathbf{gen}(X_w) \cap *(\mathbf{gen}(\mathbf{stuks}_k(X_w)) \cap \mathbf{big})>$$

Note that the shift takes place *as part of* the meaning of *big*: we are not shifting the mass noun to a count noun. And, in fact, we assume that the result of applying *big* to X is a neat mass noun: the new set of generators is $\mathbf{gen}(X_w) \cap *(\mathbf{gen}(\mathbf{stuks}_k(X_w)) \cap \mathbf{big})$, the set of X generators that are sums of *big stuks N*, not simply the set of minimal elements $\mathbf{gen}(\mathbf{stuks}_k(X_w)) \cap \mathbf{big}$. (In fact, we can *make* it a requirement on felicitous application of *big* to *X*, that the result stay neat mass.) Hence for neat mass noun N, *big N* is a neat mass noun phrase, not a count noun phrase. But *big* distributes to the generators that count as 'stuks'.

For inherently atomic neat mass nouns like *vee* in Dutch, this means that in *groot vee* (*big cattle*), *groot* distributes to *stuks vee*, the minimal elements, the individual heads of cattle. This is as it should be.

For less inherently atomic neat nouns like *kitchenware*, we expect that *big* need not distribute to the minimal elements, like it does in *vee: big* distributes to *stuks of kitchenware*, but what counts as *stuks of kitchenware* is context dependent. Thus, *the teapot, the pan, and the cup and saucer* may all count as big, even though *the cup* itself and *the saucer* itself don't. Hence the distribution of *big* is predicted to be contextual in exactly the way that the interpretation of *stuks* (in Dutch) is contextual.

For the inventory reading, a bit more work will have to be done concerning the set of generators (because it is not necessarily a subset of $\mathbf{gen}(X_w)$). It is reasonable to assume it to be at least $(\mathbf{gen}(X) \cup \uparrow(\mathbf{gen}(X)) \cup \mathbf{big}$. Here too we can make it a requirement that the interpretation of the complex stays neat mass.

In sum, *big* is distributive on neat nouns through classifier $\mathbf{stuks}_k$. This predicts that distributive adjectives treat inherently atomic neat nouns the same as count nouns (distribution to minimal elements), but other neat nouns show the distribution expected from their possible $\mathbf{stuks}_k$ interpretations.

8.4. Neat comparison.

Barner & Snedeker (2005) present experimental data to show that for children and adults neat nouns pattern with count nouns, when it comes to size comparisons: both compare in terms of cardinality. We expand upon this result in this section.

Barner and Snedeker's results have direct linguistic consequences for the semantics of *most*, which involved comparison. We give the examples in Dutch, because this will allow us to use the neat mass noun *vee (cattle)*.

We look at available readings for *de meeste* (*most*).

[Note Dutch *de meeste* is ambiguous between *most* and the superlative *the most*. For the examples below, a superlative reading is hard to get out of the blue, so it can be reasonably ignored here.] (29a) means (29b):

(29) MESS MASS

a. Het meeste *vlees* wordt gegeten op zon- en feestdagen
Most *meat* is eaten on (sun and holi)-days

b. *Meer* vlees wordt gegeten op zon- en feestdagen dan op andere dagen
More meat is eaten on (sun and holi)-days than on other days

Vlees/meat is a mess mass noun and the measure involved is a *mass measure*:

more = more in volume/more in weight.... etc.

The reading which is unavailable is a counting reading, i.e.

more = more in number of generators, more in number of minimal generators.

The reason is clear: if you were to count generators or minimal generators in the denotation of mess mass nouns you would be counting wrong.

(30a) means (30b):

(30) COUNT

a. De meeste *koeien* zijn buiten in de zomer
Most *cows* are outside in the summer

b. *Meer* koeien zijn buiten in de zomer dan binnen
More cows are outside in summer than inside

In this case, the only reading available is the counting reading:

more = more in number of generators = more in number of minimal elements.

(31) a. De meeste *stuks vee* zijn buiten in de zomer
Most heads of cattle are outside in summer
b. *Meer* stuks vee zijn buiten in de zomer dan
More heads of cattle are outside in summer than
binnen
inside

We look at neat nouns: (32a) means (32b):

(32) NEAT MASS, inherently atomic.
a. Het meeste *vee* is buiten in de zomer
Most cattle is outside in the summer
b. *Meer* vee is buiten in de zomer dan binnen
More cattle is outside in the summer than inside

Barner and Snedeker's results show that the most prominent interpretation of the comparative in (32b) is similar to that of the comparative in (31b). Hence, the most prominent reading available for (32) is the *counting* reading:

more = more in **stuks** = number of minimal elements.

I say 'the most prominent reading.' Let me be more precise: I think that *the only counting reading available* for (32) is the reading on which it is equivalent to (31), and hence counting is in terms of minimal elements.

But this is for 'inherently atomic' neat mass nouns. For less inherently atomic neat mass nouns, the most prominent comparison is also in terms of counting generators, but in that case, this need not be necessarily minimal generators.

(33) NEAT MASS, not inherently atomic.
a. De meeste keukenwaar kost meer dan 5 euros
Most kitchenware costs over 5 euros
b. *Meer* keukenwaar kost meer dan 5 euros dan 5
More kitchenware costs over 5 euros than 5 euros
euros of minder
or less

Situation 1:
In this shop, *the teapot* is 6 euros, *the cup and saucer* is 4.50, and *the pan* is 12 euros. You cannot buy *the cup* separately, nor *the saucer*, and the *teaset* is just *the teapot and the cup and saucer*, no price differences there. Two items cost more than 5 euros, one item less, hence (33) is true.

more = *more* in $\textbf{stuks}_k$: not more in minimal generators, but in the choice of generators determined by $\textbf{stuks}_k$

Situation 2:
In the neighbouring shop, *the cup* is 3 euros, *the saucer* is 3 euros, you pay 5.50 for *the cup and saucer*, the *teapot* is 6 euros, *the teaset* is 11 euros.
In this shop, three items cost more than 5 euros, and 2 items less, (33) is true.

more = *more* in $\textbf{stuks}_i$, where i is an inventory context: *more* is more in terms of the whole set of generators, counting each generator independently as one.

We see that, as in the previous case of distribution, the comparison in the neat noun is in terms of $\textbf{stuks}_k$. This gives the following semantics:

Let X be a neat noun intension [+N]. For every world w and context k:

$$\text{MOST}_{\textbf{stuks}}(X_w, \textbf{P}) = 1 \text{ iff } |\textbf{gen}(\textbf{stuks}_k(X_w)) \cap \textbf{P}| > |\textbf{gen}(\textbf{stuks}_k(X_w)) - \textbf{P}|$$

For count nouns, $\textbf{gen}(\textbf{stuks}(COW\text{w})) = \text{COW}_w$, hence:

$$\text{MOST}_{stuks}(COW_w, \text{OUTSIDE}_w) = 1 \text{ iff } |\text{COW}_w \cap \text{OUTSIDE}_w| > |\text{COW}_w - \text{OUTSIDE}_w|$$

For inherently atomic neat noun *vee*, let us assume than in the context of our farm, the *vee/live-stock* consists of cows and chickens.

$\textbf{gen}(\textbf{stuks}(VEE)) = \text{COW} \cup \text{CHICKEN}$. Let us set $\text{FA} = \text{COW} \cup \text{CHICKEN}$. Then:

$\text{MOST}_{\textbf{stuks}}(\text{VEE}_w, \text{OUTSIDE}_w) = 1$ iff $|\text{FA OUTSIDE}| >$
$|\text{FA} - \text{OUTSIDE}|$

I said above that the counting reading in terms of **stuks** is the most prominent reading of neat mass nouns. I also said that for inherently atomic neat nouns it is *the only counting reading*. But it is not the only reading. This is a major reason why I am unhappy with Chierchia's (2010) classification of neat mass nouns as 'fake' mass nouns. Because neat mass nouns are not fake mass nouns, they are real mass nouns, and the evidence is that most can compare neat mass nouns in terms of the measures that are appropriate for mess mass nouns.

Suppose that there are cows and chickens, and the cows are kept outside, but the chickens are kept inside. The chickens outnumber the cows, but in terms of biomass and volume, there is less biomass and less volume of chicken.

(34) a. Wat biomassa betreft, wordt het meeste vee
As biomass concerns is most live-stock outside
buiten gehouden
kept

b. In termen van volume, wordt het meeste vee
In terms of volume is most live-stock outside
buiten gehouden
kept

(35) a. #Wat biomassa betreft, worden de meeste stuks
As biomass concerns are most items livestock
vee buiten gehouden
outside kept

b. #In termen van volume, worden de meeste stuks
In terms of volume are most items live-stock
vee buiten gehouden
outside kept

The examples in (35) are infelicitous, or rather, it isn't clear what the biomass/volume adjunct has to do with the rest of the sentence, because, clearly, *de meeste* (*most*) in (35) compares sets of *stuks of vee* in

terms of cardinality.

But this is not true in the examples in (34). In terms of cardinality, *most farm animals* are *not* kept outside, because there are more chickens than cows, and the chickens, unfortunately, are inside. Still, (34) is true, because the comparison *can* be in terms of biomass or volume.

Compare also (36) in English:

(36) a. In terms of volume, most live-stock is cattle.
b. #In terms of volume, most farm animals are cattle.

(36b) is funny, and in as much as it is felicitous it is false in the above scenario. (36a), on the other hand, is true.

We see, then, that the facts are in line with what Barner and Snedeker's experiments show, but they are more subtle. Counting comparison for neat nouns, like distributivity, is in terms of $\textbf{stuks}_k$, hence counting comparison for neat nouns is only strictly identified with counting minimal generators for inherently atomic neat mass nouns. The counting comparison is more flexible and context dependent for less inherently atomic neat mass nouns.

Moreover, neat mass nouns are true mass nouns in that mass measure interpretations are available for *most* NOUN, if the noun is a neat mass noun; this is a real difference with count nouns: mass measure interpretations are completely unavailable for *most* NOUN if the noun is a count noun.

9. FUSION AND FISSION

I assume that *packaging* as an operation from *mass entities* to *count entities* is the same operation as *group formation* (as assumed in Landman 1991): a mass or count sum is treated as a count atom, more than the sum of its parts. But bringing in packages and groups simultaneously is more complex than I am willing to deal with here. So I deal only with packaging, and assume the following picture:

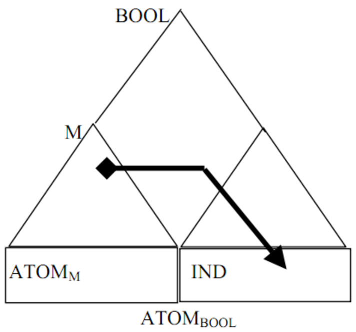

-The generator sets of *mass* predicates are subsets of the domain M.
-C = M ∪ IND.
-The generator sets of count predicates are disjoint subsets of C.

In honor of the fact that I assume one operation for packaging and group formation, I will give it a new name, and since the operation *fuses* a plurality into an atom, I will call it *fusion*:

Fusion: $\uparrow$: M–ATOM$_M$ $\rightarrow$ IND is a one-one function into IND
Fusion is an injection from M-sums into atomic packages.
$\uparrow$: ATOM$_M$ $\rightarrow$ ATOM$_M$ = {<a,a>: a ∈ ATOM$_M$}
Fusion is identity on M-atoms

Not every element of IND needs to be in **ran**($\uparrow$), the range of $\uparrow$. If *fido* is in IND, *fido* is not only more than the sum of his mass parts, but also more than the fusion of his mass parts. But there is an equivalence relation relating *fido* uniquely to the fusion of his mass parts:

$\approx$ is an equivalence relation on IND such that:
for every a ∈ IND there is exactly one b ∈ IND such that b $\approx$ a and b ∈ **ran**($\uparrow$).
For a ∈ IND, we let $a_\approx$ be the unique element of $[a]_\approx$ such that $a_\approx$ ∈ **ran**($\uparrow$).

This equivalence relation is used to relate *fido* to the sum of his M-parts.

Besides fusion, we have an operation that *splits* an IND atom into a plurality of M-elements, maps a set onto the sum of its splits, and a regular mass or count set onto the sum of the splits of its generators:

Split: $\downarrow_o$: C $\rightarrow$ M defined by:
$\downarrow_o$(b) = b if b ∈ M
$\downarrow_o$(b) = $\uparrow^{-1}$($b_\approx$) if b ∈ IND
$\downarrow_o$(X) = ⊔({$\downarrow_o$(x): x ∈ X}) if X ⊆ C
$\downarrow_o$(*X*) = $\downarrow_o$(**gen**(X)) if *X* = <X,**gen**(X)> is a regular mass or count set

In terms of split, we define an operation of *fission*:

Fission: $\downarrow$
$\downarrow$(b) = ($\downarrow_o$(b)] if b ∈ C
$\downarrow$(X) = ($\downarrow_o$(X)] if X ⊆ C
$\downarrow$(*X*) <$\downarrow_o$(X)],$\downarrow_o$(X)]−{0}> if *X* = <X,**gen**(X)> is a regular mass or count set

If *fido* is in IND, then the split of *fido* is the sum of *fido*'s M-parts, and the fission of *fido* is the set of all Boolean parts of the split of *fido*: ($\downarrow_o$(*FIDO*)]

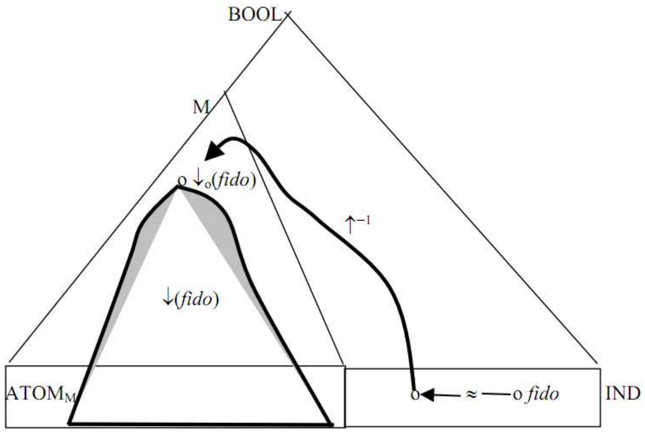

Suppose the individual dogs are d_1, d_2 and $d_2 \in$ IND. Then:

$DOG = <*\{d_1, d_2, d_3\}, \{d_1, d_2, d_3\}>$.

So **gen**(*DOG*) $= \{d_1, d_2, d_3\}$ and $\downarrow_o(\mathbf{gen}(DOG)) = \uparrow^{-1}(d_{1\approx}) \sqcup \uparrow^{-1}(d_{2\approx}) \sqcup \uparrow^{-1}(d_{3\approx})$.
For each dog d_i, go to the package of d_i's sum of M-parts, and go back to the corresponding sum of M-parts. Sum these parts together: that is the split of **gen**(DOG).

The fission of DOG, ↓(DOG), is the set of all Boolean parts of that split, and the fission of *DOG*, ↓(*DOG*), is the regular set consisting of ↓(DOG) and ↓(DOG)–{0} as set of generators.

10. FISSION READINGS

(37) a. There is *human* in this dish.
b. There is *cat* in this soup.
c. There was *dog* all over the wall.

The cases in (37) are examples of *grinding*, which we have rebaptized *fission*: in all these cases a [+C] noun is given a [–C] interpretation.

Rothstein (2009a) provides cross-linguistic evidence that fission of nouns is only possible as a *last resort mechanism* to resolve grammatical

mismatch.

Rothstein's account for the English cases in (37) is as follows:
-The singular copula in (37) is followed by a bare noun and requires number agreement. There are no bare singular nouns in English, only bare mass nouns. The bare nouns *human, cat* and *dog* are lexically count. This is a grammatical conflict.
This conflict is resolved by *fission*: *NOUN*[+C] ⇒↓(*NOUN*) [–C]

Cheng et al. (2008) point out that in Chinese, (38), which corresponds to (37c), *does not* have a fission reading, but only a plural, wallpaper reading; (38) can express that the wall is covered with doggie-wallpaper.

(38) qiáng-shang dou shì gǒu
wall- top all COP dog
There is dog all over the wall.

Rothstein's account for Chinese is as follows:
Chinese nouns are not specified for number, there is no number agreement between the copula and the noun, so the bare noun is grammatical in (38) and allows a plural interpretation. On the assumption that fission is a last-resort mechanism, it follows that (38) does not have a fission interpretation.

Rothstein (2009a) argues that in Hebrew, as in Chinese, fission readings are not possible, but they can be triggered by a mismatch in grammatical *gender* between the copula and the post-copular bare noun.

Cheng et al. (2008) point out that natural foodstuff nouns in Chinese do have mass interpretations:

(39) a. shālā lǐ yǒu zhū
salad inside have pig
There is pig in the salad.

b. shālā lǐ yǒu pìngguǒ
salad inside have apple
There is apple in the salad.

(39a) is like (38): (39a) can only be interpreted as expressing that there is a whole pig in the salad, i.e. a plate with a pig (presumably with an apple in its mouth) dressed up with lettuce leaves and other salad goods, covered in thousand island dressing.

On the other hand, (39b) can mean what the English paraphrase means: the salad has apple in it, and it doesn't have to be a whole apple, it can be apple pieces, grated apple, etc...

There is a natural account for the facts in (39), namely that foodstuff nouns like *píngguǒ* (apple) in Chinese are *ambiguous* between a [+N] reading (what hangs from the tree) and a [–N] reading (what is eaten in apple sauce). On that assumption, the mass reading we observe for foodstuff nouns is not a fission reading, but an authentic *mess mass* reading.

I want to propose something stronger here, namely, that foodstuff nouns are ambiguous, not only in Chinese, but in English and Dutch as well.

Ambiguity Assumption:

	[–N]	[+N, –C]	[+C]
English:	*meat*		*dog*
	apple		*apple*
Chinese:	*ròu* (meat)	*gǒu* (dog)	
	píngguǒ (apple)	*píngguǒ*(apple)	

With Rothstein's last resort assumption for fission readings, this predicts that food-stuff nouns have a mess mass reading in all three languages, but no fission reading.

What is the difference?

In section 1 of this paper, I argued against homeopathic semantics for mess mass nouns like *salt* and *triangle patterned wallpaper*: mess mass nouns, I argued, do not have a homeopathic interpretation, there are lexical and contextual constraints on what counts as salt and on when the salt is becoming too small to be split into two parts that both count as salt.

On the other hand, in the previous section I gave a semantics for the fission interpretation of *dog*: $\downarrow(DOG) = <(\downarrow(\mathrm{DOG})], (\downarrow(\mathrm{DOG})]-\{0\}>$, the set of all Boolean parts of the split of the set of all individual dogs. This means that the fission interpretation of *dog* is homeopathic,

in that it doesn't put constraints on what counts as fission *dog*, more than that it is mass part of the sum of all dogs.

The prediction, then, is that fission readings are *homeopathic*, closed under arbitrary mass parts. For fission readings, like those in (37a,b) repeated here, what there has to be in the dish/soup to make the statement true can be manipulated in context to an extreme degree:

(40) a. There is human in this dish.
b. There is cat in this soup.
$\exists x[\ x \sqsubseteq \downarrow_o(\mathrm{CAT}) \wedge \textit{in the soup}(x)]$

Thus, normally you will utter (37a,b) if you detect human flesh in the dish or cat flesh in the soup. But I may say (37b) with disapproval if I find a piece of fingernail in the dish, or fish a cat hair out of my soup. In a mythological context, if I, to test the Wisdom of the Gods, take something from the body of Pelops, so small that we ordinary humans would not be able to detect it, still in the Myth, Zeus will thunder at me: *there is human in this dish*, and condemn me to the Tartarus.

Thus, in context, certain parts may be regarded as too small to be considered as parts that matter, but the context can be manipulated (as in the Zeus example) to include arbitrarily small parts. And the flexibility here can be extreme.

Suppose Harold comes into the kitchen, proudly shows us the gall stone they have removed from him, and drops it by accident in the soup. You fish it out and I say:

(41) I am not going to eat that soup, it has had Harold in it.

This is, of course, funny, but so are all the other fission examples, and the thing that gets stretched in this example is: 'what counts as a contextually relevant part of Harold. Lexical mass nouns are not homeopathic. In order for (41) to be true, there has to be meat in the soup and not just something that is part of the meat.

(41) There is meat in the soup.

For instance, in the near future white calf-meat may come on the market that consists 70% of hormones. Suppose I extract the hormones from calf-meat, put them in a jar, and make a soup for the yearly dinner party of the Body-builders Club, and I scoop a considerable amount of these hormones into the soup.

The vegetarians among the body-builders may regard the soup as not suited for them, because not only do they not eat meat, but they also try not to eat additives for the production of which animals have been killed. Yet, this doesn't mean that (41) is true. For (41) to be true there has to be meat in the soup, and not just a set of chemicals derived from a meat source.

The difference between these two cases is instructive. The lexical mass nouns, by their meaning, put constraints on their denotation: if something is in the denotation of meat, only those of its parts are in the denotation of *meat* that themselves satisfy the criteria for counting as meat.

Fission interpretations, as we have seen, are different: there are no other lexical constraints for being in the denotation of the fission of *dog* than being part of the split of dogs, and contextual salience. This shows that *if* you think (which I don't) that count nouns like *dog* have mess mass interpretations, or interpretations unspecified for count and mass, like Pelletier's interpretations, or Rothstein's root-noun interpretations, these interpretations are *different* from fission interpretations. The reason for this is that a mess mass interpretation or root interpretation would put *semantic* constraints on the noun denotation, *lexical constraints*, and that is just what we don't find for fission interpretations.

We now look at (42):

(42) There is apple in the salad.

Again, genetically modified apples may come on the market that consist 70% of hormones, the same hormones as contained in the calf-meat. At the same dinner, I mix the apple-derived hormones into the salad. Just as (41) is not true, (42) is not true. For (42) to be true, it is not enough that there is part of the apples in the salad, it has to be part that *itself* counts as apple.

Thus, *apple* patterns with mess mass nouns like *meat*, suggesting strongly that also in English, foodstuff nouns like *apple* are ambiguous between a count interpretation and a mess mass interpretation.

Now look at the examples in (43):

(43) a. Er zit *kleine hond* in de salade
There is *small dog* in the salad
b. Er zit *grote hond* in de salade
There is *big dog* in the salad

(44) a. #Er zit *grote appel* in de salade
There is *big apple* in the salad
b. #Er zit *grote wortel* in de salade
#There is *big carrot* in the salad

The examples in (43) are felicitous, those in (44) are not. Why would this be?
Let us make the following assumption:

Assumption: the modified nouns *small N*, *big N* derive their fission behavior from the head noun.

With this, we argue as follows. *Small dog* and *big dog* are count noun phrases in a context where *dog* can only have a fission interpretation. By the assumption made above, *big dog* and *small dog* can also have a fission interpretation: the fission of *small dog* is the set of Boolean parts of the split of *small dogs*. This means that what (43a) and (43b) express is similar to:

(45) a. There is Chihuahua in the salad
b. There is Doberman in the salad.

We come to (44). *Apple* and *carrot* do not have a fission interpretation, they have a count interpretation and a mess mass interpretation. By the assumption made above, *big apple* and *big carrot* do not have a fission interpretation either.

But this means that *big apple* and *big carrot* in (44) can only be analyzed as:
[*big* [*apple*$_{[-N]}$]] and [*big* [*carrot*$_{[-N]}$]]. But we know that distributive adjectives like *big* are not very felicitous with mess mass nouns, and this is why the cases in (44) are not good.

Note crucially that the reading we saw for the cases in (43) is patently absent for the cases in (44). Dutch *winter wortels* (*winter-carrots*) are huge. The noun *winter wortel* itself is like *carrot* in that it is food stuff that allows both a count and a mess mass reading. We find a robust contrast between the examples in (46):

(46) a. #Er zit *grote wortel* in de salade
There is *big carrot* in the salad
b. ✓Er zit *winterwortel* in de salade
There is *winter-carrot* in the salad

If I grate a winter-carrot and put the result in the salad, (46b) is perfectly felicitous, but (46a) is terrible: (46a) just cannot mean that there is stuff derived from the split of *big carrots* in the salad.

This strongly supports the distinctions made here and argues strongly against theories in which fission is a simple operation lifting the count nature of the noun, semantically doing not much more than removing a bit of contextual restriction. On such a theory, there is no rationale whatsoever for the contrast between (46a) and (46b).

11. THE NEATNESS OF FISSION READINGS

11.1. The problem.

The fission interpretation of count nouns like *dog* as ↓(*DOG*) has a problem, as can be observed in the picture in section 9: ↓(*DOG*) is mass all right, but it is also *neat*. The reason is that **BOOL** is a complete *atomic* Boolean algebra, and ↓(*DOG*) is closed downwards, hence **min**(↓(*DOG*)) = **atom**(↓(*DOG*)), and hence **min**(↓(*DOG*)) is disjoint.

Now we have been assuming the Two Feature System, in which the features **neat** and **individuated** coincide. If so, it follows that the fission interpretation of *dog*, ↓(*DOG*), is **individuated**. But that means that the fission interpretation should allow distributive adjectives like *small* and *big*, with interpretations that distribute to the neat (individuated) generators. This means that we predict that (43a) has an *alternative* analysis:

(43) a. Er zit *kleine hond* in de salade
There is *small dog* in the salad.
$\exists x \in \downarrow(\text{DOG}) \cap {*}\mathbf{small}$**:** ***in the salad***(x)

On this interpretation, (43a) expresses that there are small generator parts of the split of *dog* in the salad. The problem is that (43a) of course doesn't have such an interpretation.

The diagnosis is that the fission ↓(*DOG*) should be [–I]. I briefly mention a few ways of solving this problem:

11.2. The Three Feature System.

We can move from the Two Feature System to a Three Feature System. In such a system, we do not make the assumption that the generators of neat nouns are necessarily individuated. (We *do* continue to make the mess mass assumption: [–N] ⇒ [–I], i.e. mess is non-individuated.) In the Three Feature System we have the following categories:

[+C +N +I]	= [+C]	count:	*boys, peas*
[–C –N –I]	= [–N]	mess mass:	*meat, cheese*
[–C +N +I]	= [–C, +I]	individuated mass:	*furniture, kitchenware*
[–C +N –I]	= [–C, +N, –I]	fission mass:	*↓(DOG)*

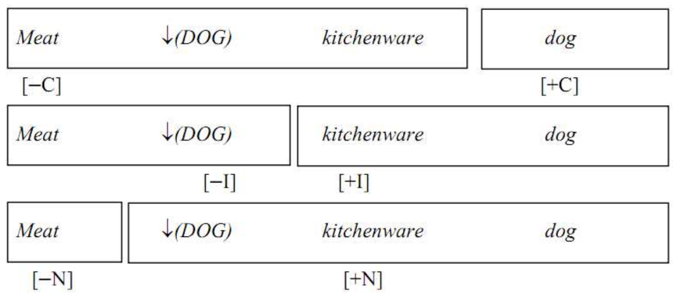

In this theory, there is a new category, [+N,–I], with neat minimal generators that are not individuated. Fission interpretations are of this category, but lexical nouns are not.

It is certainly possible to work in such a theory, but it is also a bit disappointing. The Two Feature System has a conceptual elegance that the Three Feature System lacks: in the Two Feature System the semantically relevant features are all defined in terms of the conceptual algebra of part-of structures: part-of, minimal element, generator, overlap, sum, remainder,. . .

The theory can do without the feature which has the more complex, intensional definition (*individuated*), because the two are identified extensionally. If it turns out that we can't maintain the equivalence, we have to accept that, of course, but it would be attractive if we don't have to.

Also, empirically, we have this new linguistically relevant category [–C, +N, –I], which is not lexically inhabited in any language I know of. And the question is: why not? Why aren't there languages where there are lexical nouns of the category [–C, +N, –I]?

11.3. Fission$_k$.

An obvious alternative is to change the fission operation, which produces a neat set, to an operation whose output is mess, not neat.

This is simple enough to do: let context k select a subset of fission $\downarrow_k$(X) of $\downarrow$(X):

$$\textbf{Fission}_k\text{: } \downarrow_k(X) = <\downarrow_k(X), \textbf{gen}(\downarrow_k(X))>$$

where:
1. $\downarrow_k(X)$ is a regular set
2. $\downarrow_k(X) \subseteq \downarrow(X)$
3. $\sqcup(\downarrow_k(X)) = \sqcup(\downarrow(X))$
4. **gen**($\downarrow_k$(X)) is a set of *overlapping generators* for $\downarrow_k$(X)

This is illustrated in the following picture:
Easy as it is, this proposal has a major drawback: it makes $\downarrow_k$(*DOG*) a pretty normal *mess mass* noun. And that is a problem. We have just seen that fission interpretations *differ* from mess mass interpretations, so it is not clear how the differences we have discussed would be maintained in such a theory.

11.4. Super fission.

Fission breaks down an object into its homeopathic mass set, a neat Boolean algebra. The atoms of that Boolean algebra are the ultimate minimal parts in the structure M, according to the background Boolean algebra **BOOL**.

But what is the status of those postulated minimal parts in M? And why aren't these minimal parts in M themselves ground by fission?

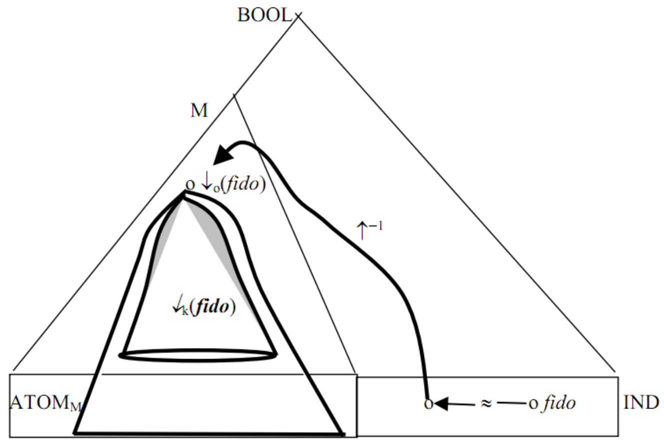

That is, the arguments that we have given, following Chierchia (1998), against atomless structures concerned the interpretations of *lexical mass nouns: mess mass nouns* and *neat mass nouns*. But that is not what we are talking about here at all, here we are talking about the question of whether *the whole structure* should be generated from a background set of 'ultimate atoms', and whether fission stops at those 'ultimate atoms'.

I propose an operation of *super fission*, which is fission that doesn't stop at the contextually provided postulated atoms in M, but breaks open such atoms.

We extend out interpretation domain **BOOL** to an interpretation domain **UNIVERSE**:

UNIVERSE = <**BOOL, SMASH**>where:

1. **BOOL** is, as before, a complete *atomic* Boolean algebra with atoms sorted into M-atoms and IND.

2. **SMASH** is a complete *atomless* Boolean algebra such that:

   1. BOOL ∩ SMASH = M
   2. for all m ∈ M: $(m]_{\textbf{BOOL}} \subseteq (m]_{\textbf{SMASH}}$

This means that **SMASH** is an atomless Boolean algebra with M as its top part, as in the following picture:

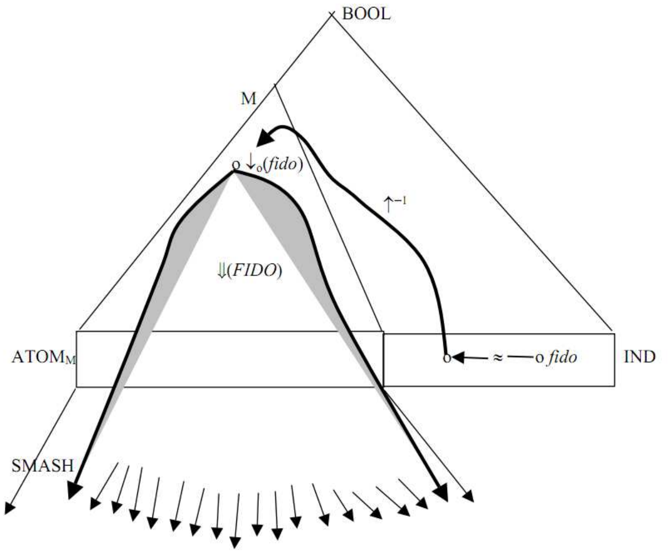

And I propose a super *fission operation*, which is like *fission* except that it takes all the Boolean SMASH-parts of the split of *dog*:

Super fission: $\Downarrow$(DOG) = <$\Downarrow$(DOG), $\Downarrow$(DOG)–{0}>
$\Downarrow$(DOG) = $(\downarrow_o(\text{DOG}))_{\textbf{SMASH}}$

The idea of superfission is that the constraints on what counts as a salient part of the superfission of the dogs are not given by the structure at all. It is only the context that decides whether something that is part in the widest sense of the split of dogs is salient enough to count as a contextually salient fission part. Using an atomless structure is to remind us that the background atoms of the structure M do not form a semantic constraint on the fission interpretation.

With this, the feature N now has three values:

neat:	[+N]	minimal generators do not overlap
mess:	[–N]	minimal generators overlap
superfine:	[# N]	minimal generators absent

The analysis changes only minimally from the Two Feature System:
-The fission interpretation $\Downarrow$(*DOG*) is superfine, which is homeopathic, and neither neat nor mess.
-For count nouns, neat nouns, and concrete mass nouns like *meat* and *salt* interpretation takes place in **BOOL** where only the values [±N] are available. This means that for such lexical nouns only the features [±C] and [±N] are available, as before, and [#N] is not. So nothing changes for these nouns.

What about *abstract* mass nouns?

Abstract mass nouns are all but absent in formal accounts of the semantics of mass nouns, and it is high time that their semantic properties are studied rigorously. I cannot at this point speculate about how they will fit into a theory like the one developed here. I do not know what the generators of *love* are (although love has arithmetic properties, as argued by Cordelia in the first scene of *King Lear*). I do not know whether denotations of abstract mass nouns are always atomic. Tarski, for one, would make a case that the mass interpretations of the abstract nouns *space* and *time* should be **superfine**, because Tarski developed the theory of atomless Boolean algebras and their standard model in the set of regular open sets as the natural background structure for three dimensional geometry.

In sum: the arguments against atomless structures concerned the interpretations of lexical mass and count nouns (excluding abstract nouns). Those arguments are accepted and maintained in the present theory. We maintain the Two Feature System, in which the intensional notion *individuated* is extensionally equated with the structural notion *neat*.

We spotted a problem: we must regard fission interpretations as *not neat*. But fission interpretations are not *mess*. We propose that, since fission interpretations are not lexically constrained anyway—which means that they are not constrained in terms of requirements on their generators—we can as well make these interpretations *ignore* the atoms that the model **BOOL** forces upon them, and hence make them atomless.

ACKNOWLEDGEMENTS

This paper has shown an exceptionally long and tortuous gestation process. The idea to look in the semantics of mass nouns horizontally at overlapping variants, rather than vertically at whether atoms are there, developed in the course of a graduate seminar on mass nouns that I taught at Tel Aviv University in the spring of 2001.

In the years that followed, I rewrote the material into a new version about every two years. However, what I was writing turned out to be more and more a kind of post-modernistic paper, consisting solely of section-long footnotes to a non-existing running text. And the problem is: I don't use footnotes, so ultimately, the more I wrote, the more it became the deconstruction rather than construction of a paper.

A good indication of the problems I had in formulating what the paper was about, is that—though versions of the paper were circulating—I didn't know how to give a talk about it.

Finally, while being on leave in Amsterdam during 2008-2010, things started to move. Susan Rothstein was working on count nouns at the same elongated table where I was sitting, and our many discussions made me clarify what my story on mass nouns was, and made me realize that I should get my act together and start presenting it.

I gave the first presentation of this material at *Palmyr IX: Logic and the Use of Language* in June 2010 at ILCC in Amsterdam. After that I presented versions at the *Tiende Internationale Bijeenkomst van Docenten Neerlandistiek* in Tel Aviv in November of 2010, at *Formal Semantics and Pragmatics: Discourse, Context, and Models* in Riga in November 2010, in a mini-course on mass and count nouns co-taught by Susan Rothstein and me at the summer school of Abralin, the Brazilian Association for Linguistics in Curatiba in February of 2011, and in departmental colloquiums at Tel Aviv University and the Hebrew University of Jerusalem in November and December of 2010.

I thank all the audiences of these presentations for their comments, criticisms and discussions (or in some cases for their own presentations); here I especially want to mention Roberlei Alvez Bertucci, Greg Carlson, Edit Doron, Alex Grosu, Manfred Krifka, Xu Ping Li, Anna Müller, Barbara Partee, Jef Pelletier, Roberta Pires de Oliveira, Dafna Rothstein Landman and Frank Veltman. Once more thanks to Barbara for her helpful comments on the final written version of this paper. This paper owes an obvious intellectual debt to Gennaro Chierchia.

Susan Rothstein reminded me that, when, in *Genesis*, Adam gives names to the animals, he doesn't have, as the Hebrew formulates it, 'a help *against* him'. In the next line God makes him fall asleep, takes a rib, etc. . . Over the last ten years, Susan Rothstein has been this paper's 'help against it'. Without her it probably wouldn't have come into existence. Susan's penetrating comments, her own work as it developed simultaneously, her dis- and encouragements alternatingly shocked it into stasis, and into feverish stages of speedy development. In the end, very little in this paper would be the way it is without her.

References

Barner, D. & Snedeker, J. 2005. 'Quantity judgements and individuation: evidence that mass nouns count'. *Cognition* 97: 41–66.

Bunt, H. 1985. *Mass Terms and Model Theoretic Semantics*. Cambridge University Press, Cambridge.

Cheng, L., Doetjes, J. & Sybesma, R. 2008. 'How universal is the Universal Grinder'. In: *Linguistics in the Netherlands* 2008, p. 50–62.

Chierchia, G. 1998. 'Plurality of mass nouns and the notion of semantic parameter'. In Susan Rothstein (ed.) 'Events and Grammar', 52–103. Kluwer, Dordrecht.

——. 2010. 'Mass nouns, vagueness, and semantic variation'. In Synthese 174 (1).

Doetjes, J. 1997. *Quantifiers and Selection*. Ph.D. thesis, University of Leiden.

Hoeksema, J. 1983. 'Plurality and conjunction'. In Alice ter Meulen (ed.) 'Studies in Modeltheoretic Semantics', 63–83. Foris, Dordrecht.

Krifka, M. 1989. 'Nominal reference, temporal constitution and quantification in event semantics'. In Johan van Benthem Renate Bartsch & Peter van Emde Boas (eds.) 'Semantics and Contextual Expression', 75–115. Foris, Dordrecht.

——. 2009. 'Counting configurations'. In Arnt Riester and Torgrim Solstad (eds.) *Proceedings of Sinn und Bedeutung* 13.

Landman, F. 1989. 'Groups I & II'. *Linguistics and Philosophy* 12: 559–605, 723–744.

——. 1991. *Structures for Semantics*. Kluwer, Dordrecht.

——. 2000. *Events and Plurality*. Kluwer, Dordrecht.

Lasersohn, P. 1988. *A semantics for Groups and Events*. Ph.D. thesis, University of Ohio, Columbus, Ohio.

Li, X. P. 1983. *On the Semantics of Classifiers in Chinese*. Ph.D. thesis, University of Ohio, Columbus, Ohio.

Link, G. 1983. 'The logical analysis of plurals and mass terms: a lattice-theoretic approach'. In R. Bäuerle, U. Egli & A. von Stechow (eds.) 'Meaning, Use and the Interpretation of Language', 303–323. de Gruyter, Berlin.

——. 1984. 'Hydras. On the logic of relative constructions with multiple heads'. In F. Landman & F. Veltman (eds.) 'Varieties of Formal Semantics', 245–257. Kluwer, Dordrecht.

Lønning, J-T. 1987. 'Mass terms and quantification'. *Linguistics and Philosophy* 10: 1–52.

Pelletier, F. J. 1975. 'Non-singular reference. Some preliminaries'. In *Philosophia 5*. Reprinted in Francis Jeffry Pelletier (ed.), 1979, *Mass Terms. Some Philosophical Problems*, Kluwer, Dordrecht.

Pelletier, F. J. & Schubert, L. 1989/2002. 'Mass Expressions'. In Dov Gabbay & Franz Guenthner (eds.) 'The Handbook of Philosophical Logic', 387–453. Kluwer, Dordrecht.

Rothstein, S. 1992. 'Types of plural individuals'. *Journal of Semantics* 15: 641–675.

——. 2009a. 'Bare noun semantics'. Talk presented at the workshop on 'Bare Nouns: Syntactic Projections and their Interpretations', Paris, November 2009.

——. 2009b. 'Stubborn distributivity, multiparticipant nouns and the count/mass distinction'. To appear in: *Proceedings of NELS 39*.

——. 2010. 'Counting and the mass-count distinction'. *Journal of Semantics* 27: 343–397.

Sybesma, R. 2009. *Het Chinees en het Nederlands zijn eigenlijk hetzelfde*. Het Spectum, Houten.

ter Meulen, A. 1980. *Substances, Quantities and Individuals*. Ph.D. thesis, Stanford University.