# Systematic polysemy and the count-mass distinction\*

# Ingrid Lossius Falkum

#### **Abstract**

This paper investigates the type of systematic sense alternation that rests on the count-mass distinction in English. Computational semantic approaches often suggest treating such sense alterations as being generated on the basis of a set of lexical inference rules, yielding, for instance, for an animal-denoting term, a meat sense ('We're having rabbit for dinner'), a fur sense ('Annie is wearing rabbit'), and a general 'stuff' sense ('After a tractor had run over the body, there was rabbit splattered all over the yard'). I question this rule-based account on the basis that it doesn't provide the interpretive flexibility required to account for the data, and argue that, although the sense alternations in question clearly have a linguistic component, pragmatics is playing a crucial role in giving rise to such sense alternations. I propose an alternative analysis, where the countmass distinction is seen as a primarily conceptual-ontological distinction being reflected in a syntactic distinction at the level of NPs (rather than as a property of nouns). On this account, count-mass syntax provides an instruction to the pragmatic system that the concept encoded by the NP should be construed as denoting either an individual or an unindividuated entity. On the basis of such underspecified inputs, highly activated encyclopaedic and logical information associated with the concept, and contextual assumptions derived from the utterance situation, the relevance-optimising pragmatic system operates to yield the speaker-intended concept. Keywords: systematic polysemy, relevance theory, count-mass distinction, concepts, interpretation, computational semantics

#### 1 Introduction

There are instances of polysemy where the related senses of a word are predictable on the basis of a general pattern of sense alternation observed for words denoting objects of the same category. The generalisation involved has the following form.

(1) If an expression has a use of type A, it also has a use of type A'.

This type of polysemy goes by various names, including 'regular polysemy, (Apresjan 1974; Kilgarriff 1992, 1995), 'semantic transfer rules' (Leech 1990 [1974]), 'lexical implication rules' (Ostler and Atkins 1992), 'logical polysemy' (Pustejovsky 1991, 1995), 'transfers of meaning' (Nunberg 1996, 2004), 'sense extensions' and 'conversion' (Gillon 1999). Nouns are the most salient (and most discussed) examples of this type of category, and I will concentrate on these in what follows. Here are some examples:

- (2) a. A *rabbit* jumped over the fence.
  - b. We're having *rabbit* for dinner.
  - c. The model wore *rabbit* on the catwalk.

\* I am grateful to Robyn Carston for many valuable discussions on the topic of this paper and for constructive comments and suggestions.

<sup>&</sup>lt;sup>1</sup> As to other parts of speech claimed to exhibit a type of systematic polysemy, see for instance the recent paper by Kennedy and McNally (2010), suggesting that colour adjectives alternate between a gradable and a non-gradable sense, and Jackendoff's (2002: Chapter 11) analysis of the causative alternation (e.g. The door *opened*/John *opened* the door) as a form of polysemy.

d. After a tractor had run over the body, there was *rabbit* splattered all over the yard.

- (3) a. We have a *pine* in our garden.
  - b. This table is made of *pine*.
- (4) a. Susan decorated the cake with a *cherry*.
  - b. When the kids left, there was *cherry* all over the kitchen floor.
  - c. Jill and Joan have a *cherry* in their garden.
  - d. This table is made of *cherry*.

In (2)a, the noun *rabbit* is used to denote the animal, its meat in (2)b, its fur in (2)c, and unspecified 'rabbit stuff' in (2)d. The noun *pine* is used to refer to a tree in (3)a, and to the food from the tree in (3)b. In (4)a, the noun *cherry* is used to denote the fruit, stuff of the fruit in (4)b, the tree carrying the fruit in (4)c, and the wood from the tree carrying the fruit in (4)d. On the basis of these examples, we may identify the following lexical alternation patterns, described by means of the formula in (1):

- (5) a. If an expression has an 'animal' use, it also has a 'meat'/'fur'/'animal stuff' use.
  - b. If an expression has a 'tree' use, it also has a 'wood' use.
  - c. If an expression has a 'fruit' use, it also has a 'fruit stuff'/'tree' use.

And there are many more such lexical alternation patterns (see e.g. Lakoff and Johnson 1980: 38-39; Nunberg 1996: 117 for further examples). Much work in computational semantics has emphasised the need for an adequate formalisation of these observed regularities. The literature contains various suggestions of how they can be incorporated into the lexicon as a set of lexical inference rules, with the aim of avoiding a listing of all options for all words (Pustejovsky 1991, 1995; Ostler and Atkins 1992; Copestake and Briscoe 1992, 1996; Boguraev and Levin 1993; Kilgarriff 1992, 1995; Kilgarriff and Gazdar 1995; Blutner 1998).

On the face of it, the arguments for treating systematic polysemy as being governed by linguistic rules seem well-founded. The processes in question appear to be regular, productive and available in many languages. Their application appears to be restricted by the existence of synonymous terms in the lexicon; so-called pre-emption or blocking (e.g. the existence of the lexical form *beef* in English 'blocks' the use of the form *cow* to refer to the meat of the animal, cf. Blutner (1998, 2002)). Some of them seem to affect syntactic behaviour, for instance, lexical alternations that appear to rest on the count-mass distinction in English and other languages with count-mass syntax, gender marking on the fruit for tree alternation in Romance languages (e.g. the distinction in Spanish between *aceituna* ('olive') and *aceituno* ('olive tree'), and *guinda* ('cherry') and *guindo* ('cherry tree').<sup>2</sup>

\_

<sup>&</sup>lt;sup>2</sup> Systematic polysemy is often seen as having an affinity with metonymy, the process whereby an expression that conventionally denotes one object or property is used to denote another object or property that stands in a certain relation to it.<sup>2</sup> A famous example is Nunberg's (1978: 22) 'The *ham sandwich* is sitting at table 20', where *ham sandwich* refers to a customer in the context of a restaurant. Some of the sense alternations described as instances of systematic polysemy arguably involve metonymic relations between the different senses, (cf. (4). For this reason, formal semantic accounts have proposed to treat not only these but also more creative cases of metonymy as having a non-pragmatic component (Copestake and Briscoe 1996). There is some evidence that creative cases may have grammatical effects, e.g., 'The *french fries* is getting impatient' (Nunberg 1996), where the agreement between the subject and the VP is determined by the referent of the NP *french fries* rather than by its syntax. The cognitive grammar tradition (Lakoff and Johnson 1980; Lakoff 1987 and many others) usually makes no distinction between systematic polysemy that involves metonymic sense relations and metonymy more generally; the label 'metonymy' subsumes them both. In his pragmatic account, Nunberg (1996) takes the process of 'predicate transfer' to be responsible for the generation of systematic polysemy as well as creative metonymy. From a communicative point of view, however, such creative cases of metonymy seem largely

The focus of this paper will be instances of systematic polysemy that (in English) rest on the distinction between count and mass uses of nouns (and thus are perhaps the clearest candidates for a linguistic analysis), as exemplified by (2), (3) and (4) above. I have two aims. The first is to consider a well-developed formal account of systematic polysemy, proposed by Copestake and Briscoe (1996), and assess whether systematic polysemy can be adequately accounted for as instantiations of lexical rules. I conclude that, although such an account captures the regularities involved, it does not provide the interpretive flexibility required to handle the full range of data. Thus, my second aim is to show how the data can be reanalysed in pragmatic terms, more specifically from the perspective of relevance theory (Sperber and Wilson 1986/1995, Carston 2002). I will defend the view that, although the sense alternations that give rise to systematic polysemy clearly have a linguistic component, the contribution of the linguistic system to its generation and interpretation is less central than is often thought. My conclusion will be that even in cases of systematic polysemy, the major interpretive work is done by pragmatics.

### 2 Computational semantic accounts: the case of 'grinding'

Copestake and Briscoe (1992; 1996: 36) propose a formal semantic framework in which *sense extensions*, that is, "predictable creation of different but related senses", are represented as lexical rules. One such rule is 'universal grinding' (Pelletier 1975). The effect of this rule is to create from a count noun (denoting a physical object) a mass noun with properties appropriate for an unindividuated substance. This is a general, abstract rule that yields the mass senses of *rabbit* and *cherry* in (6) and (7) (cf. (2) and (4) above), as well as the uses of *sun* and *shopping centre* in (8) and (9) to denote unindividuated substances:

- (6) After a tractor had run over the body, there was *rabbit* splattered all over the yard.
- (7) When the kids left, there was *cherry* all over the kitchen floor.
- (8) We got quite dark from all the *sun*.
- (9) That's a lot of *shopping centre* for a small town (Nunberg and Zaenen 1992).

Copestake and Briscoe further posit a set of conventionalised sub-cases of the grinding rule, including a specialised 'meat-grinding' function that forms food-denoting mass nouns from animal-denoting count nouns, and one that forms fur-denoting mass nouns from animal-denoting count nouns. Kilgarriff (1992, 1995) proposes a similar specialised rule that yields the wood senses of tree-denoting count nouns. Together, these rules account for the examples in (10)-(12) below:

- (10) We're having *rabbit* for dinner (*chicken*, *turkey*, *moose*, *frog*, etc.).
- (11) The model wore *rabbit* on the catwalk (*mink*, *beaver*, *calf*, *lizard*, *crocodile*, etc.).
- (12) This table is made of *pine* (*cherry*, *oak*, *chestnut*, *birch*, etc.).

confined to referential uses (in a strong sense) and bear some relation to certain nicknames (e.g. 'Four Eyes is always reading math books'). I cannot get into the topic of metonymic interpretation any further here, but see Falkum (2010) for a suggestion for a relevance-theoretic approach to these cases.

<sup>&</sup>lt;sup>3</sup> Pelletier's (1975: 456-457) original thought experiment involves a machine, the 'universal grinder' (not unlike a meat grinder), which chops up "any object no matter how large, no matter how small, no matter how hard". Then, the idea is that we could take any object corresponding to any (apparent) count noun, let's say *man*, put it into the grinder and ask what is on the floor at the other end of the grinder and get the following answer: 'There is man all over the floor'. Given that "there can be made a *prima facie* case that nothing is immune from the grinder treatment", Pelletier concludes that counts nouns having physical objects as their extensions can always be given a mass sense.

So the claim is that there is a set of lexical rules (the universal grinder and its various conventionalised sub-cases) stored in the lexicons of individual speakers that are responsible for generating systematic polysemy of the kind illustrated above. In other words, the shift in meaning that is observed in these cases (e.g. from an animal sense to a food sense), is assumed to have a wholly linguistic basis. Copestake and Briscoe do point out, however, that pragmatics may contribute to further contextual specification of the denotation, for instance, by providing the information that rabbit meat usually excludes the bones while the meat of whitebait does not.<sup>4</sup>

What are the main arguments for taking this approach, and what are its advantages? Perhaps the most obvious advantage of the rule-based account is the explanation it offers for the apparent productivity of the sense alternations in question, i.e. the fact that they can be extended immediately to words that speakers encounter for the first time. For instance, if someone tells you that he saw an okapi in the Wild Animal Park in San Diego last summer, and this was the first time you heard of the animal as well as the word for it, *okapi*, you would have no difficulty understanding an utterance of '*okapi* tastes a bit like horse' (where *okapi* refers to the meat of the animal), although you would never have come across this use before. The computational semanticist view is that this suggests the existence of a productive rule, rather than several stored senses for a word (which they appear to see as the only alternative view).

However, the apparent productivity of such sense extensions need not have an entirely linguistic basis, although it is indeed a possible explanation for it. An alternative explanation would be to take the example above to be a case of the hearer exploiting his knowledge about certain regularities in the world (e.g. that animals are in general edible, that an animal's flesh is considered 'meat', etc.) in his interpretation of *okapi* as referring to the meat. Once he has learned that okapi is an animal, the inference from animal to meat would come quite naturally. This would also serve to explain the conventional cases. However, since this real-world distinction coincides with the grammatical distinction between count and mass nouns in English, there would also be a clear linguistic clue (i.e. that the noun occurs without a determiner) to this interpretation. I return to this point in section 3.2.

Another seemingly compelling argument in favour of the rule-based account is the parallel one might draw between the conventionalised cases of grinding and derivational morphological processes. In particular, what Copestake and Briscoe see as a striking similarity is that both processes appear to be 'blocked' by the existence of an underived synonymous lexical form (Aronoff 1976; Briscoe, Copestake, and Lascarides 1995). On this view, the existence of lexical forms such as *veal*, *pork* and *mutton* blocks the application of the meat-grinding rule, and thus explains why the following sentences strike us as odd:

- (13) ?Joan likes to eat *calf* (veal).
- (14) ?We're having *pig* for dinner (pork).
- (15) ?Matt is preparing *sheep* for our anniversary (mutton).

<sup>&</sup>lt;sup>4</sup> In the discussion to follow, it should be borne in mind that the primary goal of computational semantic accounts such as that of Copestake and Briscoe and others may not be to model how these sense alternations are manifested in actual communication, but to provide a formal account in which the process of operationalisation may be an end in itself. As such, the computational semantic accounts and the cognitive pragmatic account of utterance comprehension that will be pursued in this paper rest on substantially different methodological approaches. However, since the computational approach has proved to be so influential in accounts of systematic polysemy of the kind discussed in this chapter, it seems in my view appropriate to also assess its explanatory potential in accounting for how systematic polysemy is manifested in actual communication.

<sup>&</sup>lt;sup>5</sup> Clark and Clark (1979) describe this as 'pre-emption by synonymy'.

Such lexicalised forms can be analysed as lexical exceptions to the meat-grinding rule (Ostler and Atkins 1992; Copestake and Briscoe 1996), by analogy with exceptions to morphological processes. For instance, Aronoff (1976) notes that the form \*gloriosity does not appear in English (contrary to e.g. curiosity, atrocity, etc.), which he attributes to the fact that it would be synonymous with the existing form glory and therefore blocked from being generated. However, unlike most exceptions to morphological patterns, which are usually considered ill-formed and do not occur in normal language use, the derived forms of animal terms such as calf, pig, sheep and cow seem to co-exist happily with the lexicalised forms veal, pork, sheep, and beef. Consider the following examples:

- There were five thousand extremely loud people on the floor eager to tear into roast *cow* with both hands and wash it down with bourbon and whiskey (passage from Tom Wolfe's The Right Stuff (1979), cited by Copestake and Briscoe 1996: 38, my italics).
- (17) Hindus are forbidden to eat *cow* (?beef) (Nunberg and Zaenen 1992).

Copestake and Briscoe explain such cases in terms of non-synonymy of the ground and the lexicalised forms; in (16) and (17) *cow* is possible because it is not synonymous with *beef*. Their proposal is that when such uses occur, "they either convey a negative attitude to the consumption of the meat on the part of the speaker or an entailment of extended denotation, where more of the cow ... than is normally considered 'meat' is being treated as food" (1996: 38-39). They further suggest that the utterances (16) and (17) can be analysed as apparent violations of the Gricean maxim of Manner, that is, as choices of a less common or interpretable form from among different ways of expressing the same meaning, which, as they put it, "carries the (discourse) implication that the terms are not synonymous". This gives rise to the interpretation that the speaker has expressed a negative attitude towards the consumption of the meat in (16), and the interpretation that the interdiction concerns the animal as a whole (due to its status in the Hindu religion) and not just its meat in (17). This use is clearly more informative than the possible utterance containing 'beef', which would not (at least not as easily) give access to this information.

However, there are many uses of derived forms which appear to be difficult to explain in terms of apparent maxim violation. For instance, in (18) both *pork* and *pig* would be acceptable, and in (19) they seem to be used interchangeably:

- (18) Jews and Muslims don't eat *pork/pig*.
- (19) As a general rule, I don't eat *pork*. This can be awkward I often go for the veggie option, to avoid having to explain why I don't eat *pig*.<sup>7</sup>

<sup>6</sup> Nunberg and Zaenen (1992) propose a similar account, where such uses are explained in terms of apparent violation of the Gricean maxim of Quantity (roughly: say as much as and no more than the communicative circumstances require) In the cases where the speaker has chosen to use the word *cow* instead of *beef*, the hearer is entitled to infer that she has some reason for using this vaguer term, as in like (17). Unlike the computational semanticists, however, Nunberg and Zaenen do not think that the blocking phenomenon can be explained in entirely semantic terms. Blutner (1998), working within a neo-Gricean framework, treats examples like (17) above as instances of Horn's (1984) 'division of pragmatic labour' (unmarked forms tend to be used for unmarked situations and marked forms for marked situations), explained in terms of a Q-principle (corresponding to the first part of Grice's Quantity maxim (make your contributions as informative as required)) and an I-principle (subsuming the second part of Grice's Quantity maxim (do not make your contribution more informative than required), the maxim of Relation and (possibly) all of the Manner maxims).

.

<sup>&</sup>lt;sup>7</sup> Attested: http://www.inventio.co.uk/pigs.htm

This does not mean that English speakers may not perceive the uses of *pig* in the examples above as being somehow more 'marked' than the conventional expression *pork* (although considerably less so than the uses of *cow* in (16) and (17)), to the extent that they may give rise to some additional 'effects'. However, these uses seem so natural that, in my view, it is counterintuitive to assume that they represent (apparent) violations of a conversational norm (contrary to e.g. figures of speech, where this is more evident). Moreover, the uses of *sheep* and *pig* in (20) and (21) below are entirely conventional:

- (20) Kate [with a plate of food in front of her]: This roast sheep is the best I've ever had.
- (21) I love roast suckling *pig*.

Copestake and Briscoe could argue that although these uses are conventional, they carry with them an entailment that it is the whole animal that has been roasted, and not just any unindividuated portion of its meat. This may be right, but in that case it seems that the entailment cannot be explained as resulting from the use of a derived form instead of the lexicalised one, as the following derived uses may, given the appropriate context, carry the same kind of entailment (i.e. involving a whole roast chicken/turkey):

(22) We're having roast *chicken/turkey* for dinner.

This is not to say that the lexicon may not be sensitive to the frequency of use of such (quasi) synonymous forms as *cow/beef*, *pig/pork*, etc., to the extent that the use of e.g. *cow* instead of the more frequent *beef* to denote the meat of the animal may be interpreted as carrying some additional meaning. However, the effect does not have to come from a semantic restriction on the use of *cow* due to the existence of the lexicalised form. An alternative approach is to analyse the so-called blocking phenomenon in terms of conventions of use (what Morgan 1978 terms 'conventions about language') where the use of a non-lexicalised form in cases where there exists a lexicalised one (e.g. *beef*, *veal*, *mutton*) induces a greater processing load on the hearer's pragmatic system which is then offset by the derivation of extra cognitive effects, rather than treating it as a semantic restriction on the application of a lexical rule. I return to this issue in section 3.2.

Another matter pointing to a pragmatic explanation of meat grinding cases is that a speaker's choice of a ground or an unground form to refer to the meat of an animal appears to be governed by real world knowledge. Consider the examples in (23)-(25):

- (23) Wolves eat *lambs* (?lamb).
- (24) Snakes eat *mice* (?mouse).
- (25) In Africa, lions eat *wildebeest*, *warthogs*, *zebras*, *buffalo* and different types of antelope.<sup>8</sup>

Here the nouns are all used in their plural (i.e. count) senses, which is the conventional way of talking about what animals eat. The uses of *lamb* and *chicken* in (26), however, are straightforward cases of meat-grinding:

(26) Indians eat *lamb* and *chicken* (?lambs and chickens).

.

 $<sup>^{8}\</sup> http://www.catalogs.com/info/gadgets/what-do-lions-eat.html$ 

An intuitive explanation for the difference between these examples is the different ways in which humans and other animals get hold of their food. Unlike other carnivorous animals, humans usually do not hunt down the animals they eat themselves in their daily lives and, even when they do, they don't devour the uncooked flesh. It seems that the use of a ground sense somehow entails that the meat has been prepared in some human-like fashion (which is why e.g. 'wolves eat *mutton*' would sound odd). This does not seem to be the case for the 'unground' forms in (23)-(25), where although it is the meat/flesh sense that is being communicated, the count form of the nouns allows the hearer to access his representations of the whole animals (which may give access to assumptions about the animals being hunted down as prey). The question, then, is this: Given identical linguistic environments, is it reasonable to assume that the meat senses in (26) result from the application of a rule of meat grinding, while the meat senses in (23)-(25) are pragmatically derived? This asymmetry is, of course, possible, but one might ask if lexical inferences that are clearly influenced by world knowledge in this way are not better treated along pragmatic lines. <sup>10</sup>

An argument that is often given in favour of treating systematic polysemy in terms of lexical rule application is the availability of so-called default interpretations in uninformative (or 'null') contexts (Pustejovsky 1995; Copestake and Briscoe 1996). For instance, the most accessible interpretation of (27) below seems to be the one according to which *rabbit* denotes 'rabbit meat' (and thus that Sam enjoyed but later regretted *eating* the rabbit):

(27) Sam enjoyed but later regretted *the rabbit* (Copestake and Briscoe 1996: 42).

Pustejovsky (1995) analyses such cases as type coercions, where a verb that subcategorises for an NP or a progressive VP syntactically, semantically requires a complement with an eventive interpretation. In the case where this requirement is not satisfied by the surface syntactic structure, a type-coercion operator changes the denotation of the NP from an entity into an event consistent with eventive information (represented in the form of a 'telic role') stored in the lexical representation for the noun (the so-called qualia structure). Copestake and Briscoe's (1996) account of the semantic processing of (27) is similar, and involves the selection of an appropriate aspect of the meaning of the complement (in this case its telic

\_

<sup>&</sup>lt;sup>9</sup> Ostler and Atkins's (1992: 84) discuss grinding applied to words denoting food items. This is only possible, they claim, when "the units of that food are not evident", as in e.g. some egg, some crab, some salmon, some potato, typically referring to some food substance on a plate, etc. They further note that words denoting pulses appear to resist grinding, so rather than saying 'Have some pea/bean/lentil' a speaker would say 'Have some peas/beans/lentils'. This they treat as a semantically-based constraint on the rule that turns food item denotations into food mass denotations (although, as far as I can tell, they do not specify what aspect of the semantics of pulse denoting words that is responsible for the existence of the constraint). However, exceptions to this 'constraint' regularly occur. In appropriate contexts, the 'ground' senses of words denoting pulses are unproblematic, as in utterances such as 'I just love refried bean smeared on a (http://eatingindallas.wordpress.com/2009/07/26/battling-breakfast-tacos), 'I'd like the grilled salmon fillet with mashed pea', 'Can I have a side-dish of dhal, please?', etc. How would these be explained in a rule-based account? As exceptions to a constraint on the application of a rule? In my view, the tendency to prefer the plural form of the pulse denoting noun over the mass form may rather have its explanation in the size of the food items in question - we seldom eat a single pea/bean/lentil/etc. - or anyway in some other real-world fact about pulses. In other words, this tendency appears to have a pragmatic explanation; one which obviously doesn't exclude the possibility of such words occurring with a mass interpretation, as in the examples above.

<sup>&</sup>lt;sup>10</sup> The examples in (23)-(25), are, of course, not exceptions to the rule of meat-grinding, as the direction of the lexical inference goes from animal-denoting count nouns to food denoting mass nouns, and not the other way around (i.e. from a food denotation to the use of a mass term), and so the rule of meat-grinding could still be valid. However, the examples do show that the basis for animal terms taking on food denotations cannot be entirely linguistic.

role: rabbits are for eating (by humans)). This, they claim, explains the 'default' character of the eating interpretation and militates against a pragmatic analysis. As they write, "the *meat-grinding* sense of *rabbit* provides a telic role which allows the eating interpretation to be constructed. However, if the lexicon does not propose such a sense, it is unclear what it is about the context which allows pragmatic specialization of the interpretation" (1996: 42).

I have two points to make about this claim. First, I cannot see how this explains the default character of the eating interpretation, when, on Copestake and Briscoe's account, there are in fact three possible interpretations of *rabbit* made available by the lexicon: (i) the general ground sense ('rabbit stuff'); (ii) the meat-grinding sense; and (iii) the fur-grinding sense. It is not clear to me why, on this lexicon-based account, it is the meat-grinding sense that is selected (which then gives rise to the eating interpretation), and not any of the other possible senses. 11 Second, I think the argument that the availability of default interpretations in uninformative contexts shows the need for a semantic analysis considerably underestimates the fact that hearers rarely (if ever) come to the interpretation process 'empty handed', as it were. There are hardly any entirely context-free interpretations. As emphasised by relevance theory, one of the tasks the hearer has to solve in utterance comprehension is constructing a hypothesis about the contextual assumptions to be brought to bear in the process. This set of assumptions - a subset of his assumptions about the world - could include assumptions derived from the observation of the physical environment, encyclopaedic knowledge, memories and beliefs as well as the preceding linguistic context (Sperber and Wilson 1986/1995; Wilson and Sperber 2004). When the assumptions that the hearer may derive from the linguistic and extra-linguistic context are scarce, as in (27), he will have to rely more on information stored in his long-term memory in the interpretation of the utterance. The linguistic meaning of the utterance in (27) allows him to access information stored in long-term memory about the concept RABBIT, among which the assumption that 'rabbits are delicious to eat' would be highly accessible. This assumption, in combination with assumptions that the other concepts in the sentence may give access to (for instance, that 'eating delicious food is an enjoyable activity' which may have been made accessible by the concept ENJOY), leads to the eating interpretation being much more accessible that any of the other possible interpretations. This is what gives it a 'default character'. However, its availability can be given a straightforward pragmatic explanation.

#### 3 A relevance-theoretic account

In my proposal to follow for a (mainly) pragmatic account of systematic polysemy, I would like to start by considering a claim that Fodor and Lepore (2002: 117) make in their *Compositionality Papers* a propos the animal-meat alternation. Here they say, "Opportunities for polysemy arise from *how things are in the world* (or, anyhow, from how we take them to be)". Fodor and Lepore are sceptical about the possibility of words being *lexically* 

<sup>&</sup>lt;sup>11</sup> This shows a consequence of the account that Copestake and Briscoe themselves point out (1996: 56); that the lexical rules for sense extension lead to overgeneration. For instance, on their account, the sentence '*Rabbit* is expensive these days' is, given the universal grinder, the rules of animal-meat grinding and animal-fur grinding, three-ways ambiguous between the general 'rabbit stuff' sense, the meat sense and the fur sense. While the overgeneration is in itself problematic (it does not seem psychologically plausible that all the senses are generated), there is also a question of how the ambiguity is resolved, that is, how hearers decide when one rule has prevalence over the others. A possibility would be to say that all three senses are initially generated by the linguistic system and then leave it to pragmatics to select the correct one in the particular context. This would mean, however, that in the derivation of the 'default' eating interpretation of (27), pragmatics would have to intrude to yield the meat-grinding sense of *rabbit*, which in turn would allow for the eating interpretation to be constructed.

polysemous. On their view, the reason that, e.g., lamb may be polysemous between the animal and the meat is that "lamb-the-meat comes from lamb-the-animal", hence rests on a real relation among things in the world. Fodor and Lepore do not develop their position in any detail, but in my view their claim touches on what I see as the very essence of the polysemy phenomenon. I hope what I mean by this will become clear in the following sections.

The type of systematic polysemy giving rise to the animal-meat, animal-fur, tree-wood (and so forth) alternations in meaning is, as we saw in the previous section, closely linked to the well-studied grammatical distinction between count and mass expressions in English (and other languages with count-mass syntax). There is a vast literature on this topic, the details of which I cannot get into here. However, in the next section I will give a brief overview of a set of issues that have been raised concerning the count-mass distinction, before I present the relevance-theoretic account in section 3.2.

## 3.1 The count-mass distinction and systematic polysemy

According to common intuition, 'count' expressions describe individual entities (e.g. The horses were bought in the spring), or kinds of individual entities (e.g. Horses are animals), while 'mass' expressions describe portions of quantities (e.g. There's some milk in the fridge), or kinds of quantities (Milk is healthy). Syntacticians have been interested in the morpho-syntactic characteristics underpinning the count-mass distinction. The usual view is that only count nouns admit of a contrast between the singular and the plural (e.g. horse/horses); mass nouns are almost always singular (e.g. milk, rice, advice, etc.). While count nouns can be modified by cardinal numbers (e.g. two horses) and by the quantifiers each, every, many, several, few and the stressed some; mass nouns occur with the quantifiers little, much, the unstressed some, as well as with measure phrases (e.g. a bottle of milk). Syntactically, common nouns are usually distinguished into count nouns and mass nouns by means of the feature [±count] (see for instance Gillon 1992, 1999).

Many philosophers (from Quine 1960 and onwards) have been interested in the semantic aspect of the count-mass distinction, viz. the ways in which count and mass expressions differ with respect to what they denote. Two criteria, those of cumulativity and divisivity of reference, have been proposed to distinguish the semantics of mass expressions from that of count expressions. The former criterion dates back to Quine (1960: 91), who noted that "mass terms like 'water' ... have the semantical property of referring cumulatively: any sum of parts which are water is water". Although describing a characteristic feature of mass expressions, it is not distinguishing of them as it also applies to plurals (i.e. count nouns). 12 The second criterion of 'divisive reference', proposed by Cheng (1973: 286-287), states that for any stuff/entity referred to by a mass noun, any part of that stuff/entity is also denoted by that same mass noun. Although this criterion clearly has some intuitive appeal, Gillon (1992, 1999) points out that it seems to be refuted by Quine's (1960: 99) earlier observation that "there are parts of water, sugar and furniture too small to count as water, sugar, furniture": The denotations of many apparent mass nouns have parts that do not fall under the denotation of the same noun. Furthermore, there is a large class of words that appear to be count nouns, yet satisfy the criterion of divisivity of reference (e.g. stone, rock, cord, rope, etc.). So having divisive reference can neither be a necessary nor sufficient characterisation of a mass expression.

<sup>12</sup> For instance, if the animals in my garden are rats and the animals in my neighbour's garden are also rats, then

the animals in our two gardens are rats (cf. Link 1983).

The distinction between mass and count syntax raises two psychological issues, concerning (i) how the distinction maps onto conceptual representation, and (ii) how children acquire the link between count-mass syntax and non-linguistic ontological categories. Traditionally, the grammatical count-mass distinction has been linked to the conceptual distinction between objects and substances (Macnamara 1982). Physical objects are typically count (e.g. horse, man, rock, etc), while substances are typically mass (e.g. water, rice, air, etc.). However, several authors have noted that the grammatical count-mass distinction only partly corresponds to the ontological distinction between objects and substances (see e.g. Ware 1975; Link 1983; Gordon 1985; Gillon 1992; Chierchia 1998; Pelletier and Schubert [1989] 2001). In English, for instance, count nouns also refer to abstract entities (e.g. promise, emotion) and events (e.g. competition, journey). Mass nouns also refer to abstract entities (e.g. love, beauty) and groups of unindividuated objects (e.g. furniture, jewellery). A widely held view, which better captures the range of possible denotations of count and mass expressions, is that count-mass syntax corresponds to a distinction between the kinds of entities we perceive as individuals, and those we perceive as unindividuated entities (e.g. Bloom 1994b) (where the notion of individual corresponds approximately to 'discrete bounded entity').<sup>13</sup>

Already Quine (1960) addressed the problem of acquisition in connection with count-mass syntax and its link to non-linguistic ontological categories. He saw the ontological distinction between objects and substances as a cultural construction, and count-mass syntax as the means by which children learn to discriminate objects from substances. A logical consequence of this view is that there would be substantial differences in conceptual representation between children and adult speakers, and between speakers of English (and other languages with count-mass syntax) and speakers of languages that do not have count-mass syntax, as in classifier languages such as Chinese and Japanese (i.e. languages that contain morphemes whose function is to indicate the formal or semantic class to which items belong). Against such linguistic determinism, psychologists have argued that the acquisition of count-mass syntax presupposes rather than gives rise to such basic ontological categories (Bloom 1994a, 1999; Papafragou 2005; Barner and Snedeker 2005, 2006), which may indeed have an innate basis.

There has been much less focus on the contribution of pragmatics to the interpretation of an expression as having a count or mass denotation. Pelletier and Schubert ([1989] 2001: 274), following Ware (1975), reject the possibility of developing a pragmatic account of the count-mass distinction, on the grounds that "most times there is nothing that the speaker has in mind which would allow us to classify the use of a noun phrase as either +count or +mass". While I doubt that this is in fact correct, Ware's argument is worth considering.

Ware's point was the following. We cannot draw the count-mass distinction on the basis of speaker intentions, because (i) a speaker may have a determinate intention to refer to either an individual or substance of some kind, but the hearer need not recognise this intention for communication to be successful (e.g. the hearer may respond to a request for *the* 

<sup>&</sup>lt;sup>13</sup> Some scholars do not think that this distinction adequately captures the semantics of mass terms either, as there seem to be cases where these denote discrete countable objects (e.g. *furniture*, *jewellery*, *money*). An alternative view is to treat the conceptual entity picked out by a mass term as being 'unspecified' as regards quantification and measurement (Gillon 1992, 1999; Barner and Snedeker 2005, 2006; Li, Barner, and Huang 2008).

<sup>&</sup>lt;sup>14</sup> However, there is considerable controversy as to whether such languages in fact lack count-mass syntax (cf. Krifka (1995) and Chierchia's (1998) view that all Chinese nouns are mass) or whether this distinction also appears in classifier languages under a different guise (cf. Li, Barner and Huang's (2008) distinction between 'count-classifiers' and 'mass-classifiers').

candy without knowing whether the speaker is individuating or amassing), and (ii) the speaker herself may have an indeterminate intention, and so could use a definite singular NP with the intention to refer to something without a commitment to a count or mass denotation (e.g. if one Beethoven sonata was played, and the speaker says she liked *the Beethoven*, she does not have to be either using a count or a mass noun).

Although both Ware's scenarios indeed seem possible, they require that the contexts are such that making the distinction would not be considered relevant (i.e. it would be too costly in processing terms or not contribute to the cognitive effects of the utterance). For instance, (i) could occur in a context where a bowl of candy is manifest to both speaker and hearer, and where the speaker asks the hearer to pass her *the candy*. Since the NP *the candy* is ambiguous with respect to whether it is count (in the singular) or mass, the hearer cannot be certain if the speaker considers what is in the bowl a portion of candy (mass) or a group of individual pieces of candy, but this is irrelevant to his ability to respond appropriately to her request. Similarly, in the example involved in scenario (ii), where the speaker says she liked *the Beethoven*, the communicative context does not require of either the speaker or the hearer that the denotation of *the Beethoven* is specified as either mass ('the music by Beethoven that was just played') or as count ('the sonata by Beethoven that was just played'), since there would be no doubt on the part of the hearer what the speaker is referring to, regardless of whether he construes the referent in the one or the other way (or indeed, leaves the issue indeterminate). <sup>15</sup>

In addition, Ware (1975) describes a number of cases in which a count or a mass occurrence of a noun does not make a noteworthy difference to its interpretation (this concerns in particular abstract terms, e.g. 'needing much more theory/many more theories', 'finding more truth/truths', etc.), as well as cases where there seems to be no obvious real world fact to explain the tendency to use a noun with a count or mass determiner (e.g. we say 'much whiskey/gin and orange', but 'many martinis/orange blossoms'). Ware's overall conclusion (endorsed by Pelletier and Schubert), appears to be that since there are cases in which the count-mass distinction is not reflected in the speaker-intended meaning of an utterance, pragmatics is not useful in making the distinction.

Of course, if the aim is to classify a noun as being (lexically specified as) either [+count] or [+mass], which is, as I understand it, Pelletier and Schubert's aim, then I agree that pragmatics does not take us very far in the above examples. However, this aim appears to rest on the assumption that it is meaningful to speak of a count-mass distinction in the language existing independently of the real-world distinction between countable (individuated) entities and substances (unindividuated entities). I am not convinced that it is. In my view, rather than showing the inadequacy of a pragmatic explanation, Ware's examples indicate that pragmatics may in fact have quite an important role to play in distinguishing between count and mass uses of words (or, as in the above cases, in not making a distinction at all). Recall the analysis of the definite singular NP the rabbit in (27) ('She enjoyed but later regretted the rabbit'), which involved the activation of highly accessible encyclopaedic knowledge about rabbits (e.g. that rabbits are delicious to eat), which, together with assumptions activated by other concepts encoded by the sentence, gave rise to the interpretation that the rabbit was eaten, involving a mass use of rabbit. Note that

<sup>&</sup>lt;sup>15</sup> The situation where a distinction is not grasped but where the communication is still considered good enough for practical purposes is in fact quite common. Another sort of example is the use of an NP such as *The children* where 'children' can be thought of in distinction from adults or as 'children of so and so' (even when they are adult) and there can be instances where we would not need to know which of these ways the speaker was construing some particular individual children when she says 'The children will be pleased', that is, it won't matter with regard to the hearer's ability to recognise which individuals she is referring to.

the linguistic meaning of (27) is not at odds with a count interpretation of *rabbit* (e.g., that she enjoyed *having* the (live) rabbit but later regretted it, for instance, in a context where the referent has decided to buy a rabbit for her children - at first she enjoys having it (it's cute, cuddly, etc.), but later she regrets it due to the amount of work involved in caring for it). In the absence of any syntactic clues, a crucial aspect of distinguishing between count and mass uses of nouns involves applying encyclopaedic knowledge about the entities they denote (and contextual information more generally) in their interpretation.

As I briefly touched on above, many syntactic approaches take the count-mass distinction to arise from common nouns being marked as either [+count] or [-count] in the lexicon. Underlying this assumption is the idea that some nouns are count nouns, while others are mass nouns; a noun either falls into the one or the other category and the basis for this view is rarely questioned. <sup>16</sup> However, several authors have noted that, in fact, most nouns (if not all) can be both count nouns and mass nouns (Pelletier 1975; Ware 1975; Krifka 1995; Gillon 1999; Pelletier and Schubert [1989] 2001). For instance, the mass noun water can be used as a count noun, e.g. in the context of a restaurant ('Two waters, please'), abstract terms like love and hope can be used both as mass nouns and as count nouns ('She needs some love'/'She's the love of my life', 'There is not much hope left'/'My main/primary hope is that I will finish my thesis'), count nouns like sun, shopping centre and chicken can be used as mass nouns ('We had a lot of sun at the beach', 'That's a lot of shopping centre for a small town', 'Chicken is my favourite kind of food'). As Ware (1975: 384) puts it, "there are words with a double life - sometimes occurring as count nouns and sometimes occurring as mass nouns". This raises several issues: Given the ability of nouns to occur as both count and mass, does it make sense to operate with a syntactic distinction between count and mass nouns? Shouldn't the distinction rather be treated as pertaining to occurrences of entire noun phrases, where these are interpreted as being semantically count or mass, depending on whether they denote individuals or unindividuated entities (what Pelletier and Schubert [1989] 2001 term the 'semantic occurrence approach') (and in many cases, unspecified with regard to this distinction)? If so, where do our intuitions about single nouns being either count or mass (e.g. *horse* is count, *water* is mass) come from?

In the next section, I will outline a pragmatic, relevance-theoretic account of the class of systematic polysemy that rests on this ability of nouns to take on both count and mass interpretations. The account will be based on something like a 'semantic occurrence approach' to the count-mass distinction, as briefly described above. This allows for entire noun phrases to be encoded as having either count or mass denotations (or as being unspecified with respect to the distinction, as in (27) above), thus capturing what is clearly a morpho-syntactic dimension to the distinction, while not having to assume that nouns in themselves are syntactically either count or mass. <sup>17</sup> The aim will be to show that pragmatics,

<sup>&</sup>lt;sup>16</sup> In light of this, it is easy to see the motivation behind rule-based accounts of the systematic polysemy that arises from the alternation between count and mass senses of a noun. If a noun is syntactically marked as either [+count] or [-count], then a lexical rule may operate to change the value of the feature, thus rules like 'universal grinding' and 'portioning' (a rule that converts food or drink denoting mass nouns into a portion of that substance, e.g. *three beers*) (Copestake and Briscoe 1996), or 'conversion' of count nouns into mass nouns and vice versa (Gillon 1999).

<sup>&</sup>lt;sup>17</sup> Pelletier and Schubert ([1989] 2001) is a survey paper in which they identify and discuss a range of problems associated with the count-mass distinction, and the ways in which these are manifested in the various syntactic and semantic approaches to it. As they see it, there are only two viable alternatives for an account of the count-mass distinction: (i) a 'syntactic expression approach', where COUNT and MASS are syntactic features which attach to nouns, and where 'lexical extension' rules may operate to change the count/mass feature and alter the semantic representation of the noun accordingly, and (ii) a 'semantic occurrence approach', where occurrences of noun phrases are interpreted as being (semantically) count or mass (depending on whether the entity it refers

often aided by syntactic clues, has a very constructive role to play in the interpretation of systematic polysemy involving alternations between count and mass interpretations of a single word. I will also suggest a possible answer to the question of where our intuitions about single nouns being either count or mass come from.

# 3.2 A (mainly) pragmatic account of systematic polysemy

First some preliminaries. In relevance theory (Sperber and Wilson 1986/1995, Carston 2002), which is fundamentally a cognitive account of utterance interpretation, the distinction between linguistic semantics and pragmatics is seen as corresponding to different processes involved in utterance comprehension: linguistic decoding and pragmatic inference (Sperber and Wilson 1986/1995; Wilson and Sperber 1993; Carston 2002). Assuming a modular language input system in the sense of Fodor (1983), the process of decoding linguistic utterances is performed by a language perception module (or parser), which takes the linguistic stimulus as input, executes a series of grammatical computations (or mappings) and delivers a semantic representation as output (called 'logical form' in relevance theory), <sup>18</sup> which then feeds the pragmatic inferential system. <sup>19</sup> A central characteristic of Fodor's input systems (in addition to, inter alia, their automaticity and informational encapsulation) is that they deliver shallow outputs. The structured set of concepts that constitute the logical form of an utterance is rarely (if ever) fully propositional. Carston describes the logical form of an utterance as a 'template' or 'schema' for a range of possible propositions, which contain slots that have to be filled – a process that requires pragmatic inference (see Carston 2002, Chapter 1, section 1.5.2 for more detail). A consequence of the assumption that linguistic decoding and pragmatic inference combine in verbal comprehension is that there is no need to assume that all words necessarily encode full concepts, or that every concept communicated by an utterance *could* in principle have been encoded (Wilson 1998; Carston 2002).

A mentally represented concept, a constituent of the 'language of thought' (Fodor 1975, 2008), is seen as an address (or entry) in memory that may give access to three types of

to is true of stuff or things, which is determined by the syntactic constructions) (ibid. 314). The first, syntactic account has the advantage of accounting for our intuitions about many nouns being 'count nouns' and many nouns being 'mass nouns'. However, Pelletier and Schubert note that against this syntactic account it may be held that it is unclear what syntactic work the features actually do, when, "for every sentence which has a mass term in a given location, there is another sentence which has a count term in that position" (and vice versa); no constructions will be ruled out (ibid. 322). The only role of the count/mass features seems to be to direct the semantic translation process, which points to a semantic distinction rather than a syntactic one. The semantic occurrence account, on the other hand, in which Pelletier and Schubert take it that all nouns have a 'comprehensive extension', that is, being "true of kinds (of stuff or things), true of conventional kinds of servings, true of quantities (of stuff) and true of objects coinciding with quantities of stuff' (ibid. 314), and that their precise denotations are determined by the syntactic constructions they enter into (e.g. the determiner a allows semantic operators to convert this 'basic' meaning of a noun into one that is true of individuals and kinds of individuals), our intuitions about the 'count' or 'mass' nature of many nouns is not accounted for. Although I think this account is certainly on the right track, I take it a step further in what follows, where, in the pragmatic account that I propose, much of the work of deciding on the precise denotation of occurrences of nouns is left to pragmatic inference. This allows us to assume that instead of having 'comprehensive extension', many nouns encode concepts that are mentally represented as denoting (kinds of) individuals or (kinds of) unindividuated entities (the count-mass distinction is thus primarily a cognitive-conceptual one), and that pragmatics, in combination with syntactic clues, operate to yield occasion-specific senses of nouns.

<sup>&</sup>lt;sup>18</sup> This is, of course, a kind of abstraction since in actual on-line processing the decoded words/morphemes are delivered rapidly to the pragmatic processing system (which does not 'wait' to get the logical form as a whole).

<sup>&</sup>lt;sup>19</sup> In relevance theory, the pragmatic inferential system is also seen as being modular, perhaps constituting a sub-module of the more general theory of mind module (Sperber and Wilson 2002; Wilson 2005), though in a somewhat weaker sense of 'modular' than Fodor's.

information: (i) the *logical* properties of the concept (e.g. a one-way meaning postulate such as HORSE → KIND OF ANIMAL); (ii) a set of assumptions, or *encyclopaedic* information, about the denotation of the concept, that is, conceptually represented assumptions and beliefs, including stereotypes and culture-specific information about the denotation, and also, in many cases, imagistic and/or sensory-perceptual representations (e.g. HORSES HAVE MANES, HORSES ARE EDIBLE, HORSES HAVE FOUR LEGS, HORSES ARE USED FOR RIDING, HORSES LOOK LIKE THIS: [MENTAL IMAGE], etc.), and (iii) the *lexical* (i.e. phonological and syntactic) information connected with the linguistic form that encodes the concept (Sperber and Wilson 1986/1995: 85-93).

Relevance theory sees lexical interpretation as typically involving the construction of *ad hoc* concepts (i.e. occasion-specific senses), which may be narrower or broader than the linguistically encoded senses (Carston 2002, Wilson and Carston 2007). An example of narrowing is the use of *happy* to denote a particular state of happiness, for instance 'happy in a low-key, peaceful sort of a way' (represented by the *ad hoc* concept HAPPY\*<sup>20</sup>), while an example of broadening is the use of *boiling* to refer to a very high temperature (represented by the *ad hoc* concept BOILING\*). In the construction of such *ad hoc* concepts, the hearer takes the encoded concept and its associated logical and encyclopaedic information as input, together with a set of contextual assumptions, using the relevance theoretic comprehension heuristic (according to which the hearer (i) takes the decoded linguistic meaning, follows a path of least effort in considering interpretive hypotheses, and (ii) stops when the interpretation he arrives at satisfies his expectations of relevance, cf. Wilson and Sperber 2004: 613) to derive warranted conclusions about the speaker's meaning.<sup>21</sup>

Let us assume that our ability to discriminate between kinds of individuals and unindividuated entities is independent of count-mass syntax. This seems plausible, as there is little reason to think that speakers of e.g. Chinese or Japanese represent the world differently with respect to this distinction. There is also substantial evidence that early individuation capacities predate the acquisition of count-mass syntax in development (e.g. Papafragou 2005; Barner and Snedeker 2005). If we assume that the ability to distinguish between individuals and unindividuated entities is one of the means by which we carve up and make sense of the world around us, it seems reasonable that many of our concepts will be mentally represented as either individual-denoting or unindividuated entity-denoting, in accordance with how we perceive (or categorise) the things they refer to in the world (independently of their natural-language counterparts). Given the short description of a mentally-represented

<sup>&</sup>lt;sup>20</sup> By convention *ad hoc* concepts are marked by an asterisk.

<sup>&</sup>lt;sup>21</sup> Sperber and Wilson (1986/1995) define relevance is as a potential property of all types of input to cognitive processes, which may be assessed in terms of the amount of effort (of perception, memory and inference) it takes to process the input, and the positive cognitive effects the individual may derive from it (where a 'positive cognitive effect' is described as a 'worthwhile difference to the individual's representation of the world', and may include strengthening or elimination of existing assumptions, derivation of contextual implications). Other things being equal, the more cognitive effects an input yields to an individual and the less effort it takes to process it, the higher the degree of relevance of that input to that individual at that time.

<sup>&</sup>lt;sup>22</sup> In psychology, a dominant theory of the nature of human object knowledge is the *core knowledge hypothesis*, proposed by Spelke and her colleagues (Spelke 2000; Spelke et al. 1994). This hypothesis states that our knowledge of the physical world stems from innate principles for reasoning about inanimate physical objects. This knowledge shares many of the properties of a Fodorian module by being experience-independent, domain-specific and informationally encapsulated. A possible hypothesis, then, is that our ability recognise countable objects versus undifferentiated masses (like water, snow, etc) is part of this 'core knowledge'. I return to this issue shortly in the discussion of 'meaning postulates' as a possible source of a conceptual count-mass distinction.

concept I gave above, we can see the concept HORSE as being represented in the following way:

(28) HORSE

Lexical entry: +N (... etc.)

Logical entry: HORSE → ANIMAL OF A CERTAIN KIND

Encyclopaedic entry: IS USED FOR RIDING, IS OFTEN DOMESTICATED, IS POPULAR

AMONG YOUNG GIRLS, LOOKS LIKE THIS [MENTAL IMAGE], ... etc.

The concept HORSE in (28) is associated with three types of information stored in its lexical, logical and encyclopaedic entries. The lexical entry contains information about the natural language counterpart of the concept (e.g. that the lexical form *horse* belongs to the category N). The logical entry associated with the concept consists of a meaning postulate (in the form of an inference rule), which applies reliably to any representation that meets its input conditions (i.e. the concept HORSE). The encyclopaedic entry contains a set of assumptions that represent the individual's real world knowledge about horses (e.g. that horses are used for riding, are often domesticated, are popular among young girls, and so on). Similarly, the concept WATER may be represented (schematically) along the lines of (29), as associated with the tree types of information:

(29) WATER

Lexical entry: +N (... etc.)

Logical entry: WATER  $\rightarrow$  LIQUID OF A CERTAIN KIND

Encyclopaedic entry: IS USED FOR DRINKING, IS NECESSARY FOR LIFE ON THE PLANET, FREEZES AT ZERO DEGREES CELSIUS, LOOKS LIKE THIS [MENTAL IMAGE] ... etc.

On the basis of the representations in (28) and (29), there are two different ways in which we could see the concepts as being specified as denoting individuals or unindividuated entities. Either it could be seen as (i) following from the meaning postulate(s) associated with the concept, or (ii) being derivable form real-world knowledge of the denotation of the concept, stored in the encyclopaedic entry associated with it.<sup>24</sup> I will consider each possibility in turn.

On the view expressed in (i), the logical information (or meaning postulates) associated with a concept is hypothesised to be what allows us to determine that it denotes individuals or unindividuated entities. Explaining how this might work requires some elaboration. Relevance-theory has adopted the Fodorian, informational atomist stance on lexical meaning. However, unlike Fodor, who, in recent work, denies that any of the inferences a concept enters into is constitutive of its content (e.g. Fodor 1998), relevance theorists take concepts to be associated with content-constitutive meaning postulates (see, in particular Horsey 2006 for a detailed defence of the need for supplementing Fodor's informational semantics with content-constitutive rules of inference). That a concept is associated with a meaning

<sup>23</sup> Of course, a concept can be associated with more than a single meaning postulate.

<sup>&</sup>lt;sup>24</sup> Having argued that concepts, and not their natural language counterparts, may be mentally represented as denoting individuals or unindividuated entities, the lexical entries associated with the concepts would not contain any information that would allow us make such a distinction.

<sup>&</sup>lt;sup>25</sup> Horsey argues that the distinction between logical and encyclopaedic information associated with a concept should be seen as a psychological one, which does not coincide with the analytic/synthetic distinction in philosophy: "[T]he question of whether a given inference is supported by a meaning postulate or is represented

postulate means that a tokening of the concept, e.g. HORSE, will cause the activation of the associated meaning postulate(s), 'HORSE  $\rightarrow$  ANIMAL OF A CERTAIN KIND', thereby tokening the concept ANIMAL OF A CERTAIN KIND. Assuming that the concept ANIMAL is associated with the meaning postulate 'ANIMAL  $\rightarrow$  LIVING KIND', the concept LIVING KIND would equally be tokened as a result of a tokening of the concept HORSE). This is one of a large array of concepts that Horsey (2006: section 4.3.2) describes as 'intuitive', that is, concepts for which we have an intuitive grasp (e.g. HORSE, ANIMAL, WATER, LIQUID, etc.). <sup>26</sup> He distinguishes between two types of intuitive concepts: those that are 'perceptual' and those that are 'inferential'. Assuming a modular view of mental architecture (cf. Fodor 1983; Sperber 1994), perceptual concepts are the outputs of sensory stimuli modules (broadly identified with so-called 'basic level categories', cf. Rosch et al. 1976), while inferential concepts are taken to be innate concepts which are introduced by inference rules (meaning postulates), and which are not part of the vocabulary of the perceptual mechanisms. As an example, Horsey gives the concept LIVING KIND, which can be introduced via the spontaneous inference 'ANIMAL  $\rightarrow$  LIVING KIND', and which Sperber (1994), on the basis of strong commonalities in living-kind classification across cultures, suggests treating as a domain-specific cognitive module. Much research in cognitive science indicates that there are several different cognitive domains that licence spontaneous inferences about entities relevant to them, and that the domains have many of the properties of a Fodorian module. It is argued that such domains plausibly include a naïve physics (Spelke 2000), a naïve biology (Atran 2002) and naïve psychology (Leslie 1987), and that they are part of the cognitive apparatus which makes knowledge acquisition possible (Sperber 1994) (cf. the core knowledge hypothesis, mentioned in footnote 22). Along these lines, we may speculate that the concepts INDIVIDUAL and UNINDIVIDUATED ENTITY may be such inferential concepts, perhaps constituting submodules of the more general domain of naïve physics. For instance, the concept WATER, which, according to the meaning postulate given in (29) causes tokening of the concept LIQUID OF A CERTAIN KIND, may activate the concept UNINDIVIDUATED ENTITY via the spontaneous inference 'LIQUID OF A CERTAIN KIND  $\rightarrow$  UNINDIVIDUATED ENTITY'. In the case of the concept HORSE in (28) which is argued to activate the inferential concept (or conceptual module) LIVING KIND as a result of the tokening of the concept ANIMAL, there may be a further spontaneous inference activated by the concept ANIMAL, 'ANIMAL  $\rightarrow$ INDIVIDUAL'. In this way, the distinction between concepts that denote individuals and those

in encyclopaedic knowledge is an empirical question which is to be determined by psychological investigation." (ibid. 75).

The notion of 'intuitive concepts' comes from Sperber (1997), who distinguishes between 'intuitive' and 'reflective' beliefs. Intuitive beliefs are beliefs derived from perception (including introspection), or are derived directly or indirectly from these on the basis of intuitive inference mechanisms, e.g. in the form of meaning postulates attached to concepts (e.g. HORSES ARE ANIMALS/WATER IS A LIQUID). They can also be acquired via communication (someone might tell you that e.g. OKAPIS ARE ANIMALS), which, in such cases, can be seen as playing the role of experience by proxy (the individual might herself have formed this belief via perception or spontaneous inference had she been placed in a position to experience its perceptual basis). Reflective beliefs are those that are entertained by virtue of being embedded in a *validating context* (which may involve reference to authority, to divine revelation, explicit argument or proof, etc.), and may be derived, *inter alia*, from communication (resulting in, e.g., a belief of the form: MARY [WHO I FIND RELIABLE] SAID THAT *P*) or conscious thought (e.g. THERE IS EVIDENCE IN SUPPORT OF *P*). Parallel to this is a distinction between 'intuitive concepts' (i.e. those which may appear in intuitive beliefs and inferences, e.g. concepts such as HORSE, ANIMAL, WATER, LIQUID) and 'reflective concepts' (i.e. those of which we have no intuitive grasp, but understand because we have (reflective) beliefs about them, e.g. the concept H<sub>2</sub>O, and which "are introduced by explicit theories which specify their meaning and the inferences that can be drawn on their basis").

that denote unindividuated entities can be seen as arising from meaning postulates attached to the concepts.

A second option is to take the distinction between individuals and unindividuated entities to arise from encyclopaedic information stored about the denotations of concepts. As I have said above, the encyclopaedic entry attached to a concept contains information about the objects, events and/or properties which instantiate it. Typically, this entry contains information about prototypical instances of the concept. For instance, the encyclopaedic information about the denotation of the concept HORSE may contain information about what constitutes a typical horse, e.g. the kind of horse commonly used for riding (or the sport/show horse), and information about less typical instances, such as wild horses and small pony-sized horses (e.g. the Icelandic horse). For many concepts, we may also have an imagistic representation of the typical instances of the entities/properties they denote as part of the encyclopaedic entries attached to them (for instance, Jackendoff 2002 suggests that physical object words may be associated with 3D model representations in). It may be argued, then, that the fact that we perceive the concepts HORSE and WATER in (28) and (29) as denoting individuals and unindividuated entities respectively may stem from imagistic representations of typical instances of horses and water stored in the encyclopaedic entries attached to the concepts. We 'see' horses out in the world as bounded countable individuals, and water as an unbounded entity. On this view (which I tend to favour over (i), due to its intuitive plausibility), the count-mass distinction has a perceptual basis.

By arguing that the conceptual distinction between individuals and unindividuated entities can be seen as arising from either (i) the logical information attached to a concept, or (ii) an imagistic representation stored in the encyclopaedic entry of a concept, I hope to have shown that we can account for the intuition that some nouns *are* count nouns, while others *are* mass nouns, without having to assume that this is a linguistic property of nouns (thus avoiding the problems associated with a syntactic account) but rather to do with the kinds of things that they denote in the world. In my view, this offers a more psychologically plausible account: it explains that children and Chinese speakers may perceive horses as bounded countable entities and water as an unindividuated entity despite their 'lack of' count-mass syntax, and allows for many concepts, for which the classification in terms of individuals and unindividuated entities seems less relevant (as is the case for, e.g., abstract terms for emotions such as *love*, *hope*, *grief*, *anger*, and so on, plurals such as *blues*, *wares*, *clothes*, etc.), to be unspecified with regard to the distinction.

However, this approach does not deny (the obvious fact) that, in languages that have it, count-mass syntax plays a central role in determining the denotations of nouns. But rather than taking count-mass syntax as being (somehow) prior to the conceptual distinction between individuals and unindividuated entities, count-mass syntax can be seen as an instance of a conceptual (ontological) distinction being reflected in a syntactic distinction at the level of NPs. Let us assume, then, that count-mass syntax (as it is manifested at the level of NPs) maps onto semantic representation as schematised in (30):<sup>27</sup>

- (30) a. Count syntax  $\rightarrow$  (KINDS OF) INDIVIDUALS
  - b. Mass syntax  $\rightarrow$  (KINDS OF) UNINDIVIDUATED ENTITIES

\_

<sup>&</sup>lt;sup>27</sup> These mappings are thought to model the linguistic knowledge of adult language users. I remain neutral as regards the very contentious issue of how such mappings are acquired (for discussion, see Bloom 1994a, 1994b, 1999; Papafragou 2005; Barner and Snedeker 2005, 2006).

These mappings allow a hearer to infer, on the basis of syntactic clues, that the entity denoted by the noun is either an individual or unindividuated (a substance). <sup>28</sup> But instead of taking the nouns themselves to be either count or mass (which, on the basis of that syntactic feature, select specific determiners, can appear in the plural form, etc.) the view is now that it is concepts that are perceived as denoting (kinds of) individuals or unindividuated entities, and, independently of this, the language can encode NPs as denoting (kinds of) individuals or unindividuated entities.<sup>29</sup> The normal case is, of course, an overlap between these two situations (with count-mass syntax allowing for the modification of nominal denotations by specifying cardinality, quantity, portion, amount, etc.), and so the concepts encoded by nouns such as horse, man, pencil, being mentally represented as denoting individuals, are most frequently used with count syntax, while the concepts encoded by nouns such as water, rice, and *snow*, being mentally represented as denoting unindividuated entities, are most frequently used with mass syntax. However, whenever a noun that encodes a concept which is perceived as denoting a (kind of) individual is used with mass syntax, or a noun that encodes a concept which is perceived as denoting a (kind of) unindividuated entity is used with count syntax, this gives rise to polysemy. Consider (31) and (32):

- (31) a. The horses are grazing in the field.
  - b. Yesterday, I had *horse* for the first time.
- (32) a. The plants need *some water*.
  - b. [To waiter]: I'll have the hamburger and *a water*, please.

In (31)a, the noun *horse* is used in a count NP to denote individuals, and (32)a, the noun *water* is used in a mass NP to denote an unindividuated entity. There is thus a correspondence between the conceptual representations HORSE and WATER (being perceived as denoting individuals and unindividuated entities respectively) and the syntax of the NPs in which their natural language counterparts occur. In (31)a, where *horse* is used in an NP with mass syntax, the language parser outputs a concept, let's say HORSE {unindividuated entity}, which comes with an instruction to the pragmatic system that its denotation is constrained to unindividuated entities. This concept provides a highly underspecified input to pragmatic processing. The relevance-theoretic comprehension heuristic then operates to yield the speaker-intended concept (which is one among the various sub-tasks it performs in developing the logical form into a fully propositional string of concepts) resulting in the construction of the *ad hoc* concept HORSE\*, paraphrasable as 'horse meat', which is narrower than the linguistically specified concept. The hearer constructs this concept on the basis of highly activated encyclopaedic information associated with the concept HORSE (e.g. HORSES

<sup>&</sup>lt;sup>28</sup> For instance, if someone hears an utterance of 'There's much *glurf* these days' without having ever come across the noun *glurf* before, he or she can infer that the denotation of the NP *much glurf* is an unindividuated entity.

<sup>&</sup>lt;sup>29</sup> Krifka (1995) also considers the possibility of not distinguishing syntactically between count and mass nouns, but concludes that it must be done in order to account for what appears to him to be a syntactic difference, that mass-denoting nouns can be straightforwardly used as names of kinds (e.g. 'Wine contains alcohol'), while count-denoting nouns need a definite article or pluralization to perform this task (e.g. 'The bear/bears can be aggressive'). However, another way to see this is that count-denoting nouns can in fact be used without a determiner to name a kind, only that this would force a mass interpretation upon them (for instance, an utterance of 'Bear is delicious' would be possible given the appropriate context). So it seems that it is not whether the noun is syntactically defined as count or mass but rather how it occurs in the sentence that is the determining factor here. Occurring without a determiner and plural marker yields a mass interpretation.

ARE EDIBLE) and contextual information derived from the utterance situation (e.g. that the speaker went to a fancy restaurant last night). Similarly, in (32)b, where *water* is used with count syntax, the language parser outputs a semantic representation that is specified as an individual, let's say WATER{individual}. The pragmatic processor then takes over to construct an *ad hoc* concept on the basis of the encoded concept, highly activated encyclopaedic assumptions associated with it (e.g. WATER IS DRUNK FROM CONTAINERS), and other contextual assumptions derivable from the utterance situation, say, from the fact that the speaker and hearer are at a restaurant where people are served glasses of water. The communicated concept, WATER\*, paraphrasable as 'glass of water', is narrower than the linguistically specified concept, denoting a subset of the set of individual containers of water.

We are now (finally) in a position to account for the examples of systematic polysemy that we set out in the beginning of this paper in (2), (3) and (4), repeated below as (33), (34) and (35):

- (33) a. A *rabbit* jumped over the fence.
  - b. We're having *rabbit* for dinner.
  - c. The model wore *rabbit* on the catwalk.
  - d. After the tractor had run over the body, there was *rabbit* splattered all over the yard.
- (34) a. We have a *pine* in our garden.
  - b. This table is made of *pine*.
- (35) a. Susan decorated the cake with a *cherry*.
  - b. When the kids left, there was *cherry* all over the kitchen.
  - c. Jill and Joan have a *cherry* in their garden.<sup>30</sup>
  - d. This table is made of *cherry*.

In (33), the unindividuated senses of *rabbit* are indicated by the mass syntax of the NPs, but their narrower, contextually appropriate interpretations, i.e. 'rabbit meat' (33)b, 'rabbit fur' (33)c, and 'rabbit stuff'(33)d, are all derived by pragmatic inference, on the basis of encyclopaedic assumptions associated with the concept RABBIT (e.g. RABBITS ARE EDIBLE, RABBITS HAVE FUR, RABBITS ARE OF FLESH AND BLOOD, etc.). The wood sense of *pine* in (34)b and of *cherry* in (35)d arise as a result of the nouns occurring in NPs with mass syntax, encoding unidividuated pine and cherry concepts, combined with pragmatic narrowing. Importantly, this relevance-theoretic pragmatic account claims that the reason we can use *rabbit* to denote the senses in (33)b, (33)c, (33)d, and *pine* and *cherry* to denote kinds of wood in (34)b and (35)d, is first and foremost a result of what we know about the world (or what we perceive the world to be like), and not due to any such information being stored in the lexical entries for the words, although mass syntax provides an important clue to the interpretations. Examples like (27) above, however, where the syntax of the NP is unspecified as to whether it encodes a count or a mass concept, show that the choice of a count or mass interpretation for a concept can have an entirely pragmatic basis.<sup>31</sup>

However, although they are pragmatically derived, the fact that we perceive such polysemies as being 'systematic' or 'regular' is, I think, due to several (independent) factors: One concerns the syntactic distinction between count-denoting and mass-denoting NPs. The

<sup>&</sup>lt;sup>30</sup> The metonymic extension from the meaning of *cherry* in (35) above to the tree sense of *cherry* in (35) is not accounted for here, but see Falkum (2010) for a suggestion. The alternation between the tree sense and the wood sense, however, follows the pattern described above.

<sup>&</sup>lt;sup>31</sup> Presumably, this would be crucial in classifier languages such as Chinese and Japanese, which do not distinguish syntactically between count-denoting and mass-denoting NPs.

fact that our linguistic knowledge tells us that most (if not all) nouns can be used with both count and mass syntax, and that this syntactic distinction is paired reliably with a semantic distinction between individuals and unindividuated entities, gives rise to a sense of regularity. A second factor is the way in which this type of polysemy reflects certain regularities in the world. Fodor and Lepore make this point:

Suppose it's right that 'lamb' is polysemous between the animal and the meat. Surely that's because lamb-the-meat comes from lamb-the animal. Surely there just *couldn't be* a word that's polysemous between *lamb-the-animal* and (say) *beef-the-meat*? Or between *lamb-the-animal* and *succotash-the-mixed vegetable*? That there couldn't may itself sound like a deep fact of lexical semantics. But no; it's just the truism that, the less one can see what the relation between X and Y might be, the more one is likely to think of an expression that is X/Y ambiguous as homonymous rather than polysemous. (Fodor and Lepore 2002:117)

Polysemous relations of the animal-meat, the animal-fur or the tree-wood kind indeed reflect highly regular and predictable states of affairs in the world. Because we know that there is an inherent relation between an animal and its meat/fur, and between a tree and its wood, we can easily infer, upon encountering a new kind of animal or tree, that the relation also applies to this instance (and in certain cases that it does not, for instance, if the animal in question has no fur). Since words for animals and trees are usually polysemous with respect to these relations, we can infer that the word for the new instance can also be used in this way. The view, then, is that the polysemy that the language exhibits in these cases is dependent on, and arises from our knowledge of predictable real world relations, rather than being generated by the linguistic system on the basis of stored lexical information.

A third factor that contributes to the perceived 'regularity' of this type of polysemy is, of course, that some of its instances have become conventional (e.g. *lamb*, *chicken*, *oak*, etc.), as a result of frequency of use and possibly other factors contributing to the stabilisation of senses. In such cases, the conventionalised sense may have acquired a conceptual address of its own, and would be associated with its own logical and encyclopaedic entries.

An advantage of the pragmatic approach is that it allows for considerable flexibility of interpretation. Consider the examples in (36):

- (36) a. A *chicken* pecked the ground in the backyard.
  - b. John brought home a *chicken* for supper.
  - c. Corn-fed *chicken* is difficult to find these days.
  - d. Corn-fed and inexpensive *chicken* is difficult to find these days.

In (36)a, *chicken* is used in the stereotypical way to denote the live animal. In (36)b however, it is used in an NP with count syntax to refer to an individual chicken prepared to be eaten. This 'meat' sense of chicken represents a pragmatic narrowing of the linguistically-encoded meaning, by picking out a subset from among the set of chickens (i.e. the set of whole chickens prepared to be eaten). In fact, depending on the level of specificity required by the context, it could be narrowed down further to communicate, e.g., 'a dead unplucked chicken', 'a frozen plucked chicken' (like the ones you find in any Sainsbury freezer cabinet), 'a cooked chicken ready to be eaten' (as from a delicatessen), and so on. It is difficult to see how this ability of *chicken* (used in its individual-denoting sense) to take on both animal and meat senses could be analysed in an account based on lexical rules for sense extension. It is possible that the proponents of such an account would say that deriving the appropriate sense of *chicken* in (36) above is a pragmatic matter, and that they never intended their rules to

account for such flexibility in meaning. If so, this would mean that only the rule-generated meat sense (e.g. the one that results from 'meat-grinding' on Copestake and Briscoe's (1996) account) could be considered as 'semantic', while the meat sense of *chicken* in (36)b would come about as a result of pragmatic inference. This asymmetry does not arise in the pragmatic account, which I think points towards a more empirically adequate account. Concerning the mass uses of *chicken* in (36)c and (36)d, they show that a property primarily associated with the live animal ('corn-fed') can be used with the meat sense of *chicken*, and that such a property can be conjoined with a property primarily associated with the meat sense ('inexpensive'). The fact that we do not find these utterances the least incoherent may be due to logical and encyclopaedic information activated during our processing of them (e.g. a tokening of the concept CHICKEN causes tokening of the concept ANIMAL, as well as the activation of encyclopaedic assumptions about chicken-the-animal, e.g. CHICKENS EAT CORN; CORN-FED CHICKENS HAVE MORE FLAVOUR THAN OTHER CHICKENS, etc.).

Finally, consider the uses of *mink* in (37) below:

- (37) a. A *mink* lives in the attic of our mountain cabin.
  - b. I opened the cage and suddenly there was *mink* all over the place
  - c. *Mink* is expensive these days.
  - d. I wore my jeans and she wore a *mink*.<sup>32</sup>

Let us assume that mink conventionally encodes two senses, the animal sense, MINK1, and the fur sense, MINK<sub>2</sub> (where the MINK<sub>2</sub> sense has become conventional as a result of the frequent narrowing of the mass use of the animal sense ('mink stuff') into the 'fur' sense). Thus, the MINK<sub>1</sub> concept would be perceived as denoting a (kind of) individual, while the MINK<sub>2</sub> concept would perceived as denoting a (kind of) unindividuated entity. In (37), mink is used in the conventional way (in an NP with count syntax) to denote the live animal (MINK<sub>1</sub>), with a possibly narrower communicated meaning, paraphrasable as 'wild mink' (MINK<sub>1</sub>\*). In (37), mink is used in the same sense, only with mass syntax to refer to an unindividuated entity of live minks, communicating (possibly) the ad hoc concept MINK<sub>1</sub>\*\* ('farmed mink'), which is pragmatically constructed on the basis of the semantic representation MINK1 {unindividuated entity} output by the linguistic system. In (37)c and (37)d, the noun *mink* is used to denote the fur sense, MINK<sub>2</sub>, but in different ways: According to our analysis of the fur sense as being established, in (37)c, mink is used in the conventional way (in an NP with mass syntax) to denote mink fur (MINK<sub>2</sub>). However, the count syntax of the NP a mink in (37)d allows it to map onto an individual-denoting concept (MINK<sub>2</sub>{individual}), while pragmatic narrowing yields the communicated meaning, e.g. 'mink stole' (or the ad hoc concept MINK<sub>2</sub>\*). The MINK<sub>2</sub>\* concept can be seen as denoting a subset of the set of all individual mink furs. Although (37)b and (37)d are not conventional uses of the noun mink, they support the analysis that all that count-mass syntax encodes is an instruction that the concept output by the linguistic parser is specified as being either an individual or unindividuated.

Returning to the 'blocking' cases discussed in section 2, these can be analysed within the current, relevance-theoretic account as involving an increase in the effort of processing on the part of the hearer, which is offset by an increase in the cognitive effects he may derive from the utterance. According to the Communicative Principle of Relevance every utterance (or other ostensive stimulus) "creates a presumption of its own optimal relevance" (Sperber and Wilson 1986/1995: §3). To be 'optimally relevant', the utterance should be (a) at least relevant enough to be worth the hearer's processing effort, and (b) the most relevant one

<sup>&</sup>lt;sup>32</sup> Lyrics of 'She Ain't Pretty' from the album *Snow in June* (1990) by the band Northern Pikes.

compatible with the speaker's abilities and preferences. So, a speaker aiming at optimal relevance should try to do at least the following: achieve enough cognitive effects to make the utterance worth processing (cf. (a)), and avoid causing the hearer any unjustifiable effort in achieving those effects (cf. (b)). This ban on processing effort has two consequences. First, it implies that the first interpretation that the hearer finds satisfactory (for instance, if there is a highly accessible interpretation which is relevant in the expected way), this is the only one he is justified in choosing. A second consequence is that the extra processing effort demanded by any indirectness in an utterance should be offset by extra (or different) effects, which would not have been achieved by use of a more direct utterance. Consider again the examples in (16) and (17), repeated in (38) and (39) below:

- (38) There were five thousand extremely loud people on the floor eager to tear into roast *cow* with both hands and wash it down with bourbon and whiskey (Copestake and Briscoe 1996: 38).
- (39) Hindus are forbidden to eat *cow* (beef?) (Nunberg and Zaenen 1992).

In both cases, the use of *cow* with mass syntax instead of the conventional expression *beef* demands an extra (albeit relatively minor) effort of processing of the hearer, which is offset by extra cognitive effects (compared to the alternative utterances containing *beef*): In (38), it allows the hearer to draw implications about the speaker's somewhat derogatory attitude towards the people consuming the meat, and in (39) about the status of the cow in Hindu religion. As to the utterances in (13)-(15), repeated below in (40)-(42), it can be argued that the reason the utterances seem odd is that, without a more specific context, the extra processing effort that they require of the hearer is not offset by any extra (or different) cognitive effects that he may derive from them:

- (40) ?Joan likes to eat *calf* (veal).
- (41) ?We're having *pig* for dinner (pork).
- (42) ?Matt is preparing *sheep* for our anniversary (mutton).

Finally, in (43) and (44) ((18) and (19) above), the reason the uses of *pig* seem less marked than the one in (40) may be that in both cases the 'pig' sense has already been primed by the (linguistic) contexts (by the subject *Jews and Muslims* in (43) and the general subject matter in (44), 'the speaker does not eat pork'). In this way, the uses of *pig* do not necessarily require much extra effort to process, and hence the effects that might be drawn from them are quite minimal:

- (43) Jews and Muslims don't eat *pork/pig*.
- 'As a general rule, I don't eat *pork*. This can be awkward I often go for the veggie option, to avoid having to explain why I don't eat *pig*.'
- (45) Kate [with a plate of food in front of her]: This roast sheep is the best I've ever had.
- (46) I love roast suckling *pig*.

As to the utterances in (45) and (46) ((20) and (21) above), the uses of *sheep* and *pig* to refer to the meat of the animals are entirely conventional; there would thus be no extra processing effort required, and no extra effects obtained.

#### 4 Conclusion

I have argued in favour of a mainly pragmatic account of the systematic sense alternations that rest on the distinction between count and mass uses of nouns in English, giving rise to examples such as 'John shot a rabbit/had rabbit for lunch'. My claim has been that computational semantic approaches, which usually treat such sense alternations as being generated by the application of lexical inference rules (where the effect of the rules is to change the value of a linguistically marked +count or +mass feature on a noun). I argued that such rule-based accounts do not provide interpretive flexibility required to handle the variety of interpretations that this kind of sense alternation may give rise to. Instead I have proposed that our intuitions about some nouns being 'count' and other nouns being 'mass' may stem from our mental representations of the concepts that are encoded by them, which may be perceived (by virtue of meaning postulates attached to the concepts or imagistic representations in their encyclopaedic entries) as denoting individuals or unindividuated entities (or indeed, be unspecified with regard to this distinction). I further suggested that the count-mass distinction as it is manifested in language could be seen as an instance of this conceptual-ontological distinction being reflected in a syntactic distinction at the level of NPs (rather than at the level of individual nouns). The cases of systematic polysemy would arise whenever a noun that encodes a concept perceived as denoting an individual is used with mass syntax, or when a noun that encodes a concept perceived as denoting an unindividuated entity is used with mass syntax. The pragmatic process of ad hoc concept construction then operates to yield context-specific senses of these nouns (e.g. to specify whether a mass use of rabbit communicates 'rabbit meat', 'rabbit fur', 'rabbit stole', 'rabbit remains' etc.).

#### References

- Apresjan, J. D. 1974. Regular polysemy. Linguistics 14 (2):5-32.
- Aronoff, M. 1976. Word Formation in Generative Grammar, Linguistic Inquiry Monographs 1. Cambridge Massachusetts: The MIT Press.
- Atran, S. 2002. Modular and cultural factors in biological understanding: an experimental approach to the cognitive basis of science. In *The Cognitive Basis of Science*, edited by P. Carruthers, S. Stich and M. Siegal. Cambridge: Cambridge University Press. 41-72.
- Barner, D., and J. Snedeker. 2005. Quantity judgments and individuation: evidence that mass nouns count. *Cognition* 97 (1):41-66.
- ——. 2006. Children's early understanding of mass-count syntax: individuation, lexical content, and the number asymmetry hypothesis. *Language Learning and Development* 2 (3):163-194.
- Bloom, P. 1994a. Possible names: The role of syntax-semantics mappings in the acquisition of nominals. *Lingua* 92 (1-4):297-329.
- ——. 1994b. Syntax-semantics mappings as an explanation for some transitions in language development. In *Other Children, Other Languages: Issues in the Theory of Language Acquisition*, edited by Y. Levy. Hillsdale, New Jersey: Erlbaum. 41-76.
- ——. 1999. The role of semantics in solving the bootstrapping problem. In *Language, Logic and Concepts: Essays in Memory of John Macnamara*, edited by R. Jackendoff, P. Bloom and K. Wynn. Cambridge, Massachusetts: MIT Press. 285-310.
- Blutner, R. 1998. Lexical pragmatics. Journal of Semantics 15 (2):115-162.
- ———. 2002. Lexical semantics and pragmatics. In *Semantics*, edited by F. Hamm and T. E. Zimmermann. Hamburg: Helmut Buske Verlag. 27-58.
- Boguraev, B., and B. Levin. 1993. Models for lexical knowledge bases. In *Semantics and the Lexicon*, edited by J. Pustejovsky. Dordrecht: Kluwer Academic Publishers. 325-340.
- Briscoe, T., A. Copestake, and A. Lascarides. 1995. Blocking. In *Computational Lexical Semantics*, edited by P. Saint-Dizier and E. Viegas. Cambridge: Cambridge University Press.
- Carston, R. 2002. Thoughts and Utterances: The Pragmatics of Explicit Communication. Oxford: Blackwell Publishers.

Cheng, C. Y. 1973. Response to Moravcsik. In *Approaches to Natural Language*, edited by J. Hintikka, J. Moravcsik and P. Suppes. 286-288.

- Chierchia, G. 1998. Reference to kinds across language. Natural language semantics 6 (4):339-405.
- Clark, E. V., and H. H. Clark. 1979. When nouns surface as verbs. Language 55 (4):767-811.
- Copestake, A., and T. Briscoe. 1992. Lexical operations in a unification-based framework. In *Lexical Semantics* and *Knowledge Representation*, *Proceedings of the First SIGLEX Workshop*, edited by J. Pustejovsky and S. Bergler. Berlin: Springer-Verlag. 101-119.
- ——. 1996. Semi-productive polysemy and sense extension. In *Lexical Semantics: The Problem of Polysemy*, edited by J. Pustejovsky and B. Boguraev. New York: Clarendon Press. 15-67.
- Falkum, I. L. 2010. *The Semantics and Pragmatics of Polysemy: A Relevance-Theoretic Account.* PhD Thesis. University College London.
- Fodor, J. A. 1975. The Language of Thought. Cambridge University Press: Harvard University Press.
- ——. 1998. Concepts: Where Cognitive Science Went Wrong. Oxford: Oxford University Press.
- ———. 2008. LOT2: The Language of Thought Revisited. Oxford: Oxford University Press.
- Fodor, J. A., and E. Lepore. 2002. *The Compositionality Papers*. Oxford: Clarendon Press.
- Gillon, B. S. 1992. Towards a common semantics for English count and mass nouns. *Linguistics and philosophy* 15 (6):597-639.
- ——. 1999. The lexical semantics of English count and mass nouns. In *The Breadth and Depth of Semantic Lexicons*, edited by E. Viegas. Dordrecht: Kluwer Academic Publishers. 19-37.
- Gordon, P. 1985. Evaluating the semantic categories hypothesis: The case of the count/mass distinction. *Cognition* 20 (3):209-242.
- Horn, L. R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In *Meaning, Form and Use in Context*, edited by D. Schiffrin. Washington DC: Georgetown University Press. 11-42.
- Horsey, R. 2006. The Content and Acquisition of Lexical Concepts. PhD Thesis, University of London.
- Jackendoff, R. 2002. Foundations of Language: Brain, Meaning, Grammar, Evolution. Oxford: Oxford University Press.
- Kennedy, C., and L. McNally. 2010. Color, context, and compositionality. Synthese 174 (1):79-98.
- Kilgarriff, A. 1992. Polysemy. PhD Thesis, University of Sussex.
- ——. 1995. Inheriting polysemy. In *Computational Lexical Semantics*, edited by P. Saint-Dizier and E. Viegas. Cambridge: Cambridge University Press. 319-335.
- Kilgarriff, A., and G. Gazdar. 1995. Polysemous relations. In *Grammar and Meaning: Essays in Honour of Sir John Lyons*, edited by F. R. Palmer. Cambridge: Cambridge University Press. 1-25.
- Krifka, M. 1995. Common nouns: a contrastive analysis of Chinese and English. *The generic book*:398-411.
- Lakoff, G., and M. Johnson. 1980. Metaphors We Live by. Chicago, Illinois: The University of Chicago Press.
- Leech, G. 1990 [1974]. Semantics: The Study of Meaning. Harmondsworth: Penguin Books.
- Leslie, A. M. 1987. Pretense and representation: The origins of "theory of mind.". *Psychological Review* 94 (4):412-426.
- Li, P., D. Barner, and B. H. Huang. 2008. Classifiers as count syntax: Individuation and measurement in the acquisition of Mandarin Chinese. *Language Learning and Development* 4 (4):249-290.
- Link, G. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In *Meaning, Use and the Interpretation of Language*, edited by R. Bäuerle, C. Schwarze and A. von Stechow. Berlin: Walter de Gruyter. 303-323.
- Macnamara, J. 1982. Names for Things: A Study of Human Learning. Cambridge, Massachusetts: The MIT Press.
- Morgan, J. 1978. Two types of convention in indirect speech acts. In *Syntax and semantics, Vol. 9: Pragmatics*, edited by P. Cole. New York: Academic Press. 261-280.
- Nunberg, G. 1978. The Pragmatics of Reference. Bloomington, Indiana: Indiana University Linguistics Club.
- ——. 1996. Transfers of meaning. In *Lexical Semantics: The Problem of Polysemy*, edited by J. Pustejovsky and B. Boguraev. Oxford: Clarendon Press. 109-132.
- ——. 2004. The pragmatics of deferred interpretation. In *Handbook of Pragmatics*, edited by L. R. Horn and G. L. Ward. Oxford: Blackwell. 344–364.
- Nunberg, G., and A. Zaenen. 1992. Systematic polysemy in lexicology and lexicography. In *Proceedings of Euralex II*, edited by H. Tommola, K. Varantola, T. Salmi-Tolonen and J. Schopp. Tampere: University of Tampere. 387-398.
- Ostler, N., and B. T. S. Atkins. 1992. Predictable meaning shift: Some linguistic properties of lexical implication rules. In *Lexical Semantics and Knowledge Representation, Proceedings of the First SIGLEX Workshop*, edited by J. Pustejovsky and S. Bergler. Berlin: Springer. 87-100.

Papafragou, A. 2005. Relations between language and thought: Individuation and the count/mass distinction. In *Handbook of Categorization in Cognitive Science*, edited by H. Cohen and C. Lefebvre. Amsterdam: Elsevier. 256-277.

- Pelletier, F. J. 1975. Non-singular reference: some preliminaries. *Philosophia* 5 (4):451-465.
- Pelletier, F. J., and L. K. Schubert. [1989] 2001. Mass expressions. In *Handbook of Philosophical Logic. Vol.* 10, edited by D. Gabbay and F. Guenthner. Dordrecht: Reidel, 2nd edition. 249-335.
- Pustejovsky, J. 1991. The generative lexicon. Computational Linguistics 17 (4):409-441.
- . 1995. *The Generative Lexicon*. Cambridge, Massachusetts: The MIT Press.
- Quine, W. V. O. 1960. Word and Object. Cambridge, Massachusetts: The MIT Press.
- Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive Psychology* 8:382-439.
- Spelke, E. 2000. Core knowledge. American Psychologist 55:1233-1243.
- Spelke, E. S., G. Katz, S. E. Purcell, S. M. Ehrlich, and K. Breinlinger. 1994. Early knowledge of object motion: Continuity and inertia. *Cognition* 51 (2):131-176.
- Sperber, D. 1994. The modularity of thought and the epidemiology of representations. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, edited by L. A. Hirschfield and S. A. Gelman. Cambridge: Cambridge University Press. 39-67.
- ——. 1997. Intuitive and reflective beliefs. *Mind & Language* 12 (1):67-83.
- Sperber, D., and D. Wilson. 1986/1995. Relevance: Communication and Cognition. Oxford: Blackwell.
- ——. 2002. Pragmatics, modularity and mind-reading. Mind & Language 17 (1/2):3-23.
- Ware, R. X. 1975. Some bits and pieces. Synthese 31 (3):379-393.
- Wilson, D. 1998. Linguistic structure and inferential communication. In *Proceedings of the 16th International Congress of Linguistics (Paris, 20-25 July 1997)*, edited by B. Caron. Amsterdam: Elsevier Sciences.
  - ——. 2005. New directions for research on pragmatics and modularity. *Lingua* 115:1129-1146.
- Wilson, D. and Carston, R. 2007. A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In *Pragmatics*, edited by N. Burton-Roberts. London: Palgrave. 230-259.
- Wilson, D., and D. Sperber. 1993. Linguistic form and relevance. Lingua 90:1-25.
- ——. 2004. Relevance theory. In *The Handbook of Pragmatics*, edited by L. R. Horn and G. L. Ward. Oxford: Blackwell. 607-632.