

Contents

1	Genericity – the standard view	5
1.1	Rules and their truth conditions	5
1.2	The standard theory of genericity	8
1.3	Predicate genericity	10
1.4	Nominal genericity I: the standard theoretical point of view . . .	14
1.5	Nominal genericity II: the two-way distinction approach	17
2	Kinds as sortal concepts – from a cognitive perspective	21
2.1	Overview	21
2.2	Things in the world	21
2.3	Categorisation and individuation	23
2.4	The development of kind-based object individuation	25
2.5	Linguistic input as trigger?	30
2.6	Speculating about developmental stages	32
2.7	Conclusions I: kinds as sortal concepts	34
2.8	Conclusions II: what does it mean for a kind to exist?	36
3	Kinds as ontological primitives – from a semantic perspective	39
3.1	Overview	39
3.2	Object reference and kind reference	40
3.3	Kinds as taxonomic categories	43
3.4	Kinds and concepts	45
3.5	Kinds and properties	47
3.6	Chierchia’s (1998) notion of kind	50
3.7	Lexical and formal semantics	53
4	Spatiotemporal localisation – type-token theories of genericity	57
4.1	Overview	57
4.2	Three-class systems	58
4.3	Two-class systems	62
4.4	On the sense of the term “referential”	66

4.5	Conclusions	67
5	Indefinite reference to kinds – the pumpkin crusher puzzle	69
5.1	Overview	69
5.2	Reference to subkinds and the pumpkin crusher	69
5.3	Indefinite NPs in Carlsonian object-level predications	72
5.4	Indefinite NPs in Carlsonian kind-level predications	74
5.5	Two kinds of kind-level predicates	76
5.6	<i>To invent</i> versus <i>to be extinct</i>	78
5.7	How the taxonomic reading comes about	80
5.8	Well-establishedness	84
5.9	Conclusions	85
6	Lexical classes of predicates – normal and Carlsonian kind-level predicates	87
6.1	Overview	87
6.2	Predicate classes according to the standard theory	88
6.3	Outlining the alternative	90
6.4	The representative object interpretation	91
6.5	Categorising statements as representative object interpretations	95
6.6	No reference to objects without reference to kinds	99
6.7	Carlsonian stage-level predicates	103
6.8	Conceptual restrictions on spatiotemporal localisation	106
6.9	What is special about Carlsonian kind-level predicates	108
6.10	Carlsonian kind-level predicates as existence predicates	111
6.11	Conclusions	113
7	Overt and covert articles – a kind-based DRT approach to the semantics of noun phrases	117
7.1	Overview	117
7.2	The grammatical role of overt articles	118
7.3	Forming complex predicates	123
7.4	Internal and external modification	125
7.5	Attributive adjectives	129
7.6	On the semantic function of NPs	132
7.7	The grammatical null article	134
7.8	The requirement of earliest application	136
7.9	Existentially interpreted bare plural NPs	140
7.10	The semantic contribution of a token-restricting modifier	143
7.11	Chierchia’s and Krifka’s explanations	146
7.12	Object-level arguments in kind-level positions	149

7.13 Conclusions 155

Chapter 1

Genericity – the standard view

1.1 Rules and their truth conditions

Every now and then we are confronted with sayings like the following. Please note the different types of noun phrases (NPs) forming the subject expressions in these sentences:

- (1) a. *Der Mensch lebt nicht vom Brot allein.*
the human live not of bread only
- b. *Ein Apache kennt keinen Schmerz.*
a apache know no pain
- c. *Die Männer sind alle Verbrecher.*
the men be all law-breakers
- d. *Böse Menschen kennen keine Lieder.*
bad people know no songs

What is the intended meaning of these sentences? They are used by speakers (in this case: of German) because they express a certain message, but which one? It is obvious that we cannot derive the message by decomposing the sentence structure into its syntactic constituents and reconstructing the sentence meaning from the meanings of its constituent parts. The sentences, as under (1), are usually not uttered in order to communicate their literal meanings. The utterance meaning of (1a) is not to express that someone who lives on bread only will get ill and die in the end. (1b) does not tell us that Apaches were insensitive to pain. (1c) does not encode the message that every male adult is a law breaker. And, of course, (1d) is not meant to express that bad people do not know music. Clearly, to judge whether these sentences are true or false, we must not interpret them literally.

What then are the truth conditions of these sentences? Intuitively speaking, such sayings express *rules*. Whether or not a speaker of German accepts the

sentences as true depends on whether or not she accepts the expressed rules. Let us look at the examples one by one. The rule expressed by (1a) is a statement about the human nature saying that a human being does not only have material needs, but also spiritual needs. (1b) is typically used to teach children that they must not, or need not, cry. The message is: people who do not show pain are brave. The rule of (1c) is a statement about the nature of men telling us that a man is morally defective by nature. (1d) expresses that a person who sings songs is likely to be a good person. Using the notation of Cohen (2001), we can represent the rules, i.e. the meanings of the sentences, roughly as follows¹:

- (2) a. $!(\text{human being}(x) \Rightarrow \text{has spiritual needs}(x))$
 b. $!(\text{not show pain}(x) \Rightarrow \text{brave}(x))$
 c. $!(\text{man}(x) \Rightarrow \text{immoral}(x))$
 d. $!(\text{sing}(x) \Rightarrow \text{good}(x))$

Thus, sayings are sentences expressing certain rules of everyday life. Like a law of physics is true if the processes of the physical world support the law, so can a rule of everyday life count as true if the processes of the world of everyday life support the rule. If this is the case, we say, following Cohen's example, that the rule is "in effect". Unlike the physical world, however, the world of everyday life is not objective but must be relativised to people's experiences, feelings, etc. This is why the rules of everyday life are not laws of nature, but rather social norms (cf. Cohen 2001:196). This is also the reason why there can be some event that contradicts a certain rule of everyday life without, at the same time, being a miracle²³.

Sayings like those under (1) express rules allegedly supporting the course of our everyday life. The characteristic of a rule designated by a saying seems to be that it is not reconstructable from the meanings of the syntactic constituents of the sentence. It is, however, not only sayings which are used to report on rules of everyday life. There are also rule-expressing sentences whose meanings *are* reconstructable from their constituent parts. It makes sense to distinguish two classes of such rules. The first class contains rules which are based on direct observations. Thus, in this case, we deal with statistical generalisations:

- (3) a. *The koala lives on eucalyptus leaves.*

¹Cohen distinguishes between *potential rules* that he represents as implications and *active rules*. The exclamation mark symbolises an operator mapping an implication onto an active rule.

²I take a "miracle" to be an event that contradicts the laws of nature.

³The fact that rules of everyday life cannot be verified by consulting the objective world is why they can be (mis-)used for ideological purposes (political propaganda, advertising, etc.).

- b. *An Indian rides without saddle.*
- c. *The good lawyers are expensive.*
- d. *Children like to sing.*

Investigating the actual life of particular koalas in Australia, one is soon led to draw the generalisation that koalas live on leaves of the eucalyptus trees only. Similarly, observing the riding technique of Indians, one is sooner or later led to the conclusion that it not just happened to be the case that the particular Indians observed were riding without saddle, but rather that Indians in general ride without saddle. Correlating the fees of particular lawyers with their success, one will soon convince oneself that, as a rule, quality has its price, even in the world of justice. Observing the behaviour of particular children, it simply jumps into the eye that children love to sing. Thus, the sentences under (3) express induced generalisations:

- (4) a. $!(\text{koala}(x) \Rightarrow \text{live on eucalyptus leaves}(x))$
- b. $!(\text{Indian}(x) \Rightarrow \text{ride without saddle}(x))$
- c. $!(\text{good lawyer}(x) \Rightarrow \text{expensive}(x))$
- d. $!(\text{child}(x) \Rightarrow \text{love to sing}(x))$

On the other hand, the rules of the second class do not result from abstracting across particular things and their behaviour in the world. It is rather rules which are being “supplied” to a new category of things when such a category (e.g., a new breeding, a new invention, a new fictional character) is being created. Rules of this second class are constitutive of the newly created category, therefore *defining* it:

- (5) a. *The 50-liter cow gives 50 liters of milk a day.*
- b. *A mermaid is a woman with a fish’s tail.*
- c. *The good get to heaven.*
- d. *Bishops move diagonally.*

The paradigm case is considered to be (5d) (cf. Carlson 1995, Cohen 2002). The particular way in which observable bishops move along a chess board is defined by a certain rule of the game of chess. The sentence expresses this rule. While statistical generalisations follow the principle “the behaviour determines the rule”, definitions follow the principle “the rule determines the behaviour”. The rule expressed by (5a) says that a new kind of cow “successfully” bred under economic pressure can give up to 50 liters of milk per day. Because of this property, which was, after all, the purpose of the breeding, this kind of cow is called “50-liters cow” or also “turbo cow”. (5b) is what we may call a linguistic definition. *Mermaid* is, like *bachelor* or German *Schimmel* (‘white

horse’), a nominal, the meaning of which is composed out of the meanings of other linguistic expressions. The rule expressed by (5b) reproduces the way the meaning of *mermaid* is composed out of the meanings of *woman* and *fish’s tail* (which, in itself, is composed out of the meanings of *fish* and *tail*). Nobody has ever observed a particular person getting to heaven. Therefore (5c) can express no statistical generalisation. Here, instead, we deal with a rule provided by the Christian religious believe defining the fate of “good people” as opposed to “bad people”, i.e. sinners. The rules expressed by the sentences under (5) can be represented as follows:

- (6) a. $!(50\text{-liter cow}(x) \Rightarrow \text{give 50 liters milk per day}(x))$
- b. $!(\text{mermaid}(x) \Rightarrow \text{woman}(x) \ \& \ \text{have a fish's tail}(x))$
- c. $!(\text{good person}(x) \Rightarrow \text{get to heaven}(x))$
- d. $!(\text{bishop}(x) \Rightarrow \text{move diagonally}(x))$

To recapitulate, the examples (3) and (5) illustrate two classes of rules. While the rules under (4) are statistical generalisations, the rules under (6) are definitions. In the present thesis, I presume the following:

It is linguistically irrelevant whether a rule is being induced on statistical grounds or whether it comes into existence as a definition. Linguistically relevant is only that there are sentences, namely generic sentences, which express rules (of either class). I shall argue that any such rule is a *kind predication*.

My view is partly contra to the standard view on modern research on linguistic genericity. Therefore, it makes sense to use the rest of this introductory chapter to give at least a brief characterisation of the standard view. In the chapters to follow, I will then continue to discuss the notion of a kind because it is crucial for my argument. Having (hopefully) clarified what I consider a kind to be (namely a sortal concept), I will carefully outline my proposal, which is, in a nutshell: generic sentences report on rules, and rules are kind-level predications.

1.2 The standard theory of genericity

The basic tenet of truth conditional semantics is the supposition that, to know the meaning of a sentence, is to know the conditions under which the sentence is true. Every sentence dictates, by virtue of its meaning, a set of veritable circumstances to count as true. The competent speaker knows the rules of her language well enough so that she can produce sentences which display exactly those truth conditions that she wants the sentences to display. Roughly speaking, while forming sentences, the speaker *composes* sentence meanings from

lexically fixed word meanings. The competent hearer, on the other hand, knows how to *decompose* the sentence structures presented to her by the speaker. The hearer reconstructs the truth conditions (i) by segmenting the sentence structure into its syntactic parts (roughly: content words and function words), (ii) by reading the meanings of these syntactic parts from her lexical knowledge base, and (iii) by calculating the sentence meaning from the partial meanings following the same syntactic rules that the speaker made use of in order to initially produce the sentence. In other words, the ability of the competent speaker or hearer to compose or decompose sentence meanings follows the so-called “Frege principle” or “principle of compositionality”, respectively. Here is the version of Krifka (2003):

The meaning of a complex expression is a function of the meanings of its immediate syntactic parts and the way in which they are combined.

Generic sentences represent a notorious problem for truth conditional semantic approaches. Consider the following sentences:

- (7) a. *Parrots are birds.*
- b. *Parrots can learn to speak.*
- c. *My parrot Harry can speak English.*
- d. *My neighbour Harry can speak English.*

Comparing (7a) and (7b), we see that a single expression, *parrots*, denotes different sets of object entities. In (7a), it denotes every parrot in the world without exception. In (7b), by contrast, it denotes only a subset of all the parrots, indeed a minority. At the same time, both sentences are true. The fact that one and the same expression has two completely different denotations undermines the principle of compositionality. Simply put: In order for semantic composition, from word meanings to sentence meaning, to follow predictable roads, word meanings have to be stable across syntactic environments. In (7), however, the linguistic environments seem to heavily influence the interpretation of certain expressions: Being the subject of the grammatical predicate *are birds*, *parrots* has a radically different denotation than being the subject of the predicate *can learn to speak*⁴. A comparison of (7c) and (7d) reveals something similar. Even if my parrot can say no more than “hello”, I am ready to tell everybody that my parrot can speak English. If my neighbour can say only “hello” in English, however, I will not consider my neighbour to be an English speaker. Abstraction over (7c) and (7d) leads us to sentence (8):

⁴This observation is Carlson’s (1977) main argument against an analysis of bare plurals as noun phrases involving a covert quantifier G: “[I]n order to know how to construe G, we must have knowledge of the particular predicate that is being applied. No other quantifier in English behaves even remotely in a similar fashion” (Carlson 1977:75).

(8) *Harry can speak English.*

If *Harry* refers to the bird, the ability to utter a single English word may be enough to make (8) come out true⁵. If *Harry* refers to the man, a single word is definitely not enough. This shows that, here again, the linguistic context influences the interpretation, in this case of the grammatical predicate. Again, truth conditional semantics faces a problem. How to deal with it?

The standard theory (Krifka et al. 1995) distinguishes between two basic ways in which genericity manifests itself in natural languages:

In the history both of philosophy of language and of linguistics, there have been two quite distinct phenomena that have been referred to or classified as ‘genericity’. The first is *reference to a kind* [...] In this usage a generic NP is an NP that does not refer to an “ordinary” individual or object, but instead refers to a kind [...] The second phenomenon commonly associated with genericity are propositions which do not express specific episodes or isolated facts, but instead report a kind of general property, that is, report a regularity which summarizes groups of particular episodes or facts. (Krifka et al. 1995:2)

Katz & Zamparelli (2005) refer to this as “nominal genericity”, in contrast to “predicate genericity”. The subjects in (7a) and (7b) are instances of nominal genericity, the predicates in (7c) and (7d) are cases of predicate genericity. Note that, under this approach, the single observable phenomenon, called genericity, stems from two completely different sources. Nominal genericity is explained by the ontological difference between kinds and “ordinary” objects, while predicate genericity is explained by some semantic mechanism “summarising” particulars. Let us discuss predicate genericity first.

1.3 Predicate genericity

In order to refer to non-generic predications, the term “episodic predication” is commonly used. (9) shows some English examples (after Katz & Zamparelli 2005):

- (9) a. *Tweety is flying.*
b. *After he will have finished his dissertation, John will smoke a cigar.*

⁵If a single word is not sufficient evidence for you to call a parrot an English-speaking parrot, just increase the vocabulary until you agree that the bird can speak English. What is relevant for the argument is the fact that different standards of comparison are crucial for deciding whether some creature can speak English or not, be it a child or an adult, a Japanese or an Englishman, a bird or a human.

- c. *A lion is eating raw meat.*
- d. *John must smoke a pipe if he wants to play the role of Sherlock Holmes.*

The sentence *John must smoke a pipe* in (9d) can also be understood generically. In this (somewhat unnatural) case, John should have the habit of smoking a pipe to qualify for the role of Sherlock Holmes. In the more natural reading, John would have to be ready to smoke a pipe *on stage* if he wants to play the role of Sherlock Holmes. This requirement does not necessarily include that John has to be a pipe smoker in his private life. This example illustrates that episodic predications are not restricted to indicative contexts, but can also appear in modal (in this case: deontic) contexts.

Generic predications are usually called “characterising predications”. (10) shows some examples (realised within an epistemic-modal context as in (10d), *John must smoke a pipe* must be understood as a generic, i.e. characterising, statement):

- (10) a. *Birds fly.*
- b. *John will smoke when he grows up.*
- c. *A lion eats raw meat.*
- d. *John must smoke a pipe given the type of ashes in the ash tray.*

Among characterising sentences, the standard theory distinguishes further between two different classes of predications. The examples under (10) exemplify so-called “habituals” which are almost always built around some action verb, like *fly*, *eat*, *smoke* etc. The underlying assumption is that the (action) verb, by its lexical meaning, denotes a set of particular events. If such a verb is chosen by the speaker to form the predicate in some linguistic utterance, this will, by default, amount to an (episodic) interpretation where reference is made to some particular event. Given special circumstances, however, namely given a generic context, this default can be overwritten. In this case, a generic operator additionally shows up in the semantic structure of the syntactic predicate with the effect of binding the event variable of the (action) verb. As a consequence, the utterance will not be understood episodically, but generically. Note that (at least in English) the appearance of the generic operator is not accompanied by any overt morphosyntactic signals, whatsoever. Therefore, one cannot tell from the form alone whether one deals with an episodic or with a generic predicate.

The second class of characterising predicates form so-called “lexically characterising predicates”. Katz & Zamparelli offer the following examples:

- (11) a. *John knows French.*
- b. *Fido is a dog.*

- c. *Fido likes meat.*
- d. *Dogs have four legs.*

Likewise, these sentences report on some kind of “general property” (recall the quote on page 6) characterising the referent of the subject expression. With these predicates, however, it is not plausible to make a generalisation over episodes responsible for the characterising interpretation. Although it is more or less accepted (following Davidson 1967) that dynamic predicates like, for instance, *smoke*, *fly* or *eat* denote smoking-events, flying-events, eating-events, it is by far not clear whether this supposition can be extended to cover stative predicates such as *know French*, *be a dog*, *like meat* or *have four legs*. The question is: What could count as a particular knowing-French-event (or a particular knowing-French-state, respectively) which, via habitualisation, would lead us to the meaning of *know French*? It is not by chance that there is no verb that would denote such an event (or state):

There is no episodic predicate which is morphologically related to *know French* (this would be an episodic verb which denotes events and which provides evidence that someone knows French), so the sentence *Italians know French* is a lexical characterizing sentence. (Krifka et al. 1995:17)

Thus, the strategy of semantic analysis used in connection with action verbs cannot be used to deal with the stative predicates exemplified by (11). A solution to overcome this problem in compliance with the standard theoretical assumptions is based on the idea that, in lexical-characterising sentences, the predicate is formed by an expression which *automatically* assigns a characterising property to its argument. As a consequence, it will never happen that the predication is interpreted episodically. As this automatic mechanism is supposed to be rooted in the lexicon, these predicates are called “lexical-characterising predicates”. According to Chierchia (1995), the respective stative predicates contain a certain morphological feature in their lexical-semantic structures that triggers the appearance of the generic operator in their local syntactic environments. The resulting verb phrases, in turn, amount to generic interpretations.

Let us recapitulate how the standard theory deals with predicate genericity. Assuming that verbal predicates “normally” denote episodically, generically interpreted predications are treated as deviations from normal. It is always the (invisible) presence of the generic operator, usually called “GEN”, in the semantic structure which is responsible for the deviation from the normal case. The source of GEN is either the context (in the case of action verbs) or a morphological feature in the lexical entry (in the case of (most) stative predicates). How does GEN solve the problem of the truth conditions of generic sentences? The answer is: not at all. How many English words or sentences must a per-

son say so that we would consider her to be a speaker of English? Postulating a generic quantifier does not answer this question. The question could be answered only if we would specify *how* this operator quantifies. The quantificational approach to genericity owes this piece of information, however. Consequently, generic predications remain an unsolved puzzle for the standard approach characterized here. Let us consider another well-known example:

(12) *Fred smokes.*

What are the truth conditions of (12)? How frequently must Fred smoke a cigarette so that we would consider him to be a smoker? Once? Ten times? Once a day? One package of cigarettes per day? Chierchia (1995) represents the truth conditions of the sentence as follows:

$$(13) \underbrace{\text{GEN}}_{\text{quantifier}} e \underbrace{[\text{C}(f, e)]}_{\text{restrictor}} \underbrace{[\text{SMOKE}(f, e)]}_{\text{nuclear scope}}$$

The interpretation is represented as a tri-partite quantification structure. The quantifier is the generic quantifier GEN, which quantifies over events. Specifically, it quantifies over events fulfilling the conditions noted in the restrictor. The semantic structure expresses that, whenever there is an appropriate event, the event will be a smoking-event with Fred as the smoker. The latter information is noted in the nuclear scope. The variable C in the restrictor is the mysterious part. It represents the so-called “felicity conditions” of an event:

[W]e must bear in mind that each activity or state comes with a set of “felicity” conditions. For example, in order for Fred to engage in smoking, he must feel like it; that is, he must intend, or perhaps feel compelled, to do it; he also must be in a place where there is enough oxygen, he must not be asleep or disabled, and so on. So in evaluating (12), we have to look at worlds similar to ours where the felicity conditions for smoking are met. These felicity conditions are what provides a value for the variable C in the restriction. (12) is true iff in all the worlds maximally similar to ours where the felicity conditions for Fred’s smoking are met, he does smoke. (Chierchia 1995:195; numbers adjusted)

Chierchia certainly manages to formalise the truth conditions of the sentence. Honestly speaking, however, we do not gain all too deep insights into the nature of generic predications. What Chierchia’s analysis boils down to is: the characterising sentence *Fred smokes* is true if Fred is smoking whenever nothing hinders him from doing so, neither physical nor psychological barriers. Unfortunately, the interesting thing about characterising statements is not covered at

all by this approach: whether or not I consider Fred to be a smoker depends on whether *his* behaviour fulfills *my* criteria for being a smoker.

Any attempt to specify *objective* truth conditions must fail because there simply are no objective criteria determining the concept ‘smoker’. All there is are real smoking-events and real people who sort people into smokers and non-smokers on the basis of whether or not they participate in the real smoking-events. If a person draws the conclusions that Fred smokes, then this has nothing to do with the fact that Fred sometimes does not feel like smoking or that he cannot smoke when he is asleep. Classifying or not classifying Fred as a smoker has rather to do with the *subjective* criteria of the classifying person, and these subjective criteria may vary from person to person.

Imagine two mothers each having a teenage son called Fred. Each Fred was secretly smoking a cigarette when his mother caught him in the act. Seeing her son smoking a cigarette, one mother’s (the orthodox one’s) whole world came crashing down around her because, for her, it meant that her Fred is a smoker. The other mother (the liberal one) reacted in a much more relaxed way. Aware that almost every teenager wants to try out smoking, she did not conclude from the one episode that her Fred was a smoker. This scenario shows that it depends on your personal criteria whether you accept the sentence *Fred smokes* as true when confronted with a single episode of Fred smoking. The definition of smoker depends on your point of view.

Let us recapitulate. Why is it so difficult for truth conditional semantics to cope with the meaning of generic predications? Because truth conditional semantics attempts to translate linguistic meanings into real, objective conditions. This fails, because the meaning of generic expressions is the result of (cognitively) abstracting away from real conditions.

1.4 Nominal genericity I: the standard theoretical point of view

In the domain of nominal genericity, the standard theory (Krifka et al. 1995) follows a completely different approach to account for the peculiarities of generic expressions. While the genericity of a generic predicate is viewed as the result of quantifying over what the non-generic predicate denotes, the genericity of a generic noun phrase is viewed as the result of the noun phrase denoting not ordinary objects, but kinds. In other words, nominal genericity is traced back to an ontological difference. To understand the distinction between “ordinary” entities and “non-ordinary” entities, consider example (14). In response to question (14a), the mycologist Dr. Ewald Gerhardt gives the very precise answer (14b)⁶:

⁶In an interview in the Frankfurter Allgemeine Sonntagszeitung, 6 Nov 2005, Nr. 44, p.16.

- (14) a. *Gibt es denn viele giftige Pilze?*
Are there many poisonous mushrooms?
- b. *Artenmäßig nicht, wir haben vielleicht zehn bis zwölf wirklich tödlich giftige Pilzarten, aber die können individuenmäßig in einem Wald häufig vorkommen.*

Not as far as kinds are concerned, we have approximately ten to twelve actually deadly poisonous mushroom kinds, but as far as individuals are concerned, they may come in large numbers in a forest.

Noticing the ambiguity of the question, Dr. Gerhardt gives answers to both readings. One may understand (14a) as asking for the number of kinds of poisonous mushrooms, or as asking for the number of particular poisonous mushrooms (=mushroom objects). Dr. Gerhardt explains that there are relatively few poisonous mushroom kinds (given the number of known mushroom kinds), but relatively many poisonous mushroom objects (given the number of mushroom objects in a forest). This ambiguity, noticed by the mycologist, has also been noticed by linguists: noun phrases may determine their denotation either in the domain of objects or in the domain of kinds of objects.

Among semanticists, there is no agreement on the notion of kind reference as opposed to object reference. To set the stage for discussion, let me briefly introduce the standard-theoretical point of view, i.e. the position maintained by Krifka et al. (1995). The standard theory follows Carlson (1977) in assuming that there are certain predicates which syntactically select for kind terms semantically assigning a property to the respective kind. The prime example of such a Carlsonian kind-level predicate is probably *be extinct*. The standard theory builds further theoretical generalisations on top of the important observation that definite and indefinite singular noun phrases are interpreted differently if they appear in the argument slot of a kind-level predicate like *be extinct*:

- (15) a. *The Neanderthal man is extinct.*
b. *A Neanderthal man is extinct.*
c. *The dodo became extinct in the 17th century.* [Krifka 2001]
d. *A dodo became extinct in the 17th century.*

While (15a) is a statement about the kind ‘Neanderthal man’, (15b) must be understood as a statement about a subkind of the kind ‘Neanderthal man’. Virtually the same holds for the pair (15c) and (15d): the definite refers to the kind ‘dodo’, the indefinite refers to a subkind of the kind ‘dodo’. On the basis of observations like these, the standard theory distinguishes between two modes of kind reference. “Non-taxonomic kind reference” is reference to the kind named by the common noun forming the head of the noun phrase. “Taxonomic kind

reference” is reference to some subkind of the kind named by the head noun. In English, non-taxonomically referring kind terms can be formed by definite singular NPs or by bare plural NPs (16). Apart from that, they are formed by bare singular mass noun phrases and by kind names like, for instance, *homo sapiens*. Taxonomically referring kind terms can be realised as indefinite singular NPs, modified definite singular NPs, as a syntagma made of a number word plus a noun, as a demonstrative plus a noun, or as a quantifier plus a noun (17). This is, at least, Krifka et al.’s position, illustrated by the following examples:

- (16) a. ***The panda*** will become extinct soon.
 b. ***Pandas*** will become extinct soon.
- (17) a. *The dolphin* is ***a whale***.
 b. ***The whale which was most recently put under protection*** is the blue whale.
 c. ***Two whales***, namely the blue whale and the fin whale, were put under protection.
 d. ***This whale***, namely the blue whale, is nearly extinct.
 e. ***Every whale*** (from the pygmy whale to the blue whale) is protected by law.

From the standard theoretical point of view, this distribution of forms onto functions represents the English version of a general principle, which is valid across languages. Another language may show another pattern, dependent on its specific grammatical rules. Krifka (2001) describes the general principle underlying the observable language-specific patterns as follows. Note that, as in the case of predicate generic, the non-generic use is considered to be the basic use:

[T]ypically, natural languages use expressions that basically apply to specimens of a kind (common nouns, also called appellativa) also to refer to the kind itself. [...] Also, natural languages use expressions that basically apply to specimens of a kind to refer to subkinds. This re-use of common nouns appears to be widespread in human languages. But the specific ways how languages make use of common nouns to refer to kinds (or subkinds) depend on general features of their grammar, and can vary considerably. (Krifka 2001:1-2)

We see that, according to Krifka, a noun used for denoting at the kind-level is something like a recycled noun. Originally (“basically”), a noun denotes at the object-level. Once introduced, however, nouns may be “re-used” to denote at the kind-level. Of course, it is not nouns, but noun phrases which act as referring expressions with the specific morphosyntactic shape that it acquires when being determined by language-specific grammatical rules. Krifka assumes that the

specific morphosyntactic structure “accompanying” a common noun to form a noun phrase serves the primary purpose of object reference. Because the basic function of morphosyntactic structure is, accordingly, denotation at the object-level, it can be anticipated that the (generally possible) re-use of noun phrases at the kind-level conflicts with the original function in particular cases. This happens in the case of indefinites. Reference to the kind ‘Neanderthal man’ by means of the indefinite NP *a Neanderthal man* is impossible, and one must resort to another interpretation, namely to the taxonomic interpretation. To refer to the kind ‘Neanderthal man’, the definite singular NP *the Neanderthal man* or the bare plural *Neanderthals* must be chosen instead.

1.5 Nominal genericity II: the two-way distinction approach

In recent years, some scholars have cast doubt on the validity of the standard view. In particular, Zamparelli (1998) and Dayal (2004) act as advocates of an “emancipation” of the kind domain. In their view, a common noun may *basically* denote either a set of objects or a set of kinds. Thus, object reference is no longer viewed as the primary mode of reference. Instead, the kind domain and the object domain are equally accessible as possible referential domains.

This approach enables a unified semantic analysis of the English definite article. The standard theory commits itself to a distinction between a normal (non-generic) use of the definite article and a special (generic) use of the same form. On the one hand, there has to be a lexical entry for the normal definite article. Its lexical meaning can be described by a function mapping a set of objects onto its maximum element. On the other hand, there has to be a second lexical entry, homomorphic with the first one, but different in its meaning. This second definite article is somehow associated with a generic meaning⁷.

In Zamparelli’s and Dayal’s system, by contrast, there is only one definite article listed in the lexicon. Its meaning corresponds to a function which maps a set of entities, be it objects or kinds, onto its maximum element. Obviously, this approach presupposes that every common noun has the basic semantic potential to characterise either a set of objects or a set of kinds. What set of kinds may a common noun denote? Katz & Zamparelli (2005) propose the following two assumptions:

- (i) Nouns can denote (more or less salient) sets of subkinds.
- (ii) The set of subkinds includes the superkind.

⁷According to a popular view, the generically used definite article displays no meaning at all, but functions as a syntactic expletive (cf. Vergnaud & Zubizarreta 1992; Longobardi 1994).

Given these assumptions, the semantic composition of a generically used definite singular NP is as follows. The common noun characterises, in accordance with (i) and (ii), the set of subkinds including the superkind. Syntactically combining with it, the definite article maps this set onto its maximum element which is, by stipulation, the superkind. The resulting definite NP denotes – in accordance with the facts – the (super)kind that is named by the common noun. Consider an example:

- (18) a. $[[tiger]] = \{ \text{'tiger'}, \text{'Caspian tiger'}, \text{'Sumatra tiger'}, \text{'Bengali tiger'}, \text{'Amur tiger'}, \dots \}$
 b. $[[the\ tiger]] = \text{'tiger'}$

This analysis covers instances of the kind-referring noun phrase *the tiger* in examples like the following:

- (19) a. *The tiger is a tiger.*
 b. **The tiger is a subkind of tiger.*
 c. *The tiger is a (subkind of) mammal.*
 d. *The tiger is nearly extinct.*

Furthermore, the approach extends easily to definite plural NPs. The only requirement is that the morphological plural marker is the syntactic manifestation of an operator mapping a set of singularities (in this case: subkinds) onto a set of pluralities (e.g., Chierchia 1998). Plausibly enough, the maximum plurality, i.e. the plurality that accumulates the highest number of singularities, can be viewed as the maximum element out of the set characterised by a plural noun. Then, this plurality of subkinds represents the referent of a subject NP as in (20):

- (20) *The tigers belong to the cat family.*

Thus, the “2-way distinction approach” (Zamparelli 2002), advocated by Zamparelli and Dayal, allows for a compositional semantic analysis in which the semantic contribution of the definite article is uniform across nongeneric (object-level) and generic (kind-level) usages. In Dayal’s words:

The proposal I make in connection to the definite singular generic is that there is nothing special about the determiner. It is the common noun that has two possible denotations, one in the object domain, the other in the taxonomic domain. (Dayal 2004:396)

Note that, once we take reference to the kind named by the head noun of a definite NP to be reference to the superkind, Krifka et al.’s distinction between

non-taxonomic and taxonomic kind reference becomes, strictly speaking, redundant. Every kind of kind reference is taxonomic reference! Reference to the superkind is simply a species of taxonomic kind reference. This is why Dayal calls the kind domain taxonomic domain.

In a way, Dayal's proposal exceeds Zamparelli's. Dayal wants to analyse not only the definite article, but downright *every* determiner within the 2-way distinction approach:

All determiners can combine with both meanings of the common noun. In the first case they yield the familiar readings, in the second case the taxonomic readings. The singular kind is simply an instance of a taxonomic noun phrase. It is therefore to be expected that it can occur with definite as well as indefinite determiners in English, depending on its status as a familiar or novel entity in the discourse. (Dayal 2004:396)

Accordingly, every morphosyntactic NP-type in English should be capable of displaying kind-level readings. In support of her programmatic proposal, Dayal presents the following examples:

- (21) a. *The whale is a mammal.*
b. *A lion roars.*
c. *The crustaceans evolved simultaneously.*
d. *Crustaceans can evolve simultaneously.*

Dayal's expectation, namely that NP-reference to the kind named by a singular head noun ("singular kinds") is possible also with indefinite NPs, opposes the standard-theoretical point of view according to which, when interpreted at the kind-level, indefinite singular NPs can only refer to subkinds. Recall (15b) as opposed to (15d). In other words: Dayal owes us a precise answer to the question why in (15b) and (15d), reference to the kind 'Neanderthal man' or 'dodo', respectively, by means of an indefinite NP, is impossible. In addition, a proponent of the standard theory will not be convinced by example (21b), because the subject *a lion* can alternatively be viewed as an object-level expression appearing within the scope of a generic quantifier. More needs to be said to defend Dayal's position. In chapter 5, I address these open questions related to indefinite kind-level reference. We will find good reasons to adopt Dayal's perspective in the end.

Chapter 2

Kinds as sortal concepts – from a cognitive perspective

2.1 Overview

Current semantic theories explain the generic use of linguistic expressions, in particular nominal expressions, by postulating that reference can be made not only to “ordinary” objects, but also to a special sort of “abstract” entities, viz. kinds (see chapter 1). In this dissertation I make a point of treating kinds as sortal concepts. The goal of the present chapter is to motivate this position. I shall report on some recent psychological studies informing about the genesis and the purpose of kinds aka sortal concepts: the system of kinds is something like a “mental catalogue” serving to categorise and individuate objects (sections 3 to 7). Being mental by nature, kinds “exist” in a different way than real objects exist. Objects exist in space and time, kinds do not. In other words, objects are particulars, kinds are non-particulars. Interestingly, the meaning of the natural language predicate *exist* turns out to be restricted to spatiotemporal existence. This explains why we can say that some particular exists (I, for example, am existing). But what does it mean to say that some non-particular exists? Finding the correct answer to this question constitutes the frame of the chapter (sections 2 and 8).

2.2 Things in the world

Sentences are uttered for the purpose of communicating states of affairs. This is at least the presumption from which I start in the present thesis. The speaker forms a sentence in order to inform the hearer of some interesting enough con-

stellation of things in the world¹. That the meaning of a sentence entails information about a constellation of things implies that, in general, the meaning of a sentence consists of at least two semantic components. The first component informs about a thing (or a number of things), the second component informs about the constellation of this thing (or these things) with respect to other things in the world. That is to say, the first component informs about the topic of the utterance. Choosing a certain thing as the topic implies its existence within the domain of possible referents, for trivial reasons. The second component informs about a certain property. By choosing this property, the speaker asserts that the world is such that the topic entity carries this property. On the formal plane this amounts to the Aristotelian picture according to which a declarative sentence consists of two constituents: one expression, the subject, symbolising the topic, and one, the predicate, symbolising the property.

Thus, a state of affairs communicated by means of a sentence consists of at least (i) a thing and (ii) a property carried by the thing. Now, if we draw our attention to particular sentences, e.g. in English, we diagnose a puzzling contrast: what is formally one and the same expression in some cases refers to some real, spatiotemporally localised thing and in other cases to some abstract, not spatiotemporally localised thing. For the time being, let us, in the latter case, speak of “reference to the thing-as-such”, independent of any specific spatiotemporal embedding. The sentences under (1) are usually understood as talking about a concrete dog or a concrete Apache, respectively. By contrast, the sentences under (2) are usually interpreted as relating to the dog-as-such or the Apache-as-such, respectively.

- (1) a. *The dog belongs to the FBI.*
b. *An Apache is riding through the prairie.*
- (2) a. *The dog belongs to the order Carnivora.*
b. *An Apache rides without saddle.*

Within modern semantic theories, spatiotemporally localised things are usually called “objects”, whereas things-as-such lacking any concrete spatiotemporal localisation are usually called “kinds”. For example, Krifka (1995) terminologically distinguishes between “real objects” and “abstract kinds”. A complication results from the fact that the very pioneer of kind-research takes a different stand: Carlson (1977) considers “objects” *not* to be spatiotemporally localised entities. Instead, they form together with “kinds” the domain of abstract

¹More precisely: a *declarative sentence*. I ignore question sentences (formed to bring about that the hearer informs the speaker of some constellation of things in the world) and imperative sentences (formed by the speaker to cause the hearer to manipulate on some constellation of things in the world).

“individuals”. These abstract “individuals” contrast with the spatiotemporally localised entities which he calls “stages”². Within the present work I follow Krifka’s terminology.

It is more or less standard among semanticists to trace the contrast observed with regard to (1) versus (2) back to an ontological contrast: there are *basically* two sorts of things that a speaker may choose as the topic of an utterance, viz. objects or kinds. I subscribe to this point of view. There is no general agreement, however, about the nature of kinds as opposed to objects. What is a kind? A related question is: What does it mean for a kind to “exist”? Consider the following examples collected from the internet:

- (3) a. *Five years passed before scientists were certain that **the platypus** actually existed and was not a fake.*
- b. *The unravelling of the mystery of Brownian motion a century ago by Einstein brought to an end all debate on whether **atoms** really existed or they were merely mental constructs.*
- c. *“Some parents compare their kids to **a kind of perfect child** that doesn’t really exist”, says Ken West, Ph.D., professor of counseling at Lynchburg College in Virginia.*
- d. *Researchers insist that **centaurs** actually existed.*

What does it mean for a kind to “exist”? This is not a genuinely linguistic question, but without addressing it, it will be difficult to appreciate the particular kind-based theory I am proposing to explain semantic contrasts as in (1) and (2). Therefore, there is virtue in explaining the notions of kind and object that I take for granted in this thesis.

2.3 Categorisation and individuation

Carlson in his dissertation describes the following situation:

You are on a picnic and have begun to eat. Out of the bushes pops a ground squirrel, which you throw a scrap of food to. It eats and disappears into the bushes. A few moments later, from another direction, a ground squirrel pops out of the bushes. Since all ground squirrels look pretty much alike (at least to me), there is no way of telling whether or not this second appearance of a ground squirrel is another one, or the same as before. In any event, you feed the second one (the first one?), and it scampers off into the bushes. This process is repeated several times, with only one ground squirrel appearing at a given time, and all

²A detailed discussion of Carlson’s theory will be given in chapter 4.

appearances seeming quite alike. At this point, you might become curious as to whether or not you have been witnessing appearances of the same ground squirrel, or of several. (Carlson 1977:113)

It is obvious that “appearances” are spatiotemporal by nature: if something appears to you this will always happen embedded within a specific situation (space-time segment). This may be, for instance, a picnic at a certain place at a certain time. Recognising an appearance as a squirrel is the result of your cognitive performance of *object-categorisation*. If the appearance is in fact a squirrel, you have successfully, i.e. in line with the facts, categorised and not been misled. To categorise an appearance means to identify a spatiotemporal object individual as an instance of a certain kind.

Another cognitive performance is it to identify two instances of the same kind as one and the same object. The question raised by Carlson’s picnic scenario – whether you are each time confronted with one and the same squirrel or with different squirrels – is the question of *object-individuation* (cf. Xu 2005). If you take each squirrel appearance to be an appearance of the same squirrel, you have individuated a squirrel object. If it is in fact always the same squirrel, then you have successfully individuated an object of the world around you.

As said above, categorisation means to identify the identity of an object as an instance of a (familiar) kind. Now, individuation means to *reidentify* the identity of an object as a (familiar) object.

No need to say that the cognitive capacity to reidentify objects is of central relevance for us. An important part of the world in which we live is made up of objects whose existence time extends far beyond the period of time during which these objects typically appear to us. As a consequence, we may meet such objects more than once in our lifetime (i.e. in our existence time):

[H]uman adults are committed to a world populated with individuals (people, objects, or events) that persist – we may encounter Joe Schmo on both Monday and Wednesday; we may sit in the same office chair everyday of the week; and we may go to the first and ninth innings of the same baseball game (Xu 2005:4)³

Carlson (1977) uses the picnic scenario in order to introduce the ontological categories on which he bases his semantic theory. The appearances of a squirrel are called “stages”. For Carlson, stages are the only spatiotemporal entities – the world manifests itself in space and time in form of stages. “Objects”, by contrast, are considered to be the result of a cognitive abstraction. Summarising a series of stages to constitute an individual of a higher level of abstraction,

³When I quote Xu (2005), I give page numbers in accordance with the paper available on Fei Xu’s web page.

human beings construct things that form an ontological domain in its own right besides the spatiotemporal domain of stages. The domain of kinds is the result of another cognitive abstraction whereby a series of stages and objects is summarised to form a new thing, viz. a kind.

Which stages exactly are summarised to form an object? This is, according to Carlson (1977:104), not a linguistic problem and need not bother the semanticist. Relevant for the (truth conditional) semanticist is not *how* ontological domains evolve, but only *that* they do. There is a stage-level domain, an object-level domain and a kind-level domain, and sentence designations, i.e. states of affairs, are describable as constellations of stages and/or objects and/or kinds. That is to say, in the Carlsonian ontology (“world”) relevant for linguistic reference there exist side by side stages, objects and kinds. What a sentence like, for instance, *the platypus exists* expresses is that in the “world” relevant for linguistic reference there exists a thing of the ontological sort kind, viz. the kind ‘platypus’.

My own basic assumptions differ from Carlson’s assumptions. I presume that the real world manifests itself in space and time in form of *objects*. Objects come into existence at a certain point of time, exist for a period of time along some continuous spatial path, and cease to exist at a later point of time. Ourselves objects, we navigate through this world of objects during our whole lifetime. As advantageous turned out to be an “object recognition system” enabling us to identify (categorise) and reidentify (individuate) the objects of the real world.

2.4 The development of kind-based object individuation

A series of psychological studies indicates that there is at least one kind which already at a very early stage, at the age of two months, belongs to the human cognitive equipment⁴: the kind ‘object’ (vgl. Carey 1995:127). It is the availability of this kind that enables the child to recognise objects within the flow of sensations it is exposed to⁵. Accordingly, to recognise an object is to cat-

⁴Spelke and Carey, among others, suspect that humans are born with such an object concept. However, this conclusion is no necessity, given the present state of science. It may also be the case that the child acquires the kind ‘object’ via interaction with its environment during the first two months of its life or even before that, in the prenatal phase. (Spelke et al. 1992:627). I therefore want to take an agnostic stand. Linguistically relevant is anyway only the fact that the child possesses the capacity to identify objects prior to the developmental stage when it shows the first (proto)linguistic behaviour.

⁵I put it as if objects existed in the world independently of human conceptualisation. May the philosophers discuss whether this really is the case. Relevant for us is only that,

egorise an appearance as an instance of the kind ‘object’. The object concept (=the kind ‘object’) provides knowledge about possible objects. Possessing the object concept, the child “knows”, for instance, that objects must move along spatiotemporally continuous paths and that two objects can never occupy the same space (Spelke 1990, Spelke et al. 1992)⁶.

Thus, from early on human beings are able to identify (physical) objects existing in the world they live in, i.e. to categorise objects. On the basis of this capacity, human beings go on to learn to trace the identity of an object through space and time, i.e. to individuate an object. Evidence that an object appearing at a certain point of time is the same object as an object that had already appeared beforehand comes from basically three sources: from information concerning the spatiotemporal localisation of the “later object” relative to the “earlier object”, from information concerning properties of the “later object” in comparison to properties of the “earlier object” (size, colour, texture, etc.), and, finally, from information concerning the kind membership of the “later object” compared to the “earlier object” (Xu, Carey & Quint 2004; Xu 2005). If the “later object” occupies the same space, carries the same properties and belongs to the same kind as the “earlier object”, there is good reason to conclude that both objects are in fact one and the same object.

Spatiotemporal information may suffice to detect the identity of two objects perceived at two different times: if it is evident that the “later object” is connected to the “earlier object” via a spatiotemporally continuous path, one can be sure to deal with one and the same object. But what if, as in the case of Carlson’s picnic scenario, the task of object individuation cannot be mastered on the basis of spatiotemporal information? To answer this question, Xu (2005) reports on a series of behavioral experiments the results of which she interprets as follows: human beings have at their disposal basically *two strategies of object individuation*. The first of these is an “object based system”, as she calls it. In this case the question of whether one deals with one or two objects is being resolved based on spatiotemporal information (and, to a lesser degree, also on perceivable properties). The second is a “kind based system”. In this case it is kind membership that provides the clue as to whether or not two object appearances are appearances of the same object (again, information about perceivable properties may influence the decision to a certain degree).

The reason why Xu carefully distinguishes between these two strategies is that, as the experiments show, infants younger than 10 months are able to indi-

to an important degree, semantic patterns are determined by a “commonsense ontology” (cf. Diessel 1993, Bach 1986) that involves the autonomous existence of objects.

⁶The kind available from early on is the kind ‘physical object’ (Xu 1997), to be more precise. This kind is the prerequisite to develop basic-level categories. More on that in section 6.

viduate objects on the basis of spatiotemporal evidence, yet not on the basis of kind membership. By contrast, at the age of 12 months children (like adults) also make use of kind information in order to individuate objects. Let us have a closer look at the relevant studies.

If sufficient spatiotemporal evidence for the nonidentity of two objects is presented to ten-month-olds (i.e. if two objects are shown at the same time), they, like twelve-month-olds and like adults, will conjecture the existence of two objects. I schematise the respective type of experiment⁷:

1. An infant is presented two objects on a stage
2. A screen is put up hiding the objects from the infant's sight
3. From behind the left side of the screen an object appears and disappears back behind the screen
4. From behind the right side of the screen an object appears and disappears back behind the screen
5. Steps 3 and 4 are repeated several times
6. The screen is removed
7. 1 On the stage, there is one object
7. 2 On the stage, there are two objects

Figure 2.1 illustrates the temporal course of the steps 3 to 4 (to be read from top to the bottom).




Figure 2.1: experimental schema

⁷See Xu (2005); Xu, Carey & Quint (2004) for details.

The result 7.1, that there is only one object on the stage, is a surprise from the perspective of an adult. The result 7.2, that there are two objects, by contrast, is to be expected from the adult's perspective. The fact that ten-month-old infants pay considerably more attention (a longer looking time) to the result 7.2, demonstrates that its expectations are in accordance with the adult's expectations. Arguably, the infant makes use of the same kind of information:

[E]ven very young infants employ spatiotemporal criteria in the service of object individuation (Xu 2005:7)

[I]nfants as young as 2 months of age use spatiotemporal information in establishing object identity. When spatiotemporal discontinuity is detected, very young infants establish a representation of two numerically distinct objects (Xu, Carey & Quint 2004:156)

What happens if the experiment is run leaving out step 1, i.e. withholding spatiotemporal evidence for the nonidentity of two objects? Note that Carlson's picnic scenario instantiates exactly this type of experiment (the bushes take over the role of the screen):

1. An infant is presented a screen on a stage
2. From behind the left side of the screen an object appears and disappears back behind the screen
3. From behind the right side of the screen an object appears and disappears back behind the screen
4. Steps 3 and 4 are repeated several times
5. The screen is removed
6. 1 On the stage, there is one object
6. 2 On the stage, there are two objects

Xu, Carey & Quint (2004) conducted a series of experiments of this type with twelve-month-old infants systematically varying the kind of left-hand object appearances and the kind of right-hand object appearances. Let me briefly summarise the results: In those cases where the objects appearing to the left of the screen and those appearing to the right of the screen were equal with regard to all of their perceivable properties, twelve-month-old children showed no increased attention, neither when they were confronted with outcome 6.1, nor when confronted with outcome 6.2. Apparently, both results were in equal measure compatible with the children's expectations. Xu, Carey & Quint observed the same behaviour also in cases where the object appearance to the left

differed from the object appearance to the right with respect to one single property. Whether, for instance, a green ball appeared to the left of the screen and a red ball (of the same size and texture) to the right, or whether a small ball appeared to the left and a bigger ball (of the same colour and texture) to the right, the twelve-month-old always reacted in much the same way independent of whether the outcome was 6.1 or 6.2.

In the next series of experiments the complexity of property difference between the left-hand and the right-hand object appearances was increased. But even if it was a small soccer ball (5 cm in diameter) decorated with orange, green and white hexagons that appeared to the left and a large bright red ball (9 cm in diameter) covered with glitters that appeared to the right – the children did not show any sign of surprise when, after the screen had been taken away, they encountered the small soccer ball alone. The infants likewise failed to establish representations of two objects when it was a small cup painted bright yellow with green and red vertical stripes and an orange handle that appeared to the left-hand side of the screen and a large semitransparent cup with a bright orange rim, decorated with red, green and yellow squares and circles and with a blue handle that appeared to the right-hand side of the screen.

There was one study, however, in which the twelve-month-old children did show significantly increased attention under outcome condition 6.1. This was observable when, for example, the object appearance to the left was a ball and the one to the right was a cup, even though the cup and the ball were of equal size and had the same colour⁸.

Xu, Carey & Quint reason that children at the age of 12 months seem to conjecture the existence of two objects in lack of spatiotemporal evidence in those cases where the differences in gestalt are such that an adult would classify the objects as belonging to two different basic-level categories:

[T]hey only succeeded in establishing a representation of two distinct objects when the shape differences were (what for adults) cross-basic-level-kind differences and failed to do so when the shape differences were (what for adults) within-basic-level-kind differences.

(Xu, Carey & Quint 2004:180)

In an earlier study Xu & Carey (1996) showed that at the age of 10 months, children are not yet sensitive for differences in kind membership: neither outcome 6.1 nor outcome 6.2 leads to increased attention, even if the objects are as different as, for instance, a rubber duck and a ball that appear to the left and to the right of the screen. From all these observations the researchers conjecture

⁸“The cup was 8 cm tall and 9 cm at its widest; the ball was 8,5 cm in diameter. The cup was a sippy cup with two handles. Both the cup and the ball were painted with pink and yellow stripes” (Xu, Carey & Quint 2004:176)

that at an age of approximately 12 months, children begin to mentally represent basic-level kinds (Xu, Carey & Quint 2004:180). Xu summarises the picture in the following words:

On the one hand, I am sympathetic to the view that human infants are born with a mechanism – the object based attention system – that carves up the world into distinct units. On the other hand, I also suggest that infants’ worldview undergoes fundamental changes: They begin with a world populated with objects and substances. By the end of the first year of life, they begin to conceptualize a world populated with sortal kinds (and perhaps substance kinds as well). In this new world, objects are thought of not as “qua object” but rather “qua dog” or “qua table”. (Xu 2005:33)

2.5 Linguistic input as trigger?

The task of developmental psychology is not only to give a pure description of developmental stages, but also to identify the driving processes behind these developments. As far as the genesis of the object based system of individuation is concerned, Xu subscribes to Spelke’s (1990) position considering it to be innate (but compare footnote 4). As far as the kind based system is concerned, Xu suggests the following:

Many have noted that infants begin to comprehend and produce their first words by the end of the first year, and many of their first words are nouns for object categories. I suggest that it is not a coincidence that along with acquiring their first words, infants also begin to develop a kind-based system of individuation. Some recent studies from my laboratory provide some initial evidence that perhaps language plays a causal role in this process. (Xu 2005:24)

The studies mentioned here are once more variations of the already known type of experiment (figure 2.1). The stage is again set by the evidence (Xu & Carey 1996) that, in lack of spatiotemporal information, ten-month-old infants do not individuate two objects on the basis of unequivocally distinct kind membership. The new studies were conducted with nine-month-olds. The innovation was that this time the different kinds of objects appearing to the left-hand and to the right-hand of the screen were accompanied by linguistic utterances. Once, under the so-called two-word condition, the utterance accompanying left-hand side appearances differed from the one accompanying right-hand side appearances (e.g., “Look, a duck!” versus “Look, a ball!”). Once, under the one-word condition, they were equal (e.g., each time “Look, a toy!”).

Interestingly, under the two-word condition but not under the one word condition, children looked for a significantly longer time at the (from an adult's perspective) unexpected outcome with only one object on the stage. Apparently, children took the formally different linguistic utterances to be hints indicating the nonidentity of the objects. This behaviour was observed also when the accompanying words were not the "correct" kind labels: even if different nonsense words (e.g., "Look, a fendle!" versus "Look, a toma!") were presented, the children seemed to use this as evidence that they were confronted with two different objects. Thus, what these studies seem to show is that even nine-month-olds are able to individuate two kind-different objects (in the absence of spatiotemporal information) if only the two different kinds of objects are accompanied by two different linguistic sound patterns⁹. In the face of these results Xu comes up with the hypothesis that linguistic acts play a causal role in the evolution of kind representations in the child's mind.

To recapitulate, let me try to give a summary of this process (that, admittedly, leaves open many questions): from early on, the child is equipped with knowledge about the general kind 'object'. This sort of knowledge enables the child to identify the objects in its environment. Knowledge about specific kinds (basic-level categories, cf. Rosch et al. 1976) is acquired later on when the child recognises that, often, the objects it perceives are not accidentally accompanied by certain words. At this time, the child stores pairs of simultaneously perceived object gestalts and word forms. It notices that object gestalts and word forms are not arbitrarily correlated. Comparing the series of object percepts accumulated under one word form with the series accumulated under a different word form, the child realises that the respective series are tied together not only by the word form: also the object percepts of a single series show a certain similarity with respect their gestalt¹⁰. To give a trivial example, all the objects registered under the word *ball* have the property of being round – in contrast to all the objects registered under the word *duck* (which are not round, but "duck-shaped").

As soon as the child realises that the object appearances accumulated under a certain word are perceptually similar with respect to each other (relative to the object appearances accumulated under another word), the identified similarity becomes the standard of comparison for the child's own categorisations. From

⁹When nonlinguistic sounds accompanied the object appearances, the effect did not occur. Neither with sounds like, for instance, "Look, [car alarm sound]" versus "Look, [spaceship sound]", nor with emotional articulations such as [positive:] "Ah!" versus [negative:] "Ewy!" (cf. Xu 2005:26).

¹⁰"Similarity" among objects under some concept can be measured in terms of relative distance within a conceptual space (cf. Gärdenfors 2000).

now on, the child expects to hear the respective word in connection to an object percept which is “similar enough” to the standard of comparison associated with the word. That is to say, the child has formed a hypothesis about possible objects suited to be accompanied by the word form. Often, the child will make experiences supporting the hypothesis. Often enough, however, it will be confronted with word-object-pairs that enforce an adjustment of the hypothesis. One could also say that, from the particular objects a word is *in fact* being used with, the child induces a theory of the *potential* usages of the word. In the face of a series of, for instance, balltokens, the child establishes the mental representation (=concept¹¹) of a balltype.

2.6 Speculating about developmental stages

According to the interpretations of Xu and her colleagues, the experiments show that, at different developmental stages, the child has qualitatively different kinds of knowledge about its environment, specifically about the relationship between objects and language. In the present section I allow myself to speculate about at least three such developmental stages, somewhat exceeding Xu’s own interpretation (compare footnote 13). Each stage is characterised by the child having a particular kind of “look at the context” within which it encounters real-world objects.

Stage 1: Already during the early first stage, the child is registering object appearances together with “accompanying appearances”. Let us in this connection speak of the “scenery” of an object appearance. The scenery of an object appearance includes other (at the given moment less relevant) objects existing at the same time and almost at the same place. The scenery furthermore includes accompanying smells, sounds and the like, and, therefore, also accompanying linguistic utterances. From the perspective of the child at this first stage, there is no regularity at all between the perceived object appearance and the elements of the scenery.

Stage 2: The second stage begins as soon as the child makes a certain discovery: while the form of other elements of the scenery seems to be independent of the kind of perceived object, the form of accompanying linguistic utterances is not. Strangely, it all too often happens to be the case that certain words accompany certain objects – too often to be by chance! The child concludes that there must be some correlation between word gestalts and object gestalts. This conclusion demarcates the step into the second developmental stage.

Stage 3: Once sensitive to the linguistic companions of an object appearance,

¹¹“A concept is the accumulated knowledge about a type of thing in the world” (Barsalou 2000)

the child starts to systematically collect (memorise) pairs of word and object gestalts. One might want to say that the child consciously investigates the conditions of use of linguistic forms. After enough object percepts in actual use have been collected under a certain linguistic form, the child goes on to form a hypothesis about possible usages of the form. Putting forward hypotheses about potential word usages is the characteristic of the third developmental stage.

Thus, while at the second stage the child knows *that* words are symbols, at the third stage it has knowledge of *what* a word is a symbol for. This “kind knowledge” results from abstraction across the different object percepts that have been accumulated under a single word form during the second stage. This way, every word ends up being a symbol for a kind of object. Successively, the number of kind symbols increases. The totality of kind symbols forms the nominal vocabulary of the child with every common noun being a symbol for a kind. In other words: if, by virtue of its properties, an object satisfies the conditions set by the kind, the respective common noun can be used in order to linguistically refer to the object.

It is well known that the first words which are acquired designate perceptually defined kinds (basic-level categories, cf. Rosch et al. 1976). Given the picture sketched above, this is not by chance: the first objects identified by a child are *physical objects*, after all. Recall from section 4 that the very first objects are identified as instances of the kind ‘physical object’:

A physical object is defined as any three-dimensional, bounded entity that moves on a spatiotemporally continuous path (Spelke, 1990; see also Jackendoff, 1983). The English word ‘object’ has multiple senses; this concept of physical object corresponds to one of the senses.
(Xu 1997:369-370)

Thus, the first words learned by the child sort the world of physical objects into classes of objects corresponding to basic-level categories. From the fact that sorting objects into *different* classes is the very purpose of establishing basic-level categories follows an important feature: basic-level kinds are equipollent. This means that each basic-level kind has at least one exclusive property, i.e. one property that no other basic-level category has. It is this feature, that basic-level kinds are equipollent, that allows for the inference of the existence of two different objects from the presence of two different basic-level kinds. Moreover, this feature is the reason that it is unexpected that an object transforms between basic-level categories¹².

Against this background, how can the results of the studies reported above be interpreted? As soon as the child has basic-level kinds at its disposal, it will

¹²Even though this happens in nature every now and then: from a caterpillar to a butterfly, from a tadpole to a frog, etc.

infer the existence of two distinct objects whenever it is successively confronted with two object appearances belonging to two different basic-level kinds. An object which carries all of the properties making up a ball, but which carried all of the properties making up a cup beforehand, does not fit into the child's metaphysics, so to speak. As a consequence, the child will be surprised when encountering such a "thing". In the described studies, this manifests itself in the longer looking time attested for twelve-month-old children. Prior to this developmental stage, when the child has not yet formed hypotheses about potential cups, balls, ducks etc., the child will, accordingly, not be surprised about such metamorphoses. This explains the observed behaviour of ten-month-old children. It was demonstrated that in certain cases already nine-month-olds *are* surprised by metamorphoses from one kind to another, even though they can not yet have mentally represented basic-level kinds. This was observed when the respective two kinds of objects were presented together with two different linguistic sound patterns (involving two different nouns). A possible explanation would be that in these cases, the child does not know the meaning of the nouns, i.e. the symbolised kinds, but does know that different words stand for different kinds¹³. This piece of information leads them to expect two different kinds of objects (whatever their exact content) and, hence, two different objects.

2.7 Conclusions I: kinds as sortal concepts

Against the background of the fact that linguistic expressions may basically denote objects or kinds, I reported on different experimental studies conducted by developmental psychologists. The immediate purpose was to demonstrate the genesis of kinds, the main purpose was to determine the particular notion of kind that I presuppose in the present thesis. Let me summarise my position which can be characterised by the slogan "kinds are concepts" (cf. Krifka 1995).

The real world against the background of which linguistic patterns evolve is of concrete, spatiotemporal nature. Think of it as a constellation of objects undergoing permanent changes in time. It is possible to distinguish between different types of changes: a new object *comes into existence*, for example,

¹³This corresponds to the second developmental stage within the speculative "3-step-model" sketched above. In contrast to that, Xu prefers the view according to which children know right from the beginning that words are symbols for kinds. That is to say, from her point of view it is not necessary to draw the distinction that I drew between the first and the second stage: "I am inclined to endow the human infant with certain expectations about words at the beginning of language acquisition such that in a "dubbing event" (e.g., "That is a dog!"), a word (in this case "dog") serves as a pointer to a particular sortal-kind in the environment" (Xu 2005:32-33).

when I bake a cake; an object *changes a property*, for example, when I colour a fence; an object *changes in relation* to another object, for example, when I take a cup out of a cupboard; an object *ceases to exist*, for example, when I eat up an apple. Right from the beginning of our own existence, we perceive the objects in our environment qua objects, i.e. as instances of the general kind ‘physical object’. Later on, at the age of about one year, we learn to sort (physical) objects into kinds of objects, following the linguistic practice of the social group within which we grow up. This goes together with the development of our first nominal lexicon consisting of linguistic symbols for basic-level kinds: by observing the way a particular common noun is used we establish for it a mental representation (concept). This process of inducing a concept presupposes sufficient knowledge about the extension of the noun. The necessary amount of knowledge is, roughly speaking, accumulated by again and again encountering actual pairs of linguistic utterances and constellations of objects (cf. Bartsch 1998). By successively establishing more and more kind representations, a new domain besides the real object domain evolves within our commonsense ontology, viz. the abstract and in a sense “unreal” kind domain.

The acquisition of nominals standing for basic-level kinds such as ‘dog’, ‘chair’, ‘tree’ etc. sets the stage for the acquisition of nominals standing for upper-level kinds such as ‘mammal’, ‘artifact’, ‘plant’, etc. (an upper-level kind covers the union of the sets of objects covered by different basic-level kinds), and likewise for the acquisition of nominals standing for lower-level kinds such as ‘poodle’, ‘armchair’, ‘apple tree’ etc. (a lower-level kind covers a subset of the objects covered by a basic-level kind). This way a nominal lexicon unfolds consisting of symbols for kinds which form, by virtue of the properties they represent, taxonomies. This is the reason why semanticists such as, for instance, Dayal 2004 call the kind domain “taxonomic domain”.

To recapitulate: Confronted with the real world of objects, human beings develop during their cognitive genesis the mental “world” of kinds. The reason why an English noun phrase like *the dog* or *an Apache* may either refer concretely, i.e. to a concrete thing as in (1), or generically, i.e. to an abstract thing-as-such as in (2), is that – as a matter of principle – elements of both of these two “worlds” may be chosen as the topic of linguistic utterances.

In section 2 I said that sentences are uttered in order to communicate states of affairs and that states of affairs are “constellations of things in the world”. This statement needs to be qualified now. For one thing, a sentence may be uttered to express a certain relation among objects of the real world (4a). Another sentence may be uttered to express a certain relation among kinds (4b). However, as we will see, there are also sentences like (4c) expressing a certain relation between a kind and an object, and, at least, sentences like (4d) expressing a relation between an object and a kind.

- (4) a. *The potato is lying on the table.*
 b. *The potato is a relative of the tomato.*
 c. *The potato reached Europe (a long time ago).*
 d. *The potato (over here) looks like a face.*

2.8 Conclusions II: what does it mean for a kind to exist?

Back to the question raised in section 2: what does it mean if someone asserts a certain kind to exist?

- (5) *The platypus actually exists.*

The answer at hand is: a sentence like (5) expresses the proposition that in the world there exists the kind ‘platypus’. This answer is premature, however, because it is left open what “in the world” is supposed to mean. We saw that the ontology underlying linguistic structures is systematically split into a real spatiotemporal object domain on the one side and an abstract conceptual kind domain on the other side. Given the distinction between an object world and a kind “world”, however, the adverbial *in the world* allows for at least three interpretations: (i) *in the world* relates to the object world, (ii) *in the world* relates to the kind world, (iii) *in the world* relates to the union of object world and kind world. A closer look reveals that (i) is no real option. By definition, a kind cannot exist in the object domain. Possibility (ii) is, at first sight, plausible. It faces a problem, however, when it comes to negated existentials:

- (6) *The platypus does not exist.*

Let us assume that (6) expresses, in line with (ii), that within the kind domain there exists no kind ‘platypus’. Let us furthermore assume (contra to fact) that the sentence is true. Given these assumptions the subject noun phrase cannot be analysed as referring to the kind ‘platypus’, simply because there is no such kind to refer to. But then, what is it that the subject of (6) refers to? What is the topic if someone utters (6)? Exactly the same problem arises with respect to possibility (iii): how can something that is not accessible as a possible referent (because it does not exist in the ontology) become the topic of an utterance? How can the speaker refer to something that nowhere exists?

In order to find the only plausible answer to the latter question, we have to see through the special content of the existence predicate. To anticipate, what a sentence like (5) expresses is the proposition that *in the object world there exist instances of the kind ‘platypus’*. Thus, the existence predicate is a predicate that selects for a kind term ascribing to the denoted kind the property of having object instances in the real world. In this respect, I subscribe to the

view advocated by McNally (1997). Let me briefly summarise her position. McNally puts the peculiarity of the existence predicate as follows:

The existential predicate in English is interpreted as a property of a *description of an entity*, specifically the property that the description is instantiated by some entity at some index. (McNally 1997:4)

In a subsequent passage we are told what “description of an entity” is supposed to mean:

... description of an entity, by which I intend an object that corresponds most closely to Chierchia and Turner’s notion of a *nominalized function* (Chierchia and Turner 1988; it is equivalent to Chierchia’s 1984 notion of the *entity correlate of a property*) (McNally 1997:4)

For Chierchia (and Turner), the “entity correlate of a property” is the abstract individual resulting from a so-called nominalising function applying to an ordinary property, i.e. to a property function. Being abstract entities (i.e. entities not localised in space and time), entity correlates of properties, like kinds, are different than real objects which are spatiotemporally localised by nature. Moreover, also like kinds, they are related in a systematic way to the elements populating the object domain. As for every real existing object it can be specified whether or not it instantiates a given kind, so for every object it can be specified whether or not it has a given (entity correlate of a) property.

The notion of a kind is more restrictive than the notion of an entity correlate of a property, however. According to Chierchia (1984), each of the subjects in (7) designates an entity correlate of a property. Only in the case of (7c), however, does it make sense to speak of kind reference. Therefore, the set of kinds is best viewed as a subset of the set of entity correlates of properties¹⁴:

- (7) a. *Red is a colour.*
- b. *Smoking is dangerous.*
- c. *Whales are mammals.*

For McNally (1997), the existence predicate, selecting for an expression denoting an entity correlate of a property, ascribes to this abstract entity the property of having instances. The peculiarity of the existence predicate lies in the fact that it establishes a relation among entities of two different ontological levels. Exactly this is what solves the problem that arises in connection with negated existentials. Uttering (6), the speaker chooses the entity correlate of the property of being a platypus to be the topic of her message and asserts the nonexistence

¹⁴A second subset is formed by the set of action types as exemplified by the referent of the subject of (7b) (cf. Chierchia 1998); see also chapter 3.

of objects in the real world (“at some index”) having this property of being a platypus.

I subscribe to the essence of McNally’s proposal. However, there is reason to assume that the existence predicate does not select for (entity correlates of) properties, but rather for kinds. At least, this could explain why (8c) is acceptable, but not (8a) and (8b):

- (8) a. **Red exists.*
- b. **Smoking exists.*
- c. *Whales exist.*

We therefore note that sentence (6) expresses the proposition that the kind ‘platypus’ is lacking object instances. Uttering this sentence, the speaker makes the kind ‘platypus’ the topic and asserts the nonexistence of objects in space and time that would instantiate this kind. Uttering (5), by contrast, the speaker asserts that there are entities in the real world of objects instantiating a certain entity of the “conceptual world” of kinds, viz. the sortal concept ‘platypus’. And so: What does it mean for a kind (sortal concept) to exist? For a kind (sortal concept) to exist means to have object instances in the real world.

Chapter 3

Kinds as ontological primitives – from a semantic perspective

3.1 Overview

From a referential-semantic point of view, the meanings of natural language expressions correspond to and, therefore, are always translatable into denotations. Denotations are, roughly speaking, those aspects of the world that a linguistic expression “talks about”. This implies that, directly or indirectly, the way the substance making up the world is structured manifests itself in linguistic structures. In other words, linguistic categories are expected to reflect, at least partly, ontological categories. For this reason, there can be no referential-semantic theory without some underlying ontological theory or, to put it differently, by proposing a referential-semantic analysis, the linguist automatically commits herself to certain ontological assumptions. Actually, the connection between ontological and semantic categories is trivial: if you assume that the meaning of the proper name *Maria* corresponds to the object individual Maria, you commit yourself to the assumption that the ontology includes a domain of object individuals containing an element named Maria.

The aim of the present chapter is to discuss the nature of kinds from a referential semantic viewpoint. I will focus on three aspects: (i) on the taxonomic organisation of the ontological kind domain (sections 2 and 3); (ii) on the relationship between kinds and concepts (section 4); (iii) on the relationship between kinds and properties (sections 5 and 6). I will conclude this chapter with a general remark concerning kinds as the object of scientific investigation and characterise kinds as the “bridgehead” between cognitive psychological and linguistic research.

3.2 Object reference and kind reference

Carlson (1977) turned our attention to the fact that any adequate referential-semantic analysis of such linguistic expressions like *be extinct*, *be widespread*, *be rare*, *be numerous* etc. requires the ontology to include a kind domain. I follow Carlson in considering kinds to be ontological primitives. Thus, besides an ontological domain of objects (containing elements like, for instance, Maria) there is an ontological domain of kinds¹. Zamparelli (1998) and Dayal (2004) share this position and propose that every English common noun is systematically ambiguous. It may determine its denotation either within the domain of objects or within the domain of kinds. Let us have a closer look at how.

With respect to object denotation, it has been argued (Chierchia 1998) that singular noun phrases (NPs) require their referents to be singularities while plural-NPs require their referents to be pluralities. Let *f*, *b* and *s* be the three (dog-)objects Fido, Barky and Spotty. Relative to this little universe, the singular noun *dog* denotes within the realm of singularities (atoms), represented by line 1 of the following figure:




Figure 3.1: dog objects

Pluralisation is viewed as the morphological operation that maps the denotation of the expression it applies to from sets of singularities onto sets of pluralities. Accordingly, the plural noun *dogs* denotes within the realm of pluralities, indicated by the lines 2 and 3 of figure 3.1.

This being said, let us now turn to kind denoting nouns. Dayal (2004) speaks about the ontological domain of kinds as the “taxonomic domain”. The reason for this terminology is that the kind domain is internally structured by the subkind-of relation leading to a taxonomic architecture (cf. Krifka et al. 1995:76). Let *F*, *B* and *S* be the three (whale-)kinds ‘fin whale’, ‘blue whale’ and ‘sperm whale’, let *W* be the kind ‘whale’, *D* be the kind ‘dog’ and *A* be the kind ‘animal’. These six kinds form, given the conditions of the actual world,

¹In addition, Carlson’s (1977) ontology contains a domain of stages.




Figure 3.2: taxonomy

the taxonomy given in figure 3.2². Assuming that the world relative to which the truth of a sentence is evaluated entails an ontological kind domain which is organised as in figure 3.2, the following statements count as true:

- (1)
 - a. *The kind 'fin whale' is a whale kind.*
 - b. *The kind 'sperm whale' is a whale kind.*
 - c. *The kind 'blue whale' is a whale kind.*
 - d. *The kind 'fin whale' is an animal kind.*
 - e. *The kind 'sperm whale' is an animal kind.*
 - f. *The kind 'blue whale' is an animal kind.*
 - g. *The kind 'whale' is an animal kind.*
 - h. *The kind 'dog' is an animal kind.*

In other words: the taxonomic structure given in figure 3.2 represents the conditions that must be satisfied for the sentences under (1) to be true. Alternatively, we can formulate the truth conditions of the sentences as follows, using **T** as a symbol for the subkind-of relation (Krifka et al. 1995:77):

- (2)
 - a. $[[[1a]]]=1$ iff there is a relation **T**(F,W).
 - b. $[[[1b]]]=1$ iff there is a relation **T**(S,W).
 - c. $[[[1c]]]=1$ iff there is a relation **T**(B,W).
 - d. $[[[1d]]]=1$ iff there is a relation **T**(F,A).
 - e. $[[[1e]]]=1$ iff there is a relation **T**(S,A).
 - f. $[[[1f]]]=1$ iff there is a relation **T**(B,A).
 - g. $[[[1g]]]=1$ iff there is a relation **T**(W,A).

²Arrows symbolise subkind-of relations. Transitivity holds: if an arrow leads from F to W and another from W to A, then there is also one leading from F to A. For ease of exposition, those arrows which are computable on the basis of the principle of transitivity are not represented.

h. $[[[(1h)]]]=1$ iff there is a relation $\mathbf{T}(D,A)$.

The denotation of a kind-level singular noun covers the set of all those atomic kinds that satisfy the descriptive conditions of the noun. Following Katz & Zamparelli (2005), let us assume that the extension of a nominal expression N is the kind ‘N’ as well as all the subkinds of N. To give some examples: The extension of the singular noun *dog* is, with respect to the kind domain given in figure 3.2, only the element D. The extension of *whale* is the set of the elements W, F, B and S. The extension of the noun *animal* is the set of the kinds A, W, D, F, B and S³.

Which set of kinds is covered by a plural noun? Zamparelli’s and Dayal’s idea is that - in analogy to the situation at the object-level - morphological pluralisation converts sets of singularities into sets of pluralities. Whereas the denotation of the singular noun *whale* corresponds to line 1 of figure 3.3, the denotation of the plural noun *whales* corresponds to the lines 2, 3 and 4⁴.

4.	{F,B,S,W}	...					
3.	{F,B,S}	{F,B,W}	{F,S,W}	{B,S,W}	...		
2.	{F,B}	{F,S}	{B,S}	{F,W}	{B,W}	{S,W}	...
1.	F	B	S	W	...		

Figure 3.3: whale kinds

We note that the ontology underlying linguistic patterns involves a kind domain, the internal structure of which is describable by the subkind-of relation \mathbf{T} . Think of the architecture of the kind domain as an enormous taxonomy with the kind ‘object’ constituting the superkind⁵. Except for the superkind ‘object’, every element within the kind domain is necessarily a subkind of the kind ‘object’.

³Alternatively, one could assume that the kind-level extension of N contains all subkinds of N, but not N itself. This is Dayal’s position, even though in a footnote she mentions evidence supporting the inclusion of the superkind into the denotation (Dayal 2004:426).

⁴If we followed Dayal and excluded W from the denotation of the singular noun (cf. footnote 3), we would have to subtract all those pluralities involving W.

⁵Compare chapter 2 where I argue that kinds are sortal concepts, i.e. mental representations originally developed for the purpose of object individuation. This explains *why* the kind ‘object’ forms the supercategory of the taxonomic domain.

Within the taxonomic hierarchy transitivity holds:

$$(3) \quad [\mathbf{T}(X,Y) \ \& \ \mathbf{T}(Y,Z) \Rightarrow \mathbf{T}(X,Z)]$$

For example: If the kind ‘sperm whale’ is a subkind of the kind ‘whale’ and if the kind ‘whale’ is a subkind of the kind ‘animal’, then the kind ‘sperm whale’ is also a subkind of the kind ‘animal’.

Besides the kind domain the ontology includes an object domain. Objects can be conceived of as spatiotemporal manifestations or instances of kinds. (?) propose the relator **R** to formalise the relation between object-level and kind-level individuals. **R**(x,X) symbolises that the object x is an instance (=realisation) of the kind X⁶. The following relationship holds between **T** and **R** (cf. Krifka et al. 1995:77):

$$(4) \quad [\mathbf{R}(x,X) \ \& \ \mathbf{T}(X,Y) \Rightarrow \mathbf{R}(x,Y)]$$

For example: If the object individual called Moby Dick is an instance of the kind ‘sperm whale’ and if the kind ‘sperm whale’ is a subkind of the ‘whale’, it follows that the object Moby Dick is also an instance of the kind ‘whale’.

To summarise: The entities populating the ontology come in at least two varieties, namely as kinds (or types) and as objects (or tokens). Kinds are organised by the subkind-of relation forming taxonomic hierarchies. The realisation relation **R** mediates between objects and kinds. Every object is a realisation, or instance, of some kind (at least of the kind ‘object’). However, not every kind has realisations among the objects⁷.

3.3 Kinds as taxonomic categories

Being a symbol for a kind, a common noun is a predicate characterising a set of kinds. This hybrid lexical-semantic status is due to the taxonomic architecture

⁶Here and elsewhere I will symbolise kinds by capital letters and objects by small letters.

⁷That a given kind need not have instances seems to be no matter of course. So, for example, Gatt (2004) bases his convincing analysis of the distribution of two possessive constructions in Maltese – according to Gatt, once the possessor-NP is interpreted at the object-level (“regular possession”) and once at the kind-level (“generic possession”) – on the assumption that every kind entails the existence of (at least one) object instance: $\forall x[\text{KIND}(x) \rightarrow \exists y[\text{OBJECT}(y) \ \& \ \mathbf{R}(y,x)]]$ (after Gatt 2004:207). Only the opposite implication is valid, however (cf. Dillig 1995:79): $\forall x[\text{OBJECT}(x) \rightarrow \exists y[\text{KIND}(y) \ \& \ \mathbf{R}(x,y)]]$. That there can be kinds without instances can easily be demonstrated by the truth of sentences like, for instance, *square circles do not exist*.

of the kind domain. Let us have a closer look at the way the taxonomic domain is organised.

Every taxonomic classification can be viewed as a default inheritance system (Corbett & Fraser 1999:57, see also Corbett & Fraser 1993). Within default inheritance systems, each semantic piece of information (which one can conceive of as an attribute-value pair) associated with a given category is automatically passed on to every one of the subcategories of the respective category. However, if a subcategory possesses some idiosyncratic value for some inherited attribute, the inherited value will be overwritten by the idiosyncratic one. From this it follows that, by naming a certain kind, every one of the subkinds of the named kind is at the same time characterised. Let us look at an example assuming that ‘European swan’ and ‘Australian swan’ are two subkinds of the kind ‘swan’. For ‘swan’ the attribute COLOUR is specified by the value white. This specification inherits by default to all subkinds. However, for the kind ‘Australian swan’ the attribute COLOUR is idiosyncratically valued black:

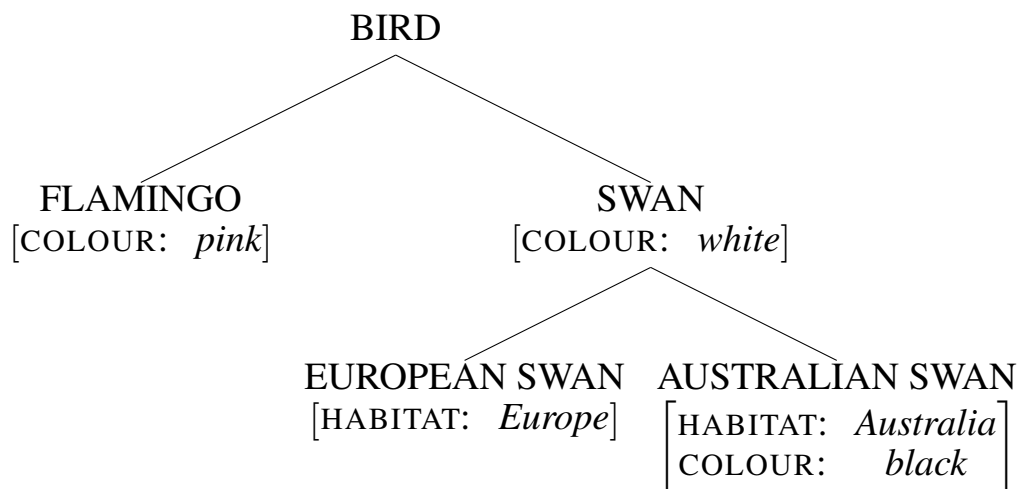


Figure 3.4: Default inheritance

Because the more specific information wins the following sentences are true relative to this taxonomic domain:

- (5) a. *Swans are white.*
 b. *In Australia swans are black.*

From the fact that a kind/category inherits every feature, i.e. attribute-value pair, from its superordinate kind (possibly overwritten by idiosyncratic values) it follows that the descriptive content of a kind always adequately describes the subkind. That is to say a subkind possesses every feature of its superkind plus at least one additional feature. The kind ‘swan’, for instance, possesses every feature of the kind ‘bird’ plus some idiosyncratic swan features. The kind

‘Australian swan’ possesses every feature of the kind ‘swan’ (and, therefore, of the kind ‘bird’) plus some idiosyncratic Australian swan features. Among these additional idiosyncratic features there may be features competing with inherited features. In this case, the idiosyncratic feature wins.

Now, if it holds that the description of a kind is always also a description of the subkinds, then the following two semantic representations turn out to be informationally identical. (10a) represents the kind-referring dimension of a common noun like, for instance, *dog* and (10b) represents its predicating dimension as a kind-characterising function⁸:

- (6) a. $[[dog]] = DOG \in D_k$
 b. $[[dog]] = \lambda X \in D_k. DOG(X)$

To sum up: Due to the taxonomic architecture of the kind domain, kind-level nominals have a hybrid lexical semantic nature. On the one hand they name a kind (10a), on the other hand they characterise a set of kinds, viz. the set of all subkinds of the named kind plus the named kind itself (10b).

3.4 Kinds and concepts

To say that a common noun is a predicate⁹ is to say that its function is to characterise a set of (kind-level or object-level) entities. Which specific set of entities it is that is characterised is determined by the common noun’s descriptive content, also called “sense”, “lexical content”, “concept” or simply “description”:

When you were asked to explain what the word *dog* means you would probably say that dogs are a certain kind of medium-sized animals with four legs and a tail, that they are often kept as pets, that they bark, that they bite etc. In other words, you will most likely give a description of dogs. This is an adequate reaction: giving a description of dogs may well count as an explanation of the meaning of *dog*. At least roughly, the meaning of such words may safely be regarded as a description of the kind of thing the word can be used for. (Lü₂ 2002:20)

We note that every common noun is associated with a description of a “kind of thing”. For certain entities this kind description is a correct description, for

⁸“ D_k ” stands for the kind domain.

⁹Whether a common noun is a derived predicate is not important at the moment. For ?) common nouns are primarily names of kinds, i.e. argument expressions, from which predicate expressions can be derived if necessary. ?) believes that different languages differ as to whether their nouns are primarily argument expressions or predicates.

others it is not. Entities correctly described by the kind description are potential referents of a noun phrase formed on the basis of the respective common noun. In this connection, it is often said that an entity “falls under the concept”. This mode of speaking implies that the descriptions at issue are concepts. This is also Lič₂’s opinion, because he goes on to write:

The meaning of the word *dog* [...] is a mental description. For mental descriptions in general, the term concept will be used. A concept for a kind, or category, of entities is information in the mind that allows us to discriminate entities of that kind from entities of other kinds.
(Lič₂ 2002:20)

It is easy to see that this is exactly the notion of a sortal concept as discussed in chapter 2. Note the somewhat irritating situation that arises when kinds are taken to be ontological primitives, i.e. entities, side by side with objects. The entities possibly falling under a sortal concept may be objects or may be kinds. If a certain kind falls under a certain concept, this must mean that a kind is a correct description of a kind. Replace “thing” by “kind” in the last sentence of the first Lič₂ quotation above:

The meaning of such words may safely be regarded as a description of a kind of kind the word can be used for.

If this seems paradoxical to you, as it seemed to me, convince yourself that the “kind of kind” is nothing but the superkind within a taxonomic hierarchy. The kind ‘whale’, for instance, is the kind of the kind ‘blue whale’, of the kind ‘sperm whale’, of the kind ‘dolphin’ etc. We see that treating the meaning of the common noun *whale* as a predicate characterising the set of subkinds of the kind ‘whale’ including the kind ‘whale’ itself is the referential semantic version of Lič₂’s position according to which the meaning of *whale* is the conceptual description of the whale kind. Compare also Barsalou’s definition in this connection:

A concept is the accumulated knowledge about a type [=kind, OMR] of thing in the world. (Barsalou 2000)

Strictly speaking, kinds and conceptual descriptions of kinds are thus different creatures. One might want to say that kinds are *reifications* of conceptual descriptions¹⁰. The relationship between kinds and concepts may be represented as in figure 3.5.

¹⁰“Reification, also called hypostatization, is treating an abstract concept as if it were a real, concrete thing” (Wikipedia)




Figure 3.5: kinds as reified concepts

The important thing to note is that even though kinds and concepts are strictly speaking different notions, they encode the same piece of information. This reminds of the Fregean idea according to which properties have a double nature which was incorporated into formal semantics by ?) and ?). The next section is dedicated to this topic.

3.5 Kinds and properties

Adjectives are property words par excellence. Pre-theoretically speaking, the meaning of the adjective *red*, for instance, is the property ‘red’; the meaning of *round* is the property ‘round’; the meaning of *cool* is the property ‘cool’ etc. To give some substance to this not very impressive semantic assumption, it makes sense to take the pragmatics of property words into consideration. An adjective designates what it designates, a property, in order to be used to characterise a set of entities. This is the referential effect of the property. From this perspective, we can think of the property ‘red’ as some sort of mental pattern determining the standard of comparison relative to which it is decided whether or not a given entity has the property of being red. Comparing the real colour of the entity with the pattern red one will either conclude that the entity is red or that it is not red.

This referential-semantic perspective is adopted by formal semanticists who consider properties to be (mathematical) functions. In the extensional version, the property ‘red’, for example, is a function taking entities as input and delivering one out of the two truth values true or false as output. Other adjectives like *round*, *cool* etc. are treated accordingly¹¹. This formal semantic approach has interesting consequences. If the linguistic meaning of an expression is identified with its referential effect, then there is virtually no difference in meaning between adjectives, (non-relational) nouns and (non-transitive) verbs. Traditionally, adjectives are classified as property words, nouns as object words and verbs as action words. From a formal semantic point of view, there is no reason to sort these expressions into different classes¹².

¹¹This practice does not extend to all adjectives, however. Exemplars like *apparent* or *former* require a different approach (cf. Partee 1995, Heim & Kratzer 1998).

¹²I abstract away from Davidsonian arguments in the argument structure of verbs.

The meaning of a noun like *dog* and the meaning of a verb like *walk* can, like the meaning of an adjective, be viewed as a pattern determining whether a given entity is a dog or whether it is walking. There indeed seems to be no reason not to treat (one-place) adjectives, nouns and verbs on a par: the meaning is always (a function from situations into) a function from entities into truth values. Using semantic types this means that adjectives like *red*, nouns like *dog* and verbs like *walk* are of type $\langle e,t \rangle$ (extensional) or $\langle s \langle e,t \rangle \rangle$ (intensional), respectively¹³.

However, the described formal semantic practice to model properties as functions runs into a fundamental problem for which Chierchia (1984) and Chierchia & Turner (1988) developed a technical solution. Let me briefly recapitulate the problem using own examples:

- (7) a. *Chomsky is cool.*
 b. *Cool is cool.*

The grammatical subject of (7a) denotes an object entity and is, accordingly, an expression of semantic type $\langle e \rangle$. The grammatical predicate of (7a) denotes a property and is of semantic type $\langle e,t \rangle$. Semantic composition succeeds: the subject expression can semantically saturate the argument slot of the predicate expression. The result is the sentence given in (7a) which denotes a proposition, i.e. which is of semantic type $\langle t \rangle$. The problem arises in connection with sentence (7b). What is in this case the semantic type of the grammatical subject? Note that the subject of (7b) and the predicates of (7a) and (7b) are all of the same form, yet they cannot be of the same meaning (of the same semantic type). Were they of type $\langle e,t \rangle$, the semantic composition of (7b) would fail. What could be a solution to this problem? If $\langle e,t \rangle$ was the semantic type of the grammatical predicate in (7b), the subject would have to be either $\langle e \rangle$ or $\langle \langle e,t \rangle, t \rangle$. If $\langle e,t \rangle$ was the semantic type of the grammatical subject, the predicate would have to be of semantic type $\langle \langle e,t \rangle, t \rangle$.

Chierchia and Turner's solution is to consider the subject of (7b) to be of semantic type $\langle e \rangle$, albeit of a special ontological sort. The idea is that the ontological domain of entities is partitioned into the sort of ordinary entities (to which, for example, the entity referred to by the subject of (7a) belongs) and into another sort of entities named "nominalised functions" by Chierchia and Turner. This is a telling name: a nominalised function is conceived of as an image of a property function. According to Chierchia and Turner, every property function has such an image in the entity domain, resulting from the application of a so-called nominalisation operation¹⁴. An example for a nominalised function is

¹³"e" stands for entity, "s" for situation and "t" for truth value.

¹⁴Be careful not to get confused: the term "nominalised function" is a little misleading because a nominalised function is not a kind of function, as the term seems to suggest,

the referent of the subject expression of (7b). Thus, for every property function (semantic type $\langle e, t \rangle$) there is a property entity (semantic type $\langle e \rangle$).

The class of expressions of semantic type $\langle e \rangle$ includes (with “o” indicating ordinary objects and “nf” indicating nominalised functions) expressions of type $\langle e_o \rangle$ and expressions of type $\langle e_{nf} \rangle$. Formally, there are two operations defined mediating between the two property variants. The operator “ \cap ” that maps property functions onto property entities is called nominaliser (also: down-operator). Its counterpart “ \cup ”, called predicativiser (also: up-operator), maps property entities onto property functions:




Figure 3.6: the double nature of properties

To repeat, the sentence (7b) poses a problem for the formal (truth conditional) approach to semantics because it is impossible that in its two manifestations, the expression *cool* realises one and the same meaning. The problem disappears if property functions were systematically related to property entities (=nominalised functions). Assuming this, we would find within the denotation of the predicate *cool* not only object individuals like Chomsky, but also nominalised functions like ‘cool’.

Because within the formal semantic framework, not only adjectival meanings, but also verbal and nominal meanings are properties (see above), Chierchia and Turner’s proposal covers also the subject expressions in the following sentences:

- (8) a. *Sabre-toothed tigers are cool.*
 b. *Smoking is cool.*

Consider (8b) first. By means of \cap , a verb like *smoke*, basically a predicate (semantic type $\langle e, t \rangle$), can be converted into a term (semantic type $\langle e \rangle$). This term denotes the property entity ‘smoke’. According to Chierchia and Turner (1988:294-295), in many languages this operation manifests itself in form of infinitive morphology. By the same token it is to be expected that the nominal predicate in (8a) can be converted into a term designating the property entity ‘sabre-toothed tiger’.

There is a puzzle about (8a), however. Why does the subject show the plural form? As it stands, there is no explanation for the fact that the nominalised

but an entity.

function ‘sabre-toothed tiger’ cannot be referred to by the corresponding singular nominal:

(9) **Sabre-toothed tiger is cool.*

?) tries to solve this puzzle by viewing property entities denoted by nominals as special properties: kinds. This particular notion of a kind excludes the application of the nominaliser to property functions that characterise only singularities. As a consequence, bare singular nominals cannot be used to refer to kinds. Let us look at this in more detail.

3.6 Chierchia’s (1998) notion of kind

Chierchia (1998) is concerned with a subset of property entities, namely with those property entities that are denoted by nominals. These are called “kinds”. Thus, for Chierchia, a kind is the image of a nominal property within the entity domain. The image of a verbal property is called an “action type” by Chierchia. As I understand it, kinds and action types are species of property entities. Given this, the image of an adjectival property could be viewed as a property proper:




Figure 3.7: Property entities according to Chierchia

An example of a property proper term is the subject *cool* in (7b), while an example of an action type term is the subject *smoking* in (8b). The subject of (8a) exemplifies a kind term and so do the subjects in the following sentences:

- (10) a. *Blue whales are whales.*
 b. *Whales are mammals.*

If *whales* in (10a) and *mammals* in (10b) are both predicates of semantic type $\langle e,t \rangle$, *whales* in (10b) must be an argument expression of type $\langle e \rangle$ ¹⁵. According to Chierchia (1998:349), such an argument *whales* is derived from the

¹⁵Alternatively, one could assume that (10b) translates into a generic quantification structure (e.g., Greenberg 2002). In this case the subject nominal could be analysed as a predicate of type $\langle e,t \rangle$ restricting the quantificational domain.

predicate *whales* by means of the down-operator \sqcap . Chierchia assumes that kind terms like *whales* or *dogs* resemble proper names like *Moby Dick* or *Fido* in that both sorts of expression designate entities. Unlike a proper name, however, a kind term typically designates a spatially discontinuous entity. Both kind terms and proper names are semantically analysed as functions mapping situations onto entities:

It seems natural to identify a kind in any given world (or situation) with the totality of its instances. Thus, the dog-kind in our world can be identified with the totality of dogs, the scattered entity that comprises all dogs, or the fusion of all dogs around. In our framework, this entity is modeled by the set of dogs. This means that we can model kinds as individual concepts of a certain sort: functions from worlds (or situations) into pluralities, the sum of all instances of the kind. (Chierchia 1998:349)

Within Chierchia's (1998) theoretical system, the down-operator \sqcap takes a function from situations into sets of entities as input and delivers a function from situations into the maximum entity of the respective input set (with the latter entity being usually "scattered" in space). The input and the output of the down-operator are viewed as informationally equivalent:

[K]inds and [...] properties can in a way be seen as two modes of packaging the same information. Using Frege's (1891) metaphor, properties are unsaturated, something that (at a world) is true or false of individuals. Kinds are saturated, something that at a world has concrete, if possibly spatiotemporal discontinuous, manifestations. (Chierchia 1998:354)

The relationship between kinds and properties is visualised as follows by Chierchia (1998:349):




Figure 3.8: kinds and property functions

Importantly, the nominaliser \sqcap is not applicable to any property. A property function can be transformed into a kind only if the set of objects characterised by the property shows a "sufficiently regular behaviour". This requirement is fulfilled by natural kinds such as 'dog', 'lion' or 'tree', by artifacts such as

‘table’ or ‘car’, and also by “complex things” like ‘intelligent student’ or ‘spot of ink’. Unlikely is a kind ‘broken shoe that Leo left behind’ because broken shoes that Leo left behind are not known for showing a sufficiently regular behaviour:

NPs can be modified, and not every modified NP is going to be associated with a kind. This will depend on whether it picks a class of objects that display a sufficiently regular behavior. What counts as sufficiently regular is determined by the shared knowledge and beliefs in the community of speakers (and is thus subject to a certain degree of variation). (Chierchia 1998:372)

Showing a sufficiently regular behaviour is not a sufficient condition for kinds, however. In addition, a Chierchia kind must not be associated with an instantiation set whose cardinality is fixed to one:

Moreover, kinds [...] will generally have a plurality of instances (even though sometimes they may have just one or none). But something that is necessarily instantiated by just one individual [...] would *not* qualify as a kind. (Chierchia 1998:350)

Two important consequences follow from this latter restriction. First, proper names like *Moby Dick*, *Fido* etc. cannot be kind terms because such proper names designate singletons. Secondly, bare singular nominals cannot be kind terms. Why is this? Recall from section 2 that, according to Chierchia, only atomic objects are in the extension of singular nominals. If a singular nominal characterises more than one entity relative to a situation (like, for instance, *dog* relative to the situation in figure 3.1), there will be no unique maximum entity and the down-operator is not applicable. Consequently, bare singulars cannot be used as kind terms. If a singular nominal characterises only one (atomic) entity relative to a situation, as does for example *moon* relative to the actual world, then there is a maximum entity. Nevertheless, *moon* cannot be used as a kind term because of the stipulation that kinds must not be associated with only one object. So this is how Chierchia explains why bare singulars (in English) cannot be kind terms.

Krifka (2004) draws attention to the fact that every adequate theory should allow for kind terms associated with empty instantiation sets (e.g., denoting extinct kinds). It is not clear why kinds associated with singletons should be excluded but kinds associated with an empty set of objects should be acceptable.

Let me summarise: Because predicates like *rare*, *widespread*, *extinct* etc. call for a domain of kinds within the ontology, the question of the nature of kinds arises. It soon becomes clear that the notion of a kind is closely related to the

notion of a concept. Chierchia (1998) directly models kinds as concepts¹⁶ (note that Chierchia uses the term “individual” in the sense of object individual):

[W]e can model kinds as individual concepts of a certain sort: functions from worlds (or situations) into pluralities, the sum of all instances of the kind. (Chierchia 1998:349)

Chierchia’s proposal in a nutshell: A common noun in English may either denote a property function or an entity. If it denotes an entity it functions as a kind term. Abstracting away from all the technical details, Chierchia’s notion of a kind contrasts in one important respect from the notion of a kind that I take for granted in this dissertation. For Chierchia, kind terms and proper names are of the same semantic type, $\langle s, e \rangle$. The difference is only that the entity designated by the proper name is (usually) spatially continuous, whereas the entity designated by a kind term is (usually) spatially discontinuous. Thus, for both denotations it holds that they are (intensionalised) spatially localised entities. This sharply contrasts with my basic assumptions: kinds and objects differ from each other with respect to spatiotemporal localisation (see chapter 4). While objects are spatiotemporally localised entities (i.e. objects are tokens), kinds are not (because kinds are types). For me, kinds and objects differ as to *whether* they are spatiotemporally localised. For Chierchia, they differ as to *how* they are spati(otempor)ally localised.

3.7 Lexical and formal semantics

In the previous chapter I discussed the notion of a kind as a sortal concept from the point of view of cognitive psychology. In the present chapter I addressed the same topic from the point of view of semantics. Being ontological primitives, kinds are available as possible referents of linguistic expressions. It was argued that kinds should be viewed as reified sortal concepts forming a taxonomically organised ontological domain. As will be demonstrated later on in the chapters 5 to 7, this view proves useful to explain a number of otherwise puzzling linguistic phenomena. To conclude this chapter, I would like to emphasise the role that kinds play as, so to speak, bridgeheads between the conceptual system and the grammatical system, i.e. between psychology and linguistics.

For a formal semanticist, the meaning of a content word, e.g. of the noun *dog*, is a function characterising a set of entities relative to a situation. For a lexical-conceptual semanticist, the meaning of a content word is a mental pattern (con-

¹⁶...where “concept” is meant not mentalistically, but rather mathematically (cf. Löbner 2003:355). Compare in this connection: “[A] concept [...] is a function from possible worlds to individuals” (de Swart 1998:213)

cept) that may be used to characterise entities relative to a situation. As a referential theory, formal semantics is not at all interested in the internal structure of concepts. For a formal semanticist, this is not a linguistic but a psychological topic falling into the realm of cognitive psychology. A lexical-conceptual semanticist would disagree. She would say that as soon as a concept manifests itself as a symbol in language, investigating the structure of the concept/kind is the task of linguists.

Now, as discussed at length, speakers use language not only to refer to real, spatiotemporal entities (objects), but also to refer to abstract, conceptual entities (kinds). What are the implications for the two approaches to meaning? From the formal semantic point of view, when used at the kind-level, the meaning of the noun *dog* amounts to a function characterising a set of kinds. From the lexical-conceptual semantic point of view, lexical nouns are symbols for kinds. Are these two approaches compatible with each other? Of course they are. They differ from each other only with respect to the respective topic of investigation.

To see this, let us take on the more abstract perspective of a general scientist. For the general scientist, the topic of investigation is everything that exists in the real natural world. This includes, among many other things, people interacting with each other and with their environment. On closer inspection, it turns out that people are equipped with certain cognitive capacities enabling them to successfully navigate through the natural world, including the capacity to individuate objects. Object individuation operates on top of a system of kinds sorting the real world of objects¹⁷. This mentally represented conceptual system differs from person to person to a greater or lesser degree. If our general scientist specialises in investigating these cognitive capacities, she will be called a psychologist. Moreover, as a result of interacting with each other over large periods of time, various forms of communication developed within human populations. Among them is a very powerful form of communication: languaging. If our general scientist specialises in investigating the nature of language, she will be called a linguist. If she further concentrates on the meaning aspect of language, she will be a semanticist.

What I want to point out in this dissertation is that the system of kinds serves not only as the basis for object categorisation and individuation performances. It is also the lexical fundament on top of which the grammatical processes operate. Therefore, the realm of cognitive psychology and the realm of linguistics overlap when it comes to kinds. On the one hand, kinds are sortal concepts belonging to the general conceptual system. On the other hand, kinds are ontological primitives, i.e. possible referents of linguistic expressions.

The task of cognitive psychology is to investigate the structure of concepts.

¹⁷More precisely, *kind-based* object individuation (see chapter 2).

The task of lexical semantics is to investigate the structure of those concepts that linguistically manifest themselves as kind symbols in the mental lexicon of speakers. Reconstructing the taxonomy of kinds can be viewed as one goal of lexical semantics. The task of referential semantics is to investigate the regularities governing linguistic referencing to ontological entities. The goal is to answer questions like these: How are morphosyntactically complex expressions build up from basic lexical expressions so that they can be used to refer to objects, kinds, events, etc.? What sort of morphosyntactic complexation is suited for what sort of reference?

Chapter 4

Spatiotemporal localisation – type-token theories of genericity

4.1 Overview

Carlson's (1977) famous theory of genericity is based on a distinction between two sorts of ontological entities: stages and individuals. Whereas stages are conceived as spatiotemporally bounded entities, individuals are viewed as abstract entities which are not localised in space and time. Owing to the influence that Carlson's framework had upon the field of genericity research, one goal of the present chapter is to introduce the basic ingredients and the principal mechanism of Carlson's theory. More important, however, is something else: to work out that Carlson's (1977) particular theory is but one instantiation of a more general kind of theory. Like the Carlsonian version, every realisation of this general type of genericity theory is characterised by a distinction between spatiotemporal entities (tokens) on the one side and non-spatiotemporal entities (types) on the other. I will refer to these theories as type-token theories. Within type-token frameworks, generically used expressions are explained as denoting non-spatiotemporal entities, i.e. types.

To achieve this goal, I will introduce and compare different instances of type-token theories that have been proposed in the literature. In section 2 we will encounter two theories that work with a three-way ontological partition: On the one hand, there is Carlson's (1977) classic theory differentiating between spatiotemporal stages and non-spatiotemporal individuals, the latter being further subdivided into objects and kinds. On the other hand, there is its Russian counterpart Šmelev (1996)¹. Šmelev's theory bases upon the Carlsonian original

¹That I compare semantic theories from the Russian scientific sphere with theories from the angloamerican sphere is due to the fact that parts of this chapter base on a paper presented at the workshop "Formal Semantics in Moscow" (23 April 2005, Moscow State University MGU). One aim of the workshop was "to help strengthen bridges be-

nevertheless differing – not only terminologically – in important details. Beginning with a distinction between intrinsically non-spatiotemporal classes as opposed to individuals, Šmelev separates non-spatiotemporal individuals (“abstract individuals”) from spatiotemporal individuals (“instances”).

Section 3 is dedicated to theories that get by with a two-way partitioned ontology. Bulygina (1982) distinguishes spatiotemporally localised predicates from predicates which are not spatiotemporally localised and identifies, corresponding with the predicate distinction, two sorts of noun phrases (NPs): concretely referring noun phrases and non-concretely referring noun phrases. I will argue that it is possible to reinterpret Bulygina’s noun phrase distinction as the distinction between object-referring and kind-referring noun phrases in the sense of Krifka (1995). Finally, Kratzer’s (1989/1995) proposal will be discussed and compared to Bulygina’s account. Kratzer contrasts predicates with a spatiotemporal location argument and predicates lacking a spatiotemporal location argument.

In section 4, I draw attention to a terminological difficulty: it turns out that many semanticists use the term “referential” only in connection with object-level expressions, while others speak of “referentiality” also in connection with kind-level expressions.

Finally, I will outline a type-token-theory based on only two ontological classes. Given this is right, we expect to find among linguistic expressions token-referring terms and type-referring terms as well as token-level predicates and type-level predicates. Following Krifka (1995), I take tokens to be real objects and I consider types to be abstract kinds (sortal concepts, see chapter 2) – stages can thus be dispensed with. This minimal set-up talks of “nominal genericity” (cf. chapter 1) whenever an NP is interpreted not at the object-level but at the kind-level, and of “predicate genericity” whenever a predicate ascribes its property not to an object but to a kind.

4.2 Three-class systems

Carlson (1977) suggested a systematic two-way distinction of the ontology, mainly to explain the semantic behaviour of bare plural noun phrases in English. Every entity is supposed to exist in two ontological varieties, with each one being “a spatially and temporally bounded manifestation” (Carlson 1977:115) of the other. Spatiotemporally manifestations are called “stages” and contrast with “individuals”. Carlson characterises an individual as the spatiotemporally unbounded “whatever-it-is that ties a series of stages together to make them stages of the same thing” (Carlson 1977:115). In Carlson’s ontological system, one

tween “western” and “Russian” approaches to semantics” (B. Partee).

subtype of an individual is represented by the category of objects and one by the above mentioned category of kinds:

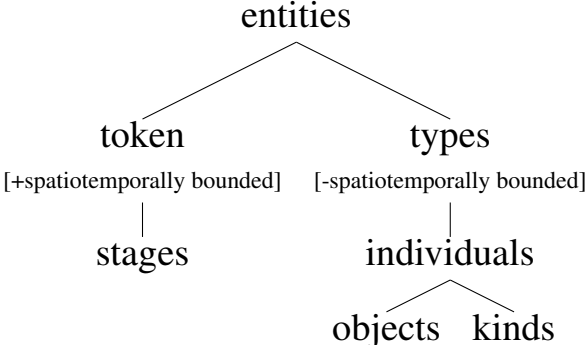


Figure 4.1: Carlson’s ontology

Concerned with the interpretation of Russian noun phrases, Šmelev (1996) follows Carlson assuming a three-partite structure of the entity domain. The crucial distinction is once again the contrast between a “concrete spatiotemporal manifestation of an object” and an “object by itself, abstracted away from concrete spatiotemporal manifestations” (Šmelev (1996:44-45). Kinds (aka klassy) are inherently abstract entities in this sense of abstractness. Figure 4.2 shows the ontological system by Šmelev:

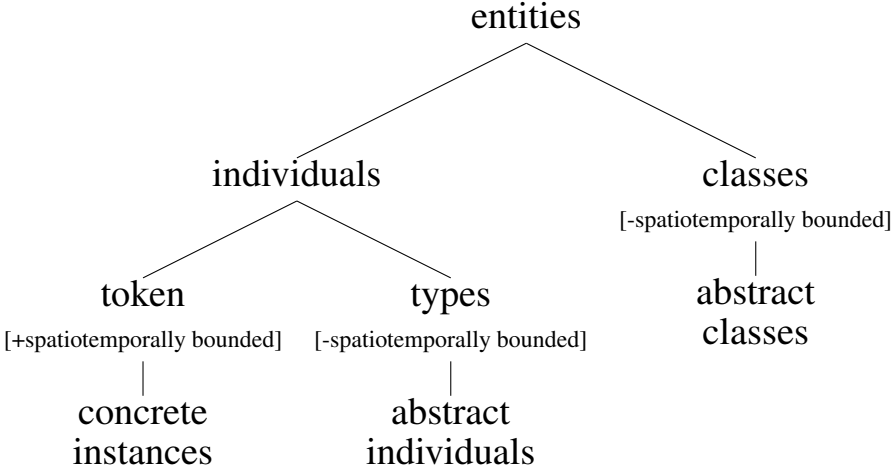


Figure 4.2: Šmelev’s ontology

Thus, Carlson as well as Šmelev assume that a noun phrase may in general be used to refer to three sorts of entities. This assumption, that the referent of a given noun phrase may belong to one out of three ontological classes, entails the prediction that the language possesses three classes of predicates. Consequently, Carlson (1977) claims that the English lexicon contains stage-level predicates

(a), object-level predicates (b) and kind-level predicates (c):

- (1) a. *be available, be hungry, be drunk, be dead, run, ...*
- b. *be altruistic, be intelligent, be a doctor, be a mammal, ...*
- c. *be widespread, be common, be rare, be extinct, ...*

The lexical meaning of a stage-level predicate is viewed as ascribing a property to some stage. By the same token, an object-level predicate ascribes a property to some object and a kind-level predicate ascribes a property to some kind. To enable the semantic composition to succeed, a noun phrase filling the syntactic argument slot of a stage-level predicate must deliver a stage, one filling the slot of an object-level predicate must deliver an object, and one filling the slot of a kind-level predicate must deliver a kind. Otherwise, the semantic composition would fail.

It is a characteristic feature of Carlson's (1977) system that there are no expressions which refer to stages by themselves, i.e. lexically. Instead, noun phrases contributing a stage always derive from a semantic operation applying covertly to an object term or to a kind term. Specifically, Carlson proposes two such covert "realisation operators" called R and R'. Applied whenever an object-level noun phrase is realised in the argument slot of a predicate selecting for stages, R changes the denotation of the noun phrase from the abstract object to some concrete spatiotemporal manifestation of it. In a similar way R' maps a kind term denotation onto some concrete spatiotemporal manifestation of the kind.

Example (2) illustrates the Carlson-style semantic composition of the sentences *Bill is running* and *Dogs are running*. Remember, following Carlson's conception, *Bill* is the name of the object Bill and *dogs* is the name of the kind DOG (x_s stands for an entity variable restricted to stages):

- (2) a. $[[\textit{Bill is running}]] = [\lambda x_s. \text{RUN}(x_s)](\text{R}(\text{Bill}))$
- b. $[[\textit{Dogs are running}]] = [\lambda x_s. \text{RUN}(x_s)](\text{R}'(\text{DOG}))$

Another feature of Carlson's theory is that (dynamic) verbs are viewed as lexically denoting at the stage-level. As a consequence, if used generically, the meaning of a verb like, for instance, *to smoke* must be shifted from the stage-level to the individual-level in order to compose with an object term or kind term. This is achieved by a covert operator called G, compare (3a,b). Apart from G, Carlson assumes another verb phrase operator G'. By shifting a lexical object-level predicate onto a kind-level predicate, G' enables its semantic composition with a kind term, compare (3c):

- (3) a. $[[\textit{Bill smokes}]] = \text{G}([\lambda x_s. \text{SMOKE}(x_s)])(\text{Bill})$
- b. $[[\textit{Sailors smoke}]] = \text{G}'([\lambda x_s. \text{SMOKE}(x_s)])(\text{SAILOR})$

- c. $[[\text{Dogs are four-legged}]] = G'([\lambda x_0. \text{FOUR-LEGGED}(x_0)])(\text{DOGS})$

With respect to the basic lexical format of a predicate, Šmelev takes on a different stand. A verb like *kurit'* is considered to be semantically underspecified with respect to whether it ascribes a property to some abstract entity (leading to a generic interpretation) or to some spatiotemporal manifestation of it (leading to an episodic interpretation). Thus, only by embedding it into a linguistic context, the meaning of the verb will be disambiguated (cf. Šmelev 1996:48). Šmelev points out, however, that a number of Russian predicates are fixed to one of the three predicted meaning categories. (4) shows examples of Russian predicates which, according to Šmelev, intrinsically assign properties to concrete instances (aka stages) (a), to abstract individuals (aka objects) (b) and to abstract classes (aka kinds) (c):

- (4) a. *byt' pjanym, byt' bolen, byt' rad, byt' v razdraženii, ...*
 'be drunk' 'be ill' 'be happy' 'be exited'
- b. *byt' pjanicej, byt' bol'noj, ljubit' Ivana, byt' umnyj, ...*
 'be a drinker' 'be chronically ill' 'love Ivan' 'be intelligent'
- c. *byt' redkost'ju, vymeret', polučit' bol'šoe razprostranenie, ...*
 'be rare' 'die out' 'spread out'

Croft (1986) criticises Carlson's system for involving two "asymmetries", noticed already by Carlson himself, that do not appear to be very attractive from a theoretical point of view. The first one concerns the nonexistence of lexical stage terms in English. While there are basic object terms (*Bill, this car over there, my uncle*) and basic kind terms (*dogs, the computer (as such), Hippopotamus Amphibius*), stage terms exist only virtually in the sense that they always result from the semantic computation of an object term or kind term within the argument slot of a stage-level predicate. The second asymmetry concerns the fact that every English object-level predicate can be used as a kind-level predicate without exception. As Carlson puts it: "whatever may be meaningfully predicated of an object may also be meaningfully said of a kind" (Carlson 1977:248). This regularity does not derive from principled grounds in Carlson's theory. Only by postulating the operator G', Carlson manages to "discover" object-level predicates as they are predicted by his three-partite system:

"[T]here appear to be no specifically object-level predicates in the language. In actuality, there are many, but every one of them may be applied to a kind via the G' operator" (Carlson 1977:280).

However, if every object-level predicate can really function as a kind-level predicate, why should a covert semantic operator like G' exist at all? Croft suspects "two many levels" in Carlson's theory and sketches a system that builds on only

two entity types (“closed-classes” versus “open-classes”) which would amount to only two respective predicate classes.

4.3 Two-class systems

Although Bulygina (1982) is not very explicit about the ontological structure of the presupposed entity domain, her approach can be viewed as a dual system similar to what is being suggested by Croft. Two classes of predicates are distinguished, one represented by predicates ascribing some accidental property, the other represented by predicates assigning some essential property to the referent of the term with which they syntactically combine. Bulygina describes the former as “temporally localised” predicates and the latter as not being temporally localised: Those eventualities described by temporally non-localised predicates are called “qualities” (kačestva), those described by temporally localised predicates are called “appearances” (javlenija). The class of predicates describing appearances (=episodic predicates) includes more fine grained lexical semantic categories such as predicates describing states, predicates describing processes, and predicates describing events (cf. Bulygina 1982)²:

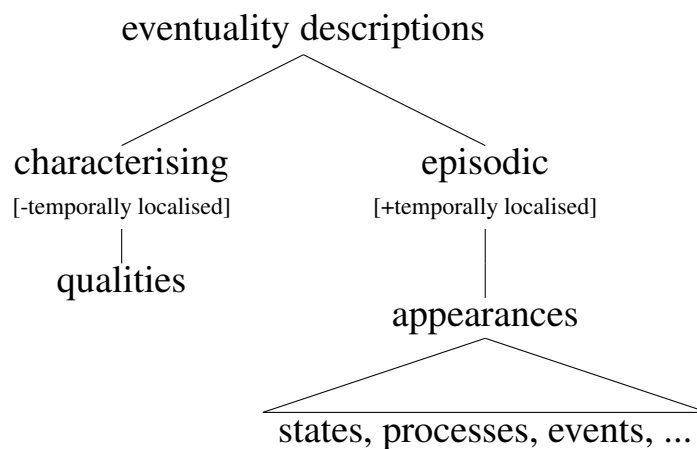


Figure 4.3: Bulygina’s predicate classification

With respect to noun phrase interpretations, Bulygina distinguishes between “referential” and “non-referential” noun phrases, the latter being sometimes also characterised as “non-specific”, “kind-referring” or “non-concretely referring” noun phrases. Bulygina furthermore coins the term “semantic agreement” to address the fact that the referential status of the subject argument expression partly

²For a similar distinction between verbal predicates describing “atemporal properties and relations” (vnevremennye svojstva i sootnošenija) and verbal predicates describing states, processes and events, see Padučeva 1996.

depends on whether the predicate is episodic or characterising (see also Mehlig 1983, Weiss 1983). An episodic predicate necessarily requires a concretely referring (“referential”) subject, while a characterising predicate allows for both a concretely referring and a non-concretely referring (“non-referential”) subject:

“Aktual’noe” upotreblenie predikata i nereferentnost’ imen vzaimno nesovmestimy, no “neaktual’nost” (vremennaja nelokalizovannost’) ne izključaet referentnosti imeni. (Bulygina 1982:29)

The “actualised” [episodic] use of a predicate is not compatible with a non-referential nominal argument, but “non-actuality” (being temporally not localised) [genericity] does not exclude the referentiality of the nominal. (own translation)

What sort of entity is it that concretely referring noun phrases or, respectively, non-concretely referring noun phrases refer to? I believe that Bulygina’s approach is perfectly compatible with the ontological distinction between “real objects” and “abstract entities” as advocated by Krifka (1995). Real objects are viewed by Krifka as specimens or, to use the term once again, spatiotemporal manifestations of abstract entities called “concepts”. Concepts, moreover, are possibly conventionalised. Krifka uses the term “kind” to refer to a conventionalised concept:

[K]inds are considered to be abstract entities that are well-established in the background knowledge of speaker and hearer and can be referred to by definite NPs like *the bear*, which [are] in the extension of kind predicates like *be extinct* or *be a mammal*, and which [are] organised in taxonomic hierarchies. [...] Similar to kinds, concepts are abstract entities related to real objects. However, they need not be well established, but could be construed from scratch. [...] [K]inds form a subset of the more comprehensive sets of of concepts. (Krifka 1995:402)

Objects are related to kinds/concepts via a relation **R** introduced by Krifka et al. (1995). The formula **R**(x,k) is meant for expressing that x is a specimen of the kind k. Krifka et al. write that **R** “is similar to Carlson’s realization relation [...] with the exception that we do not consider stages here” (Krifka et al. 1995:66). They do not say in which respect **R** is similar to R and R’. Recall from above that Carlson’s realisation operators derive terms referring to spatiotemporally bounded entities from terms referring to abstract entities and that the application of this shift in meaning is Carlson’s ultimate explanation for why the respective noun phrases are interpreted existentially. If Krifka et al.’s **R** were to serve a “similar” function, one must conclude that it likewise relates spatiotemporally unbounded entities to spatiotemporally bounded entities. Note that this implies a drastic difference between Carlson’s and Krifka et al.’s semantic conceptions.

In particular, the two conceptions differ in terms of the ontological status being ascribed to objects. For Carlson, objects are not spatially and temporally bounded entities, but abstractions (the same holds for Šmelev’s “abstraktnye individy”). For Krifka et al., by contrast, objects are spatially and temporally bounded manifestations of those abstract entities called kinds or concepts, respectively. In other words, while Carlson considers stages to be tokens and individuals (objects and kinds) to be types, Krifka et al. consider objects to be tokens and kinds to be types – and dispense with stages. Figure 4.4 summarises Krifka’s ontological assumptions:

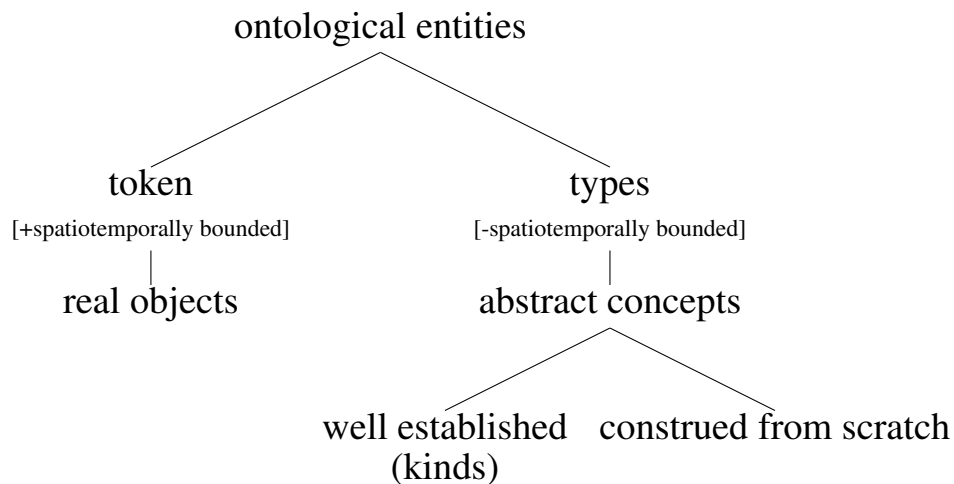


Figure 4.4: Krifka’s ontology

Bulygina (1982) likewise does not use stages as ontological primitives. Let us see how she describes the fact that predications are often limited to the stage of an individual’s lifetime:

V otlīchie ot “charakterizujuščich” predikatov, pripisyvajuščich predmetu kačestvo, [...] predikaty sostojanija opisyvajut vremennye “stadii” suščestvovanija predmeta (lica), pripisyvaja emu imenno kak individuumu priznak, aktual’nyj dlja dannogo otrezka vremeni (ili dlja neskol’kich vremennyh otrezkov), i v étom smysle – priznak prechodjaščij, “slučajnyj”. (Bulygina 1982:33)

In contrast to “characterising” predicates, which ascribe to an object a quality, state predicates [and process predicates, and event predicates (figure 4.3)] describe temporal “stages” of the existence of an object (person). They assign to the object qua object a feature, which is limited to a given time interval (or limited to several time intervals) and is, in this sense, a contingent, accidental feature. (own translation)

Thus, the feeling that in the sentences under (2), the predication is limited to a

temporal stage of the individual's lifetime is modeled along completely different lines by Bulygina and Carlson.

Carlson takes stages, i.e. time slices of individuals, to be ontological primitives. This position suggests that there is a class of lexically dedicated stage-level predicates. Those lexical predicates which do not belong to this class are the individual-level predicates. The class membership of a predicate is reflected in the predicate's selectional restrictions. Accordingly, the relevant piece of information telling whether a predication is interpreted temporarily/accidentally or not is rooted in the selectional restrictions of the predicate expression.

For Bulygina, by contrast, time slices of individuals are interpretive effects, usually resulting when a temporally localised predicate, which describes a state-of-affairs which typically holds for short intervals, combines with an individual term. In her theory, it is not the case that predicates are classified lexically into stage-level and non-stage-level predicates.

In other words, the piece of information telling whether or not the predication holds of a stage is not semantic, but pragmatic by nature. It is the descriptive content of the predicate which makes an "individual-level" usage or a "stage-level" usage the more likely option. How exactly the predicate is used in the end is determined by the context of utterance. On Bulygina's account, predicates are lexically predetermined with respect to the semantic feature of temporal localisation. Depending on whether a predicate is temporally localised (episodic) or not (generic), the respective argument must be interpreted as referring to an object (concretely referring) or as referring to a kind (non-concretely referring).

Kratzer's (1989/1995) theory resembles Bulygina's because also Kratzer considers stages to be interpretive effects rather than ontological primitives. Like Bulygina, Kratzer makes the connection of a predicate to some (spatio)temporal location responsible for its non-generic interpretation. Kratzer also distinguishes lexically (spatio)temporally localised predicates from lexically (spatio)temporally non-localised predicates. On the one hand, there are those predicates which possess a referential situation argument (unfortunately called "stage-level predicates" by Kratzer), on the other hand, there are those lacking such a situation argument (called "individual-level predicates", accordingly). The presence of a (Davidsonian) situation argument has the effect of connecting the validity of the predication to a spatiotemporal location – whatever its exact nature is:

I don't want to commit myself to a particular view with respect to the precise nature of the Davidsonian argument. It may not be an event argument. It may simply be an argument for spatiotemporal location. This is the minimal assumption necessary to explain the data (Kratzer 1995:128)

It is important to note that Kratzer’s distinction between generic and episodic predicates is not based on an ontological contrast between individuals and stages, but rather, as on Bulygina’s system, in which predicates basically come in two forms: with and without a connection to some (spatio)temporal location. Only if a predicate is connected with some concrete spatiotemporal location, the predication will be understood to hold accidentally. (5) shows a Kratzer-style interpretation of the sentence *Bill is running* (l is a spatiotemporal location variable):

$$(5) \quad [[\textit{Bill is running}]] = \exists l. \text{RUN}(\textit{Bill}, l)$$

This concludes our short survey of different ontological and lexical classifications proposed by semanticists from the angloamerican and Russian scientific worlds. Before I summarise the results let me draw attention to an important terminological problem.

4.4 On the sense of the term “referential”

The “western” linguists whose theories have been discussed above use the attribute “referential” in a different way than their Russian colleagues: for the Westerners, object terms (non-generic NPs) as well as kind terms (generic NPs) are referential expressions, whereas for the Russians only object terms count as “referential” – kind terms are said to be non-referential expressions. Needless to say that these different terminological practices are likely to lead to misunderstandings. Note, however, that using the term “referential” only for object reference was not restricted to the Russian scientific sphere³. Compare, for example, Givič₂¹n:

In the terms used here, referentiality is a semantic property of nominals. It involves, roughly, the speaker’s intent to ‘refer to’ or ‘mean’ a nominal expression to have non-empty references – i.e. to ‘exist’ – within a particular universe of discourse. Conversely, if a nominal is ‘non-referential’ or ‘generic’, the speaker does not have a commitment to its existence within the relevant universe of discourse. Rather, in the latter case the speaker is engaged in discussing the genus or its properties, but does not commit him/herself to the existence of any specific individual member of that genus. (Givič₂¹n 1978:293-294)

As you see, Givič₂¹n does not consider expressions designating a kind (“genus”) to be “referential” expressions. For him, referentiality is connected to the “ex-

³Let us therefore quit using the notion of a “western linguist”.

istence” of entities in the universe of discourse. That is to say, referential expressions entail the existence of at least one object instance of a kind (“member of that genus”). A nominal is “referential”, if its interpretation implies the existence of entities in the object domain. Thus, only object-referring NPs count as referential.

At a different place Givi_{1/2}n (1978:323) talks about non-referential (generic) expressions to relate to a “universe of types” and referential expressions to relate to a “universe of tokens”. This corresponds exactly to the position that I, following Krifka (1995), advocate in the present dissertation: the domain of entities is partitioned into a kind domain (=universe of types) and an object domain (=universe of tokens).

I believe that it is justified to speak of a “referentially” used expression whenever reference is made to some entity, be it a type or a token. Therefore, I continue to follow the terminological convention according to which also kind terms are referential expressions.

4.5 Conclusions

In this chapter different variants of what I called “type-token-theories of genericity” have been introduced. All of them share the partition of the ontology into two sorts of entities: spatiotemporally localised entities and not spatiotemporally localised entities. On the basis of this underlying ontological contrast, type-token theories are capable of explaining the contrast between generic and non-generic terms (nominal genericity) and generic and non-generic predicates (predicate genericity): a noun phrase is used generically if it refers to a not spatiotemporally localised entity; a predicate expression is used generically if it ascribes a property to a not spatiotemporally localised entity. Carlson’s (1977) particular theory is, so to speak, the prototype of a type-token theory. As Chierchia (1989) puts it:

Carlson [...] has strongly argued for the fruitfulness of a type/token distinction that cuts across the whole domain of individuals [...] So the domain of entities is articulated into individuals and kinds, on the one hand, and their stages on the other. This classification leads to a striking simple notion of genericity. A generic statement is about an “intensional” entity, namely an ordinary individual or a kind. A non generic statement is about stages, i.e. specific, spatiotemporally located, “extensional” entities. (Chierchia 1989:11-12)

The type-token-theories discussed above differ from each other in how the contrast between spatiotemporally localised entities and not spatiotemporally lo-

calised entities is implemented. The probably most interesting variation concerns the role that objects play within the different approaches.

No semantic theory can dispense with objects. The reason is, of course, that objects are those entities referred to by proper names like *Chomsky*, *Moby Dick* or *Alaska*. While Carlson (1977) classifies objects as not spatiotemporally localised entities, i.e. as types, others like, for instance, Krifka et al. (1995) take objects to be spatiotemporally localised entities, i.e. token.

Common sense tells us that an object, the bearer of a proper name, is three-dimensional and exists in time. How could it be not spatiotemporally localised then? I suspect that Carlson's counterintuitive decision to treat objects as not spatiotemporally localised entities was due to linguistic reasons:

- (6) a. *The sperm whale is a whale.*
- b. *Moby Dick is a whale.*

Both of the generic statements (6a) and (6b) contain the predicate *is a whale*. In one case, (6a), the essential property to be a whale is ascribed to a kind, namely to the kind 'sperm whale'. In the other case (6b), the same essential property is ascribed to an object, namely to the sperm whale whose name is Moby Dick. This poses a problem. If, on the one hand, generic predications are property assignments to not spatiotemporally localised entities but, on the other hand, objects are spatiotemporally localised entities, how to make sense of a sentence like (6b)? The semantic composition of this sentence ends in a semantic mismatch⁴. Carlson solves this problem by uniting kinds and objects under one ontological sort, namely under not spatiotemporally localised "individuals".

My own approach (to be presented in the following chapters) is likewise a type-token-theory. In Line with Krifka (1995) I take objects to be spatiotemporally localised ("real") entities and kinds to be not spatiotemporally localised entities. Generic NPs are kind terms, generic predicates are kind predicates. The question raised in connection with example (6) will be answered in the following chapters.

⁴One might want to postulate a type shifter saving the composition by converting the kind predicate *is a whale* into an object predicate (cf. Ditzler 1992). The problem with type shifters is, however, that their occurrences are often not predictable. As long as there is no independent motivation for the occurrence of a type shifter apart from formally "solving" the mismatch problem, type shifters have to count as ad hoc solutions.

Chapter 5

Indefinite reference to kinds – the pumpkin crusher puzzle

5.1 Overview

The present chapter deals with the semantics of generically interpreted indefinite noun phrases. According to the standard view (Krifka et al. 1995), it is generally not possible to use an indefinite like *a whale* to refer to the kind ‘whale’ – indefinite singular noun phrases can only be used to refer to a subkind of the kind ‘whale’. Following a suggestion put forward by Dayal (2004), I will criticise this position for being too categorical. As we will see, it is well possible for an indefinite noun phrase (NP) to denote the kind. The impression that kind reference is impossible arises because very often it happens that this mode of reference is blocked. However, this is not for semantic, but for pragmatic reasons.

5.2 Reference to subkinds and the pumpkin crusher

Carlson (1977) identified a class of predicates whose argument slot is reserved for kind terms. Sometimes it is the subject position which is lexically restricted to kind terms (*be extinct, die out, be widespread, etc.*), sometimes it is the direct object position (*invent, exterminate, etc.*). Krifka et al. (1995:10) make use of these Carlsonian kind-level predicates to test whether a noun phrase is capable of referring to a kind or not.

- (1) a. **The lion** will become extinct soon.
- b. **Lions** will become extinct soon.
- c. **A lion** will become extinct soon.

Their observation: while both the definite NP in (1a) and the bare plural in (1b) can denote the kind ‘lion’, the indefinite NP in (1c) cannot. Instead, (1c) must be understood in such a way that a *subkind* of the kind ‘lion’ is facing extinction. The pattern repeats itself with direct objects of kind-level predicates like *invent* or *exterminate*. Apparently, (2b) can only mean that Babbage invented a subkind of computer, say a new computer model:

- (2) a. *Babbage invented **the** computer.*
b. *Babbage invented **a** computer.*

Dayal (2004:396) doubts the conclusion that indefinite NPs cannot refer to the kind named by their head noun. As counter evidence she presents an example due to Bart Geurts:

- (3) *Fred invented **a** pumpkin crusher.*

What is it that the NP *a pumpkin crusher* refers to in this example? One possibility is, no doubt, that *a pumpkin crusher* refers to a subkind of the kind ‘pumpkin crusher’ (e.g. to the kind ‘electric pumpkin crusher’). What is special about (3), however, is that *a pumpkin crusher* may also denote the kind ‘pumpkin crusher’ itself. This runs counter to Krifka et al.’s generalisation which I reproduce here as the “K-rule”, consisting of two parts¹:

K1: A noun phrase of the form *a nominal* can be interpreted at the kind-level such that it refers to a subkind of ‘nominal’.

K2: It cannot be interpreted such that it refers to the kind ‘nominal’ itself.

Example (3) is in obvious contradiction with the second part of K-rule (=K2): if *nominal* is instantiated by the expression *pumpkin crusher*, then it should not be possible for *a pumpkin crusher* to denote the kind ‘pumpkin crusher’. Obviously, however, this *is* possible. In the face of this datum, Dayal (2004) rejects the validity of a generalisation like K2. Instead she proposes a rule which I formulate as a tripartite “D-rule”:

D1: Both noun phrases of the form *the nominal* and noun phrases of the form *a nominal* can refer to the kind ‘nominal’.

D2: Kind-referring noun phrases of the form *the nominal* refer to familiar kinds.

D3: Kind-referring noun phrases of the form *a nominal* refer to novel kinds.

Who is right, Krifka et al. or Dayal? Which one is the adequate rule, the K-rule or the D-rule?

¹The term “nominal” is meant to be a variable for simple or modified nouns.

There are two alternative explanations for (3) that can come to mind and that are, in a sense, more conservative than the theory that will be presented below. The first is that *a pumpkin crusher* actually refers to a subkind of pumpkin crusher (in line with K2) and that, since no alternative pumpkin crusher subkinds are known, this subkind matches the kind. If that was correct, however, why does (4) not allow for the kind reading?

(4) *Fred invented an ice crusher.*

Do you know of different sorts of ice-crushers? I don't. I know that there are such things as ice crushers, but I am not familiar with different kinds of ice crushers. Therefore, if the indefinite in (4) was be subkind referring by virtue of K2, it would have to direct me to the kind, by the logic of this alternative explanation. This prediction is not borne out, however, as I only get the subkind reading in (4). Of course, there is the possibility that I induce different sorts of ice crusher from scratch (e.g. the electric versus the mechanical ice crusher) and that this somehow blocks the matching of subkind and kind. But this argument does not go through because I could just as well induce from scratch different sorts of pumpkin crusher (e.g. the electric versus the mechanical pumpkin crusher), but (3) nevertheless allows for the kind reading. As can be seen from this, the problem with the first alternative is that the examples (3) and (4) should pattern on a par (unless you know of different kinds of ice crusher), but they do not.

The second alternative is that (3) exemplifies the representative object interpretation (see chapter 6). The noun phrase *a pumpkin crusher* would accordingly denote an object (the prototype model of the invention) that indirectly delivers the appropriate argument for the kind predicate. Plausible at first sight, this explanation gets into trouble in the face of examples like the following one:

(5) *Fred invented a well-selling pumpkin crusher.*

If the nominal is interpreted at the object-level, as entailed by this idea, how can it be modified by an adjective like *well-selling*? The interpretation of (5) is not that one and the same object sells again and again, but that the kind sells well, which means that one instance of the kind after the other finds a buyer. This questions the second alternative.

The conclusion to be drawn from these facts is that (3) indeed represents a case of indefinite kind reference, and calls for an explanation. In what follows, I will develop a possible explanation by elaborating on the suggestion of Dayal (2004).

5.3 Indefinite NPs in Carlsonian object-level predications

Generically interpreted indefinite noun phrases can be found not only as arguments of Carlsonian kind-level predicates. Therefore, before turning to Carlsonian kind-level predications, it is relevant to our discussion to remember how generic indefinites show up with Carlsonian object-level predications.

Since predicates like those under (6) combine with object-referring proper names, they are classified by Carlson (1977) as “object-level” predicates²:

- (6) a. *Flipper is intelligent.*
- b. *Flipper has flippers.*
- c. *Flipper is a mammal.*

Carlsonian object-level predicates can be defined as follows:

A Carlsonian object-level predicate is a predicate that syntactically selects for an object term as its argument and that semantically assigns a property to an object (namely to the referent of the argument expression).

Predicates classified by Carlson as object-level predicates allow not only for subjects referring to objects, but also for generic subjects. The latter may be formed by bare plural NPs (7), definite singular NPs (8) and, what is relevant for the present discussion, by indefinite NPs (9):

- (7) a. *Dolphins are intelligent.*
- b. *Dolphins have flippers.*
- c. *Dolphins are mammals.*
- (8) a. *The dolphin is intelligent.*
- b. *The dolphin has flippers.*
- c. *The dolphin is a mammal.*
- (9) a. *A dolphin is intelligent.*
- b. *A dolphin has flippers.*
- c. *A dolphin is a mammal.*

How can we explain the generic interpretation of the subject expressions in (7) to (9)? In Carlson’s (1977) system the genericity of the respective noun

²Carlsonian stage-level predicates likewise combine with object-referring proper names. According to Carlson, however, they do not assign a property to the object as a whole (as the predicates under (6) do), but only to some stage of it.

phrases follows from the ontological assumptions underlying his semantic theory: objects form, together with kinds, the ontological class of “individuals”. Being spatiotemporally unbounded (=abstract) entities, Carlsonian individuals contrast with Carlsonian stages which are conceived of as spatiotemporally bounded entities³. Thus, individuals including objects are by definition generic entities. It is only with stage-level predications that an utterance about non-generic entities (stages of individuals) can be made.

Krifka et al. (1995) follow Carlson in treating predicates like those under (6) as object-level predicates⁴. However, Krifka et al. proceed from ontological background assumptions which differ from those of Carlson. The system dispenses with stages, it is objects that function as spatiotemporal entities in Krifka et al.’s ontology. These “real objects” contrast with kinds which are meant to be “abstract concepts” (cf. Krifka 1995).

It is clear that, given these assumptions, the genericity of the subject-NPs in (7) to (9) can no longer be explained by property assignments to abstract entities (i.e. by “individual-level predications”). After all, objects are not abstract entities in Krifka et al.’s view. Therefore, a different explanation is called for.

For Krifka et al. (1995), the generic interpretation of the sentences results from the supposition that the semantic components distributed over the overt syntactic constituents of such a sentence are mapped onto a three-partite quantification structure induced by covert syntactic structure. The quantifier involved in this three-partite quantification structure is supposed to be the generic quantifier GEN. Let us look at an example: a semantic representation for (9a) would be something like (10)⁵:

- (10) a. *A dolphin is intelligent.*
 b. GEN_{x,s} [C(x,s) & DOLPHIN(x)] [INTELLIGENT(x,s)]

According to (10), the sentence *a dolphin is intelligent* is true if in all situations

³Compare the more detailed discussion of the Carlsonian ontology in chapter 4.

⁴Notice at this point that, different from Carlson, Krifka et al. (1995:10) consider a predicate such as *is a mammal* to belong to a special class of “hybrid” predicates. Unfortunately, Krifka et al. decide to call members of this special predicate class “kind predicates” even though they clearly differ from Carlsonian kind-level predicates in also accepting object terms as arguments. We will come back to these hybrid predicates in section 9.

⁵In the representation, “x” is a variable over objects and “s” is a variable over situations (in this case: states). The relation “C” is responsible for excluding all those situations from entering the quantification structure that do not fulfill the felicity conditions associated with the respective predicate (cf. Chierchia 1995); compare the discussion of quantificational generics in chapter 1.

fulfilling the felicity conditions of the predicate *be intelligent*⁶, every object that is a dolphin is at the same time an object that is intelligent. Characteristic of this analysis is that the indefinite NP *a dolphin* by itself, i.e. at the phrase level, is interpreted at the object-level. It is only by embedding it into the particular syntactic environment, i.e. at the sentence level, that its generic interpretation is brought about.

To sum up so far, generically interpreted indefinite singular-NPs appear in predications dealt with as object-level predications in Carlson (1977). According to Krifka et al. (1995), their genericity is the consequence of a covert generic quantifier present in the semantic structure of the respective sentence.

5.4 Indefinite NPs in Carlsonian kind-level predications

Besides object-level predicates, Carlson identifies a class of predicates which he calls “kind-level” predicates and which can be defined as follows:

A *Carlsonian kind-level predicate* is a predicate that syntactically selects for a kind term as its argument and that semantically assigns a property to a kind (namely to the referent of the argument expression).

As already mentioned in section 2, Krifka et al. (1995) observe that indefinite NPs as arguments of Carlsonian kind-level predicates are given an exceptional interpretation which they label “taxonomic interpretation”: while bare plurals (11) and definite singulars (12) denote the kind named by the head noun (in this case: ‘blue whale’), indefinite NPs can only denote a subkind of the kind named by the head noun (in this case: a subkind of the kind ‘blue whale’); recall the K-rule noted above:

- (11) a. *Blue whales are in danger of extinction.*
b. *Blue whales are rare in this region.*
- (12) a. *The blue whale is in danger of extinction.*
b. *The blue whale is rare in this region.*
- (13) a. *A blue whale is in danger of extinction.*
b. *A blue whale is rare in this region.*

An at first sight surprising example is now (14), because in this case it *is* possible to refer to the kind ‘blue whale’ by means of the indefinite *a blue whale*:

- (14) a. *A blue whale is a rare animal.*

⁶A situation involving a dolphin who is suffering from dementia does not belong to this set.

- b. *A blue whale is an animal in danger of extinction.*

If we wanted to keep on working with the K-rule, we would have to argue that the predicates *is a rare animal* and *is an animal in danger of extinction* are, unlike *is in danger of extinction* and *is rare* not kind-level but object-level predicates⁷. Then we could explain the genericity of the indefinite subjects in (13) in the same way as the genericity of the subjects in (10) is explained: by appealing to quantification over situations involving object participants. Note, however, that we would predict the sentences under (15) to be acceptable, but my informants hesitate to accept them:

- (15) a. *?Lora is a rare parrot.*
b. *??Lora is a parrot in danger of extinction.*

Back to sentences like (13) and to the question *why* in these cases reference to the kind named by the head noun is excluded. As we saw, there are two alternative explanations available. Krifka et al. (1995) offer a semantic explanation: reference to the kind named by the head noun is excluded because an indefinite singular NP lacks the semantic potential to denote the kind named by the head noun. According to this view, it is simply a matter of fact that one of the grammatical rules governing the syntax-semantics-mapping in English is the K-rule. In contrast, Dayal (2004) offers a pragmatic explanation: reference to the kind named by the head noun is possible if only the kind designated is satisfying the novelty condition. If this requirement is not fulfilled, as in (13), an alternative interpretation – the subkind reading – will be actualised. If, as in (3), the novelty condition is fulfilled, the speaker may use an indefinite noun phrase to refer to the kind named by the head noun (‘pumpkin crusher’), thereby introducing it into the discourse.

Let us note that, as arguments of predicates like *be extinct* listed by Carlson (1977) as kind-level predicates, indefinite noun phrases must be interpreted “taxonomically”, i.e. subkind referring. However, transforming these predicates into nominal predicates (*is in danger of extinction* → *is an animal in danger of extinction*), amounts to predicates whose indefinite argument is capable of denoting the kind named by the head noun. I cannot tell why this predicate shift brings about a change in the way we understand the indefinite subject but, and this is what matters here, it represents *prima facie* counter-evidence to K2.

⁷Alternatively, we might want to count them as hybrid predicates designating properties possibly true of kinds *or* objects (cf. footnote 4).

5.5 Two kinds of kind-level predicates

Now recall that we are trying to explain why (16a) (with intended reference to the kind ‘trullala’) is not acceptable, but (16b) is.

- (16) a. **A trullala has died out.*
b. *Fred invented a trullala.*

One might want to think that the sentence-initial position in (16a) is for some reason not compatible with an indefinite kind term. As (17) shows, however, this hypothesis must be dropped immediately:

- (17) *A trullala has been invented.*

A better proposal is that of Dayal (2004). Her point of departure was the widespread assumption that the use of an indefinite article is connected to a certain pragmatic novelty condition. This condition brings it about that any individual designated by an indefinite noun phrase must be understood as being newly introduced into the discourse. What is (more or less) common wisdom as far as interpretations at the object-level are concerned (e.g., Heim 1983), is supposed to be true also at the kind-level: indefinite NPs are used to introduce kinds when they have the status of novel discourse referents.

Against this theoretical background, let us assume that (1c), (13a) and (16a) (under their non-taxonomic reading) are rejected because the argument slots of the kind-level predicates *will become extinct soon*, *is in danger of extinction* and *has died out* are syntactic positions which are not capable of satisfying the novelty condition associated with indefinite noun phrases. Let us go on to assume that nothing bars an indefinite from realising the argument of a predicate like *has been invented*. Why should this be so?

There is no such thing as an exhaustive list of predicates forming the class of Carlsonian kind-level predicates. Even though semanticists agree that it is reasonable to isolate such a lexical predicate class, it is a matter of discussion which particular examples should count as kind-level predicates. The paradigm case of a Carlsonian kind-level predicate is supposed to be *(be) extinct*, inclusive of all its variants: *is extinct*, *has died out*, *will become extinct soon*, *is in danger of extinction*, etc. It is not by chance that it is exactly these examples which are typically used to determine whether a given noun phrase can be used to refer to the kind or not. Other predicates classified by Carlson (1977) as kind-level predicates are under debate. Katz & Zamparelli (2005), for example, consider *(be) common* and *(be) numerous* not to be kind predicates, but frequency predicates. Another example: in face of the fact that a sentence like *a white elephant is rare* is completely acceptable, Krifka (p.c.) is led to conclude that *(be) rare* should better be excluded from the class of kind-level predicates.

The examples under (18) show Carlsonian kind-level predications which, in line with the K-rule, disallow indefinite NPs referring to the kind named by the head noun. Notice that the judgments (stars) relate neither to subkind readings (which are always possible) nor to object-level readings (which are sometimes possible):

- (18) a. **A universal scholar is as good as extinct these days.*
 b. **A blue whale is in danger of extinction.*
 c. **An arctic blue whale was eradicated within the last 15 years.*
 d. **A brown hare is on the Red List of endangered species.*
 e. **A sea lion was discovered on the GaLApagos islands.*⁸
 f. **A wolf is getting bigger as we travel north.*
 g. **An H5N1-virus is now rampant in Europe.*

The sentences under (19) likewise show kind-level predicates with indefinite subject-NPs. This time, however, reference to the kind named by the head noun is possible. The sentences should not be interpreted in the sense of the object-level reading which is sometimes also possible⁹:

- (19) a. *A pumpkin crusher has been invented.*
 b. *A Trabant is still seen relatively frequently in Leipzig.*
 c. *A liger has been bred in two different places at the same time.*
 d. *A Royal Flush is rare.*
 e. *An indefinite NP comes in two forms in Maori.*
 f. *A surgeon fish is called surgeon fish because it heals other fish.*
 g. *A chess piece is either black or white.*
 h. *An H5N1-Virus developed in two places at once.*
 i. *A baby's dummy gets lost from time to time.*

What the examples under (19) suggest is that the K-rule is too strong a theoretical generalisation and has to be adjusted accordingly: It is not the case that Carlsonian kind-level predicates *categorically* exclude indefinite kind arguments (under the non-taxonomic reading). Note that, to draw this conclusion, it suffices to find one single counterexample. So if you are not convinced that this

⁸The intended reading is the one under which “discovery” means the first time when the species sea lion was noticed with the place of the discovery, the Galapagos islands, being focused.

⁹Note that adding *so-called* (*a so-called pumpkin crusher has been invented; a so-called Trabant is still seen relatively frequently in Leipzig; etc.*) unequivocally leads you to the intended non-taxonomic kind-denoting reading.

or that predicate in (19) is really a kind predicate, you should focus on (19a). This is the paradigm case which is definitely a kind predicate.

And so it seems to be the case that within the class of Carlsonian kind-level predicates there is a subclass of predicates containing a certain lexical-semantic component which, in effect, is responsible for the impossibility of indefinite kind-level arguments (under the non-taxonomic reading). The paradigm example *die out* is one of them.

The idea is that the lexical-semantic component requires the familiarity of the kind thereby excluding indefinite nominal arguments because the indefinite article is associated with a novelty condition. This story would be an explanation in the spirit of the D-rule (see above). The question to be answered then is: What kind of “lexical-semantic component” plays a role here?

To sum up, there are kind-level predicates allowing for indefinite nominal arguments under the non-taxonomic reading. The widespread lack of this reading can therefore not be due to a categorical feature of Carlsonian kind-level predicates, as suggested by the K-rule.

5.6 *To invent versus to be extinct*

For the time being let us call the kind-level predicates of (18) “extinct-predicates” and those of (19) “invent-predicates”. Two of the predications under (19) are better excluded from the invent-class though, namely:

- (20) a. *A surgeon fish is called surgeon fish because it heals other fish.*
b. *A chess piece is either black or white.*

The reason for this special treatment is that the sentences under (20) do not entail the existence of object instances. Sentence (20a) is true even in a situation where there are no particular surgeon fish (if, for example, surgeon fish had died out). Likewise (20b) does not require the object world to contain real existing chess pieces in order to be true. What these sentences talk about are rather *potential* surgeon fishes and *potential* chess pieces. For all the other sentences under (19), by contrast, it holds that, given their truth, real pumpkin crushers, Trabants etc. must exist or must have existed. Thus, these sentences talk about *actual* individuals.

What both extinct-predicates and invent-predicates share is that they have to be interpreted episodically in the sense that their truth conditions include conditions in the real (“actual”) world. This distinguishes them from predications like those in (20). Extinct-predicates and invent-predicates differ from each other, however, in that the former but not the latter support the K-rule. A defender of the K-rule would have to show that invent-predicates are not kind-level predicates. Let us have a closer look on the particular invent-predicates that we met:

- (21) a. *X has been invented.*
 b. *X has been bred.*
 c. *X developed.*
 d. *X is rare.*
 e. *X is frequently seen.*
 f. *X comes in two forms.*
 g. *X gets lost from time to time.*

These predications can be paraphrased as in (22). Notice that the paraphrases do not aim at giving precise and exhaustive lexical-semantic descriptions. The semantic descriptions offered are by far too simple and partly even wrong, but they will serve our present purposes:

- (22) a. 'there are instances of the kind X, because an inventing-event occurred'
 b. 'there are instances of the kind X, because a breeding-event occurred'
 c. 'there are instances of the kind X, because a developing-event occurred'
 d. 'there are instances of the kind X, little in number'
 e. 'there are instances of the kind X, in large numbers'
 f. 'there are instances of the kind X, in two forms'
 g. 'there are instances of the kind X, several of them disappearing out of their owner's reach'

Compare these paraphrases with those given in (24) for the predicates in (23):

- (23) a. *X is extinct.*
 b. *X is in danger of extinction.*
 c. *X was eradicated.*
 d. *X is rampant.*
 e. *X is getting bigger as we travel north.*
 f. *X was discovered on the GaLApagos islands.*
- (24) a. 'the number of instances of the kind X decreased to zero'
 b. 'the number of instances of the kind X will decrease to zero if the present development will continue'
 c. 'the number of instances of the kind X has decreased to zero owing to brutal force'
 d. 'instances of the kind X are very widespread, with negative consequences'

- e. ‘instances of the kind X in the north are (on average) bigger than instances of the kind X in the south’
- f. ‘it was on the Galapagos islands that instances of the kind X were noticed for the first time’

What the paraphrases are supposed to show is that it is only in (21) that the piece of information about the existence of object instances of the kind (named by the head noun) is part of the assigned property. Only invent-predicates entail the semantic component ‘there are instances of the kind X’ as part of their lexical meaning. In the case of extinct-predicates the assigned property is also concerned with instances of the kind, but it does not include the particular piece of information according to which there are instances of the kind X. Rather, by using an extinct-predicate, the speaker presupposes the existence of instances of the kind X as known to the hearer.

To sum up: Using an invent-predicate, the speaker *asserts* the existence of instances of the kind about which the predication is made. Using an extinct-predicate, the speaker *presupposes* the existence of instances of the kind about which the predication is made. By contrast, the interpretation of predications as in (20) is independent of the existence of instances of the designated kind.

5.7 How the taxonomic reading comes about

Imagine two constituents which form a complex syntactic expression. Each one of the two imposes a certain condition on the way it has to be used. Imagine furthermore that the two conditions contradict each other. The resulting complex expression will fail to be interpretable, unless some reinterpretation can be found under which the two *prima facie* contradicting conditions can exist side by side. I propose that the possible taxonomic interpretation of the sentences under (18) results from the resolution of such an interpretive conflict. Specifically, it results from resolving the conflict between the conditions of use of the indefinite article on the one side and the conditions of use of the extinct-predicate on the other side.

By virtue of their lexical meaning, extinct-predicates presuppose the existence of instances of the kind named by the respective head nouns. Knowing that instances of a certain kind exist implies, for trivial reasons, acquaintance with the kind. Therefore, upon hearing an utterance involving an extinct-predication, the hearer will receive, among others, the piece of information that the respective kind referent must belong to the stock of discourse referents that she is sharing with the speaker (that is to say, the common ground). This lexical presupposition of the extinct-predicate may now come into conflict with the conditions of use of the noun phrase filling its argument slot. When does such a

situation arise?

Morphosyntactic NP-types – by which I mean, for example, NPs accompanied by a definite article (like the English *the raven*), NPs accompanied by an indefinite article (like the English *a raven*) and bare nominals (like the English bare plural *ravens* or the Hebrew bare singular ‘*orev*) – come with their own conditions of use. Which particular conditions these are depends on the lexical meaning of the accompanying article, if present: the syntactic presence of an article restricts the spectrum of possible usages of the noun phrase which it flags.

As far as the definite article is concerned, the speaker can choose this option only if the intended NP-referent has either been previously mentioned in the discourse or is otherwise reconstructable (unequivocally identifiable) in the context. If this is not the case and the speaker nevertheless chooses a definite NP, she will force the hearer to accommodate some appropriate referent. That is to say, a definite NP presupposes the existence of its referent in the domain of discourse (cf. Geurts 1999). It is therefore reasonable for the speaker to use a definite NP if she wants to direct the hearer’s attention to some familiar entity. Definite NPs are, so to speak, special tools for reference to familiar entities.

Bare NPs can also be used to refer to familiar entities. Not flagged by any function word which could impose its lexical condition of use, a bare nominal is free of any pragmatic restrictions concerning the novelty or familiarity of its referent. Unlike a definite NP, it *need not* be used to refer to familiar entities, but it *can*¹⁰.

Indefinite NPs cannot be used to refer to familiar entities.¹¹ Here it is once again the syntactic presence of an overt article displaying its meaning that restricts the range of possible noun phrase usages. Choosing the indefinite article is an option only if the speaker wanted to newly introduce the NP-referent into the discourse while uttering the sentence in which the NP appears. NPs accompanied by an overt indefinite article are, to reinvokethe metaphor used

¹⁰In languages such as English where there is a grammaticalised definite article the bare nominal is *not* used to refer to familiar entities because the availability of the specialist blocks the usage of the allrounder (cf. Chierchia 1998, Krifka 2004).

¹¹For the sake of concreteness, I treat kind-level indefinites as referring expressions, just like definites. The alternative would be that indefinites have an existential interpretation, where the predicate applies to kinds. As far as I can see, the present proposal could also be stated in terms of this alternative – on condition that definites were also existentials (which is not implausible, see Gendler Szabó 2000, Ludlow & Segal 2004). This last requirement is necessary because the pragmatic reasoning outlined below can only work if definites and indefinites share the same semantic format. Otherwise the use of an indefinite instead of a definite would not trigger an implicature of novelty (see Hawkins 1991).

earlier, special tools for reference to novel entities. Figure 5.1 summarises the conclusions so far.

expression	lexical restriction
<i>X is extinct</i>	X must be familiar
<i>the X</i>	X must be familiar
<i>a X</i>	X must be novel
<i>X-pl</i>	no restriction
<i>the X is extinct</i>	okay
<i>a X is extinct</i>	presupposition conflict
<i>X-pl are extinct</i>	okay

Figure 5.1: lexical restrictions

Following Dayal (2004) I will assume that these lexical restrictions are in effect at the object-level as well as at the kind-level. In the second last line of the figure we find the “presupposition conflict” resulting whenever an indefinite NP fills the argument slot of an extinct-predicate: the extinct-predicate requires its kind argument to be familiar, the indefinite NP requires its referent to be novel.

To repeat this in simple words: The speaker wants to tell to the hearer that a certain kind is extinct. To understand that a certain kind is extinct the hearer must know that there have been instances of the extinct kind. Knowing that instances of the kind have existed implies knowing the kind (how else could the hearer have categorised the instances as instances of the respective kind?). Now, if the hearer knows the kind, it is not the adequate linguistic way for the speaker to refer to the kind by means of an indefinite NP. Against this background, a very simple explanation for the arising of the subkind reading suggests itself. Let us take (25) as example.

(25) *A blue whale is extinct.*

The predicate *is extinct* requires of its argument (at least) two things: first, it must denote a kind and, secondly, the previous existence of instances of the kind must be taken for granted. Thus, in any situation suited to host the utterance of (25), the hearer (as well as the speaker, of course) must know that there were instances of the kind denoted by the argument expression of *is extinct* beforehand. From the fact that to know that there were instances of a certain kind one has to know the kind (for trivial reasons) it follows that the hearer (and the speaker) must be familiar with the respective kind. Otherwise the sentence would be pragmatically infelicitous. That is to say, the sentence is adequately

uttered only in a situation in which the kind ‘blue whale’ is hearer-familiar. Now, this runs counter to the ordinary kind-level interpretation of the indefinite NP *a blue whale* under which it denotes the kind ‘blue whale’ presenting it as hearer-new. And so the semantic composition of (25) inevitably runs into a conflict between two conditions: the indefinite NP requires the novelty of the kind while the extinct-predicate requires the familiarity of the kind.

The conflict seems to be unavoidable, but it is not. To overcome it, the referent of the kind-level NP simply has to be *doubled* so that one of the now two kinds, presented as hearer-new, can satisfy the requirement of the indefinite NP *a blue whale* while the other one, presented as familiar, can satisfy the requirement of the predicate *is extinct*. I propose that this is what happens in the case of the taxonomic reading.

That it is possible to double the kind referent is the consequence of the particular architecture of the kind domain. Remember (chapter 3) that a noun like *blue whale* is first of all the linguistic symbol (name) of a kind, namely of the kind ‘blue whale’. Due to the fact that kinds are organised in taxonomic hierarchies, however, the noun *blue whale* is at the same time a predicate characterising the set of all subkinds of the kind ‘blue whale’ plus the kind ‘blue whale’ itself¹². If we consider the semantics of the indefinite article to correspond to a choice function mapping a set onto one of its elements (under the pragmatic condition that the chosen element is hearer-new), then not only the kind directly named by the nominal, but also any of its subkinds is a suitable referent of the indefinite NP. This is taken advantage of in the case of the taxonomic reading.

The predicate *is extinct* presupposes the existence of instances of the kind named by the head noun of its nominal argument. If the argument is realised by the NP *a blue whale*, the existence of blue whale instances is – from the hearer’s viewpoint – to be taken for granted. As a consequence, the kind ‘blue whale’ has to have the status of a familiar entity. Therefore, it cannot be the referent of the indefinite NP for indefinite NPs designate novel entities. The hearer is led to conclude that the NP *a blue whale* designates a kind entity other than the kind ‘blue whale’ – one which is neither previously mentioned nor reconstructable from the context. The question is whether such an alternative kind referent can be found. It can. The reason is that – as outlined above – also subkinds of the kind named by the head noun are suitable referents of an indefinite NP. While it is still the kind ‘blue whale’ which is named by the head noun *blue whale*, it is now some other (hearer-new) kind which is chosen from the kind extension of the head noun to serve the referent of the NP *a blue whale*. This way of resolving the presupposition conflict becomes possible because every instance of a subordinate kind is automatically an instance of the superordinate kind. Let

¹²Compare also the discussion in section 5 of chapter 1.

a blue whale in (25) denote, say, the kind ‘arctic blue whale’. Given this, it is possible that the requirement of the extinct-predicate – that the hearer is familiar with instances of the kind ‘blue whale’ – is satisfied while, at the same time, the requirement of the indefinite – that the referent of the NP, the kind ‘arctic blue whale’, is hearer-new – is also satisfied.

To sum up: Due to the taxonomic architecture of the kind domain, a sentence like (25) is interpretable even though the lexical-pragmatic requirements of the indefinite article and of the extinct-predicate seem to contradict each other. Letting the NP denote a subkind of the kind named by the head noun, the presupposition of the predicate (it must be known that there are instances of the kind ‘blue whale’) and the presupposition of the argument (the referent of *a blue whale* must be hearer-new) can both be met.

5.8 Well-establishedness

We saw that an indefinite NP is well capable of denoting the kind named by its head noun. The (wrong) impression that indefinite NPs can only be used to refer to a subkind arises because in the majority of cases of kind-level reference indefinite NPs indeed do denote a subkind. It frequently happens that non-taxonomic reference is blocked. The reason is the novelty condition associated with the indefinite article that blocks the non-taxonomic reading, for example, if the indefinite NP is filling the argument slot of an extinct-predicate. Filling the argument slot of an invent-predicate, by contrast, the indefinite NP may well denote the “uppermost” kind.

There is a certain observation mentioned in the beginning of this chapter which I have, for expositional reasons, ignored up to now. But now it can no longer be ignored. What I mean is the fact that indefinite subkind reference can be attested not only in the argument position of extinct-predicates, but also in the argument position invent-predicates: While (26a) relates to the kind named by the head noun *pumpkin crusher*, namely to ‘pumpkin crusher’, (26b) relates to a subkind of the kind named by the head noun *computer*. (26b) is normally understood as reporting on Fred’s invention of a new kind of computer (cf. (2)):

- (26) a. *Fred invented a pumpkin crusher.*
b. *Fred invented a computer.*

At first sight, this observation seems to undermine my proposal. A closer look, however, reveals that this is not the case. The crucial difference between (26a) and (26b) is that the kind denoted by the direct object in (26b) is “well-established in the background knowledge of speaker and hearer” (Krifka 1995),

but not so the kind denoted by the direct object in (26a)¹³. What does “well-established in the background knowledge of speaker and hearer” mean? I think a reasonable definition is the following:

A kind is well-established as soon as speaker and hearer have agreed upon categorising a given set of objects as instances of the respective kind.

Note that according to this definition, the well-establishedness of a kind goes together with knowledge of object instances sharing a certain property, namely the kind property. If this definition is on the right track, it explains why indefinite NPs whose head noun names a well-established kind can only be used to refer to a subkind. The explanation is virtually the same as in the case of extinct-predicates. The generalisation is:

Whenever the existence of instances of the kind named by the head noun must be taken for granted while interpreting the syntactically full fledged NP, the use of an indefinite NP will automatically lead to a subkind interpretation, because an indefinite NP can only be used if the NP-referent is hearer-new and, given the kind named by the head noun is familiar, only a subkind can be hearer-new.

To sum up, an indefinite NP has the semantic potential to non-taxonomically denote the kind. The novelty condition associated with the indefinite article, however, brings it about that this mode of reference is blocked whenever the existence of instances of the kind is presupposed. This is the case (i) if the indefinite NP appears in the slot of an extinct-predicate and (ii) if the head noun of the indefinite NP names some well-established kind.

5.9 Conclusions

This chapter questioned the received view according to which indefinite noun phrases cannot refer to kinds for semantic reasons. It was proposed instead that indefinites *in many cases* cannot refer to kinds, due to pragmatic constraints. One important case is when the indefinite associates with novelty and therefore clashes with semantic requirements of the predicate. The aim of this chapter was basically to elaborate on this (pumpkin crusher) case.

We also met other cases of apparent kind reference of indefinites:

(27) a. *A blue whale is a rare animal.*

¹³Krifka (1995) reserves the term “kind” for well-established categories only whereas for me, also ad hoc categories are kinds.

b. *A surgeon fish is so-called because it heals other fish.*

If this is kind reference, then it is surely not motivated by novelty of the named kinds ‘blue whale’ and ‘surgeon fish’, respectively, as these kinds are well-established. Given that the present chapter is on the right track, it should be possible to identify a pragmatic implicature other than novelty that can account for why well-establishedness does not lead to a presupposition failure in (27). This must remain a task for future research. Let me speculate that sentences like (27a) are uttered in “encyclopedic contexts” where it is obvious to all that the speaker will talk about some well-established kind. As a consequence of that, the use of the definite would lose its significance of signalling the familiarity of the kind, and the use of an indefinite noun phrase would not conversationally implicate novelty; the indefinite would become an alternative option. Such a story would at the same time apply to categorizing sentences like the following one:

(28) *A blue whale is a mammal.*

A similar story could be told with respect to (27b): drawing attention to the naming of the kind, the predicate *so-called* invites the implicature that the question of familiarity with existing instances of the kind is beyond relevance. This way it creates a context in which the use of an indefinite will not automatically go together with non-familiarity.

Here we cannot tell whether or not any successful explanation along these lines could be found. If it turns out to be impossible to explain (27) in terms of indefinite reference to kinds, then we are left with the radical alternative of excluding these sentences from the class of kind-level predications. We would be left with the case of the pumpkin crusher, which is interesting enough, as I hope to have shown in this chapter.

Chapter 6

Lexical classes of predicates – normal and Carlsonian kind-level predicates

6.1 Overview

Successful as it was in explaining bare plurals, the theory of Carlson (1977) was subject to criticism. The presumably most important objection that has been raised concerns the mismatch between Carlson's theoretical predictions with respect to indefinite singular noun phrases and their actual behavior. The problem is that in and of itself, the theory cannot account for the fact that kind-level singular indefinites denote subkinds – additional ad hoc stipulations are needed (e.g. Heim 1991:523).

If correct, the results of the previous chapter suggest a way of rescuing Carlson's theory (as far as this criticism is concerned): What if the seeming fact that singular indefinites must denote subkinds is only an impression that arises because a kind denoting indefinite is, due to pragmatic constraints, a curiosity? Then we would have to rethink Carlson's theory.

The goal of the present chapter is to spell out how a revised version of the Carlsonian predicate classification could look like if the discussion of chapter 5 is taken into account. The new predicate classification that I will present keeps with Carlson's basic assumption that there is a systematic ontological distinction between type entities and token entities. In contrast to Carlson's original theory, however, and in line with Krifka et al. 1995, I take objects to play the role of tokens while kinds alone play the role of types. Consequently, stages can be dispensed with (compare chapter 4). Predicates that in Carlson (1977) are classified as object-level or stage-level, respectively, together fall into a single lexical class. Specifically, I will propose that they are registered in the lexicon as "pure" kind-level (sic!) predicates that can, given the respective desire of the speaker, be turned into object-level predicates by means of a semantic operation called "spatiotemporal localisation". Most of the predicates classified by

Carlson as stage-level still have a special lexical status in that they, once spatiotemporally localised, refer to Davidsonian events. Somewhat surprisingly, predicates classified by Carlson as kind-level turn out not to represent normal kind-level predicates. The reason is that the property assignment of predicates like *be extinct*, *die out*, etc. involves conditions imposed on the mode of existence of object instances of the kind. Thus, the property assigned is not a “pure” kind property. I propose to analyse Carlsonian kind-level predicates as complex existence predicates.

To set the stage for discussion, in section 2 I first of all repeat the predicate classification that is standardly assumed. Section 3 is included to give a bird’s eye view of the alternative to be developed in the subsequent sections. Of central relevance for my alternative approach will be the “representative object interpretation” (Krifka et al. 1995) that object-referring NPs actualise given certain circumstances. It will be introduced in section 4. As I argue in section 5, *every time* an object term forms the subject of a kind-level predicate there is reason to call the resulting interpretation “representative object interpretation”. After all, given that kinds are sortal concepts (cf. chapter 2), objects per se are nothing but representatives of kinds in the real world. From this it follows that object reference necessarily involves kind reference. This sortalist position is defended in section 6. In section 7 I demonstrate how Carlsonian stage-level predicates integrate into the system developed so far. If all predicates are lexically kind-level predicates there will have to be some grammatical mechanism to turn basic generic predicates into episodic predicates. I call it “spatiotemporal localisation”. Even though, in principle, every kind-level predicate can undergo spatiotemporal localisation, some are for conceptual reasons more, some less likely to undergo this process. This will be pointed out in section 8. Section 9 is dedicated to the lexical peculiarity of Carlsonian kind-level predicates. In section 10, I analyse them as specimens of existence predicates. Section 11 gives a summary and concludes the chapter.

6.2 Predicate classes according to the standard theory

In this section, we briefly recapitulate the standard assumptions made with respect to generic predications involving indefinite subjects. The basic source is once again Krifka et al. (1995). The aim is, first, to highlight the predicate classification underlying these assumptions and, secondly, to point certain problems that such a classification faces.

Forming the subject of a Carlsonian object-level predicate, a by itself object-level indefinite NP will be understood generically if it is interpreted within a generalised quantificational structure (induced by covert syntactic structure):

(1) *A dog has four legs.*

Forming the subject of a frequency predicate, a by itself object-level indefinite NP will be understood generically because it likewise appears within the syntactic context of some “multiplier”, in this case of the frequency predicate:

(2) *A white elephant is rare.*

Forming the subject of a Carlsonian kind-level predicate like *die out*, *be extinct* etc., the indefinite NP is supposed not to denote at the object-level, but rather at the kind-level. Due to the K-rule¹, the referent of the indefinite NP must be some subkind of the kind named by its head noun. Thus, in this case, the reason why the noun phrase is understood generically is that the nominal determines its denotation not within the ordinary domain of objects, but within the kind domain:

(3) *A blue whale is in danger of extinction.*

Also, there is the case of generically understood indefinite NPs forming the subject of special, hybrid predicates as identified by Krifka et al. (1995):

With some other predicates, such as *be a mammal*, *be domesticated*, and *be protected by law*, the kind-referring interpretation of the subject is not the only one; indeed, a proper name referring to a particular animal can also be used as subject with these. Yet when a general term is used as subject NP, the kind-referring interpretation has at least priority over the object-referring interpretation. We call predicates which favor a kind-referring interpretation of an argument *kind predicates*.
(Krifka et al. 1995:10)

Predicates of this type combine with both object terms and kind terms. As the quotation suggests, the syntactic combination with a kind-referring expression represents the somewhat more natural option. Appearing in the argument slot of a predicate like this, an indefinite NP interpreted at the kind-level will denote the kind, not the subkind:

(4) *A parrot is a bird.*

One wonders: if the hybrid predicate *is a bird* functions as a kind-level predicate, as in (4), why does the kind term argument *a parrot* not refer to a subkind, as it would be predicted by the K-rule? This problem becomes even more evident in cases like (5) where the predicate is an explicit kind-level predicate:

¹This is the rule according to which kind-level indefinite NPs can only be used to refer to subkinds (recall chapter 5).

(5) *A cockatoo is a kind of parrot.*

Although the predicate is definitely kind-level, the subject does not refer to a subkind of ‘cockatoo’, but to the kind ‘cockatoo’ itself. This runs counter to the K-rule. Why does the K-rule not apply in this case?

Related to this is the problem that was the topic of the previous chapter: being arguments of invent-predicates, indefinite NPs need not refer to subkinds. Why is this? Either invent-predicates are not kind-level predicates at all, or there are many exceptions to the K-rule.

Having pointed to these problems, we believe to have enough reason to try an alternative. Figure 6.1 summarizes the predicate classes which are implied by the standard theory:

lexical class	examples
object-level	<i>have four legs, be intelligent, ...</i>
kind-level	<i>die out, come in different forms, ...</i>
hybrid	<i>be protected by law, be a bird, ...</i>

Figure 6.1: predicate classes according to the standard theory

6.3 Outlining the alternative

Following Carlson (1977), Krifka et al. (1995) assume a partition of the lexicon into object-level predicates and kind-level predicates. As discussed in chapter 4, however, there is a problem with such a categorial distinction. Croft (1986) points out that every predicate the lexical meaning of which is analysed as an object-level predicate by Carlson can also be used as a kind-level predicate. There is not a single exception to this generalisation. Carlson is aware of this situation: “whatever may be meaningfully predicated of an object may also be meaningfully said of a kind” (Carlson 1977:248). In the face of this, how useful is it to lexically classify predicates like *have four legs, be intelligent* etc. as object-level predicates and postulate a grammatical (generalisation) operation converting object-predicates into kind-level predicates (cf. chapter 4)?

On the one hand, whatever may be meaningfully predicated of an object may also be meaningfully predicated of a kind. On the other hand, it is not the case that whatever may be meaningfully predicated of a kind may also be meaningfully predicated of an object. It is, for example, not meaningful to say of

an object that it was in danger of extinction. To capture these facts, Carlson assumes the following:

- Natural language predicates manifest themselves in the lexicon (i) as predicates assigning properties to kinds and (ii) as predicates assigning properties to objects.
- There is a grammatical operation (“generalisation”) converting object-level predicates into kind-level predicates.

In addition, Carlson assumes that there are lexical predicates assigning properties to stages. There is, however, an alternative way to deal with the facts described above:

- Natural language predicates manifest themselves in the lexicon exclusively as predicates assigning properties to kinds.
- There is a grammatical operation (“spatiotemporal localisation”) converting kind-level predicates into object-level predicates.
- For a subset of lexical kind-level predicates it holds that they involve a semantic component which blocks their spatiotemporal localisation.

This is the general idea behind the predicate classification for which I will argue in the present chapter. Note a welcome side-effect: since lexical predicates are conceived of as being kind-level, every lexical predicate is predicted to have a generic meaning.

6.4 The representative object interpretation

Let us carefully recall the data to be accounted for. The first thing to note is that Carlsonian object-level predicates (6) combine with both object terms and kind terms. In this respect, they equal Krifka et al.’s hybrid predicates (7). Note that every one of these sentences expresses a generic statement:

- (6) a. *Fido has four legs.*
b. *A poodle has four legs.*
c. *Fido is intelligent.*
d. *A poodle is intelligent.*
- (7) a. *Fido is a dog.*
b. *A poodle is a dog.*

Those predicates that are analysed as stage-level predicates by Carlson can likewise appear with object terms as well as with kind terms actualising generic interpretations:

- (8) a. *Fido barks.*
b. *A poodle barks.*

I will postpone the discussion of Carlsonian stage-level predicates to section 7. Unlike stage-level and object-level predicates, Carlsonian kind-level predicates only combine with kind terms – at least at first sight:

- (9) a. *The sperm whale is in danger of extinction.*
b. **Moby Dick is in danger of extinction.*
c. *A poodle is a kind of dog.*
d. **Fido is a kind of dog.*
e. *The Coke bottle comes in different sizes.*
f. **The Coke bottle I bought yesterday comes in different sizes.*

We see that all sorts of predicates can be used to express generic statements thereby combining with kind terms as well as with object terms. The exception to this rule is represented by Carlsonian kind-level predicates which seem to combine with kind terms only.

That Carlsonian kind-level predicates are only compatible with kind terms is viewed as a matter of fact by the standard theory:

There are some predicates with argument places that can be filled only with kind-referring NPs. Examples are the subject argument of *die out* or *be extinct* and the object argument of *invent* or *exterminate*. The reason is, of course, that only kinds (not objects) can die out or be invented. (Krifka et al. 1995:10)

Only kinds can die out, be invented etc. The following examples once more support this view:

- (10) a. **Che Guevara died out.*
b. *?The army exterminated King Kong.*

But why are the following sentences acceptable? After all, an object term shows up in the argument slot of a Carlsonian kind-level predicate in these sentences:

- (11) a. *Weizenbaum invented Eliza.*
b. *My neighbour invented Fido, a robot dog that fetches his daily newspaper.*
c. *Wilmot bred Dolly.*

What the proper names *Eliza*, *Fido* and *Dolly* refer to is the prototype of some newly created kind. In (11b), it is even explicated what the newly created kind is: ‘a robot dog that fetches my neighbour’s daily newspaper’. Prototypes are real objects which, being the first exemplars of the new kind, are special in that they are representatives of the new kind. Given this, also the examples under (11) support the position of the standard theory: proper names are allowed in the argument slot of a Carlsonian kind-level predicate only if the bearer of the proper name plays the role of a representative of a kind, as does, for instance, the prototype exemplar of a certain invention. Krifka et al. (1995:85) speak of the “representative object interpretation” in this connection. A natural example is the following one (collected from the Internet):

- (12) *The prominent use of clownfish in the film “Finding Nemo” prompted mass purchase of the animals for children’s pets in the United States. Animal welfare activists called on customers to buy only exemplars that are certified by the Marine Aquarium Council (MAC). Only this way, they said, one can make sure that Nemo will not die out.*

Nemo, star of the film “Finding Nemo”, is not the prototype of a newly created kind. His status as a representative of the kind ‘clownfish’ is rather due to his presence in mass media. A particular fish became a symbol for a whole kind such that the name *Nemo* refers not only to the individual Nemo but also to the kind ‘clownfish’.

The representative object interpretation is not restricted to proper names. Given an appropriate contextual setting, descriptive NPs referring to objects may likewise be used in this way. Predestined to frame such usages are zoo-contexts. After all, a zoo is exactly this: a collection of object individuals representing kinds. So let us take a walk in the zoo:

- (13) a. [at the lion’s cage:] *Look kids, this is the lion.* [Krifka et al. 1995]
 b. [at the tiger’s cage:] *The tiger you see here is in danger of extinction.*
 c. [in the birdshouse, pointing to an exemplar:] *This bird was discovered by Darwin on the Galapagos islands.*
 d. [pointing at a platypus:] *Hard to believe, but this strange animal over there is indeed a kind of mammal.*

The representative object interpretation comes about if an object term appears in the argument slot of a kind-level predicate and if the referent of the object term has, for some reason, the semiotic status of a symbol for a kind. Given these preconditions, the object term can function as a kind term: reference to the kind is realised indirectly via reference to a representative of the kind.

The object individuals Dolly and Eliza possess the status of a symbol for a kind because they are the first (and only?) exemplars of a new kind. Moreover, the new kind is meaningful enough to make its way into the stock of shared knowledge of the speaker community. In this respect, Dolly and Eliza contrast with Fido from (11b) who, to the sorrow of my neighbour, never became popular. For this reason, I explicated the newly created kind for which Fido stands as a prototype: *robot dog that fetches the daily newspaper*. In zoo contexts, the kinds represented by the exhibited animals are usually named on signs fixed to the animal cages.

It is reasonable to ask why the examples under (10) are uninterpretable and thus unacceptable. Why should it not be possible to assign the respective kind-level property to a kind represented by the object individuals Che Guevara or King Kong, respectively? This objection seems justified. The sentences under (10) are rejected for pragmatic reasons, semantically they *are* interpretable. What we would have to do is think of Che Guevara as a representative of a kind, e.g. of the kind ‘full-blooded revolutionist’. Similarly, King Kong would have to represent a kind exterminated by the respective army. Under these circumstances, we would deal with interpretations in analogy to example (12).

Now that we are sensitive to the representative object interpretation we have to rethink not only the judgments of (10), but also those of (9). All of the sentences in (14) are interpretable, although, at least partly, they are certainly a burden to our imagination. We have to think of circumstances under which the objects Che Guevara, King Kong, Moby Dick, etc. are symbols for kinds:

- (14) a. # *Che Guevara is extinct.*
 b. # *The army exterminated King Kong.*
 c. # *Moby Dick is in danger of extinction.*
 d. # *Fido is a kind of dog.*
 e. # *The Coke bottle I bought yesterday comes in different sizes.*

The representative object interpretation shows two things: First, Carlsonian kind-level predicates are syntactically combinable with object terms. Secondly, this does not contradict Krifka et al.’s (1995) statement that only kinds can die out or be invented. Because under the representative object interpretation an object term functions *as a kind term*, this is, on the contrary, support for Krifka et al.’s statement.

In this section I introduced the representative object interpretation in connection to Carlsonian kind-level predications. Postponing the discussion of Carlsonian kind-level predicates until section 8, I now go on to discuss the representative object interpretation in connection with the other predicate classes.

6.5 Categorising statements as representative object interpretations

The representative object interpretation can be described as follows: (i) An object term appears within the argument slot of a kind-level predicate. (ii) By virtue of the fact that the object referred to by the object term is a representative of a kind, the object term refers not only to the object, but also, so to speak indirectly, to a kind. It is to this kind that the kind-level predicate ascribes its property. In this section, I want to draw attention to structural similarities between the above noted sentences realising the representative object interpretation and the sentences under (15) expressing statements which are sometimes called “categorising statements”.

- (15) a. *Chomsky is an anarchist.*
b. *This bird is a stork.*
c. *I am a human being.*

Uttering (15a), the speaker expresses that the object individual Chomsky is an instance of the kind ‘anarchist’. The communicative goal of the speaker is to inform about a certain categorisation of the respective object individual. In (15), the kind (of which the subject referent is claimed to be an instance) is delivered by an NP introduced by the indefinite article. Categorising sentences also come with definite articles, though:

- (16) a. *Look kids, this is the lion.* [= (13a)]
b. *Superman is the saviour of mankind.*
c. *Chomsky is the darling of the left.*

As mentioned above, Krifka et al. (1995) consider (16a) to exemplify the representative object interpretation. Does this mean that also the other sentences under (16) express representative object interpretations? If so, how about the categorising sentences in (15)? There is definitely something that a sentence realising the representative object interpretation and a categorising sentence have in common. What are the similarities, and what are the differences?

In (15) as well as in (16) objects are asserted to be instances of kinds. However, while we readily accept more than one object instantiating the kind in (15) we hesitate to accept more than one instance with respect to (16):

- (17) a. *Chomsky is an anarchist and my neighbour is too.*
b. *This bird here is a stork and that bird over there is too.*
c. *I am a human being and Socrates is too.*
d. ??*This is the lion and that is too.*
e. ??*Superman is the saviour of mankind and Batman is too.*

- f. ??*Chomsky is the darling of the left and Michael Moore is too.*

However, there are cases in which the kind referred to by the definite NP may well have two or more instances. Usually, such examples involve kinds which are expected to have only one instance but which, somewhat surprisingly, turn out to indeed have two or more instances.

- (18) a. *The gardener AND the butler are the murderer.*
b. *Schulze, Miller and Petersen are the winner of the jackpot.*
c. *Herbert is the perfect man for Augustine, and so are Klaus, John, Bill and Walter.*

The expressions *the murderer*, *the winner* and *the perfect man* refer to kinds that are expected to have singleton instantiation sets. Where does this expectation come from? From a lexical-semantic point of view, the nouns *murderer*, *winner* and *perfect man* definitely do not describe functional concepts, i.e. concepts restricted to singletons only. Nevertheless, they are frequently *used* as functional concepts. Thus, functionality arises as a consequence of pragmatics: murder is typically committed by one murderer, the goal of a competition is typically to choose one winner, and being perfect is something so special that it is unlikely that two objects could reach this standard (already one object being perfect is extraordinary).

If the (English) speaker intends to use a noun as a functional concept, she has to mark it by the definite article (Löbner 1985). In other words, the definite article instructs the hearer to interpret the nominal as unequivocally referring. Consequently, the hearer expects that there will be only one object referent, given the expression is interpreted at the object-level. The postcopular-NPs in (18), however, are not interpreted at the object-level, but rather at the kind-level. Still these NPs refer unequivocally, not to a certain object but to a certain kind. In order to be unequivocally identifiable, this kind must be pre-established in the common ground of speaker and hearer. Thus, the requirement for unequivocal reference is satisfied at the kind-level. Consequently, there is no need for unequivocal reference at the object-level. This is why the sentences in (18) are not contradictions. We can go even further saying that it is *because* the subject in (18a) refers to two objects that a kind-level interpretation of the definite singular-NP following the copular verb is the only plausible interpretive option.

In contrast, the definite descriptions in (17d-f) may be understood not only as referring to kinds, but also as referring to objects. In the latter case, this amounts no longer to categorising statements, but to identity statements informing the hearer that two apparently different entities are in fact one and the same entity. If the main sentence is interpreted as an identity statement, the subsequent sentences will cause confusion: if the object individual Superman was identical the

the object individual designated by the expression *the saviour* and so was the object individual Batman, then Superman and Batman must be one and the same person. This runs counter to world knowledge (in that case knowledge about a fictional world) and therefore, I believe, the sentences are felt to be strange. If, on the other hand, the definite descriptions in (17d-f) are understood as referring to kinds, it is not excluded that there are two or more object instances of the denoted kind (as with (18)). This interpretation may be pragmatically marked, but it is not excluded on semantic grounds.

Thus, the sentences in (15) and (16) do not differ with respect to a semantic restriction to singleton instantiation sets. The sentences differ rather with respect to whether the kind denoted by the postcopular expression is pre-established (familiar) in the current discourse or not. In (16), unlike in (15), the kind of which the subject referent is asserted to be an instance is presented as a category already established in the common ground of speaker and hearer. As a consequence, from the hearer's perspective, the new information is only the relating of some familiar object to some familiar kind.

The appropriate context to utter such sentences is therefore a context where the kinds named by the postcopular NPs as well as the objects named by the proper names represent discourse-established entities. Think of a quiz show where the candidate is asked to correctly relate the sportsman names *Muhammed Ali*, *Robby Nash* and *John McEnroe* to the sportsman categories 'tennis player', 'windsurfer' and 'boxer'. The candidate would win if she answered:

- (19) *Muhammed Ali is the boxer, John McEnroe is the tennis player and Robby Nash is the windsurfer.*

If the sentences under (16) containing kind-level definite NPs can be viewed as actualising representative object interpretations, why not also the categorising sentences under (15)? The only difference is that the objects in (15) represent kinds newly introduced into the discourse while the objects in (16) represent already established kinds.

Let us recapitulate: The first thing to note that only kinds can be invented, can die out etc. This is why, as a rule, only kind terms can fill the argument slot of Carlsonian kind-level predicates. The systematic exception to this rule is formed by object terms referring to objects that function as representatives of kinds. If the object referred to is a representative of some kind, the object term is capable of filling the argument slot of predicates such as *has been invented* or *is extinct*. This is possible because the object term not only refers to the object but (indirectly) also to a kind. Second we note that predicates such as *is a lion* or *is the lion* can combine with kind terms, but also with object terms. In the former case, the result is a statement about a kind which is said to carry

the (kind-level) property of being a lion or being the lion, respectively.² In the latter case, the result is a categorising sentence expressing a statement about an object which is said to be an instance of the kind ‘lion’. If the postcopular noun phrase is introduced by a definite article, it is explicitly expressed that the kind ‘lion’ must be interpreted as a category presupposed in the common ground.

Somewhat stretching the notion of a “representative”, I propose to consider not only categorising sentences involving definite postcopular NPs but also categorising sentences involving indefinite postcopular NPs to express representative object interpretations. Specifically, I propose the following:

- (i) **Representative object interpretations proper** result when, for a group of speakers, the denoted object is established in the real world as a symbol for a kind independently of the utterance.
- (ii) **Representative object interpretations in a wider sense** result when it is the utterance itself that serves to establish that the denoted object represents a certain kind.

Common examples for (i) are Dolly as a representative for the kind ‘cloned sheep’, Nemo as a representative for the kind ‘clown fish’ and probably also Bin Laden as a representative for the kind ‘terrorist’. (ii) is the case of categorizing sentences where the object is claimed to be a representative of some discourse-old or discourse-new category.

is presented as established or newly introduced in discourse. According to this view instances per se are nothing else than representatives of kinds manifesting themselves in the real world.

To sum up: As long as we consider predications over kinds only, we find no reason for distinguishing Carlsonian kind-level predicates from other predicates. Every predicate can be combined with a kind term to assign a (kind-level) property to the respective kind. Moreover, every predicate can combine with an object term. If a kind predicate combines with an object term the kind-level property will be assigned to a kind represented by the respective object. In the case of categorising sentences the object is identified as a representative of the kind through the utterance. In the next section I will propose that prior to the utterance the kind represented by the object is semantically underspecified and that it is the very purpose of such utterances to fix a specific kind (the terminology “categorising statement” simply hits the nail on the head).

²Statements about a kind X which is said to bear the property of being the kind Y are, for instance *Not the sperm whale, but the orca is the killer whale* or *Amanita muscaria is the fly agaric*.

6.6 No reference to objects without reference to kinds

I assume that the sentences under (20) all express kind predications. Such an analysis has become possible because, as I showed in chapter 5 contra the standard view, indefinite NPs may well refer to the kind named by the head noun of the NP. Accordingly, each sentence involves a kind term combining with a kind-level predicate:

- (20) a. *A poodle has four legs.*
b. *A poodle is intelligent.*
c. *A poodle is a dog.*

Sentences like these are uttered in order to ascribe a kind-level property to the kind denoted by the subject-NP. The denoted kind then forms the topic of the utterance. Note that, as a consequence of this analysis, the sentences do not talk about objects at all. Or, to put it differently, the proposed analysis comes with the prediction that these sentences do not impose any truth conditions on the object domain. This is a welcome consequence because it helps to understand the well-known peculiarity of generic predications that is so notoriously difficult to deal with in truth conditional semantic terms: the existence of exceptions, i.e. of object individuals not carrying the property described by the predicate expression (cf. Krifka et al. 1995, Cohen 2002; see also chapter 1). In the extreme, a generic sentence may be true even if not a single object carries the predicate property. (20a), for example, may be true even in a situation where, by accident, every real existing poodle (=every poodle object) has lost one of its legs (cf. Krifka et al. 1995:44).

I furthermore assume that the sentences in (21) express kind-level predications *even though* the subject is formed by an object term:

- (21) a. *Fido is a dog.*
b. *Fido has four legs.*
c. *Fido is intelligent.*

It is always an object term, in the case at hand the proper name *Fido*, that syntactically combines with a kind-level predicate. Sentences like these are uttered in order to inform the hearer of a categorisation that the speaker takes to be valid and relevant: the object referent of the subject expression is asserted to be an instance of a kind which is described by the predicate expression. The predicate of (21a) explicitly names the kind, ‘dog’. Let us postpone the question of how to identify the kind implicit in the predicates of (21b) and (21c), addressing another problem first: How can a kind-level property be assigned to an object? How can semantic composition succeed if the entities composed with each other belong to different ontological levels?

It seems that my analysis runs into a contradiction. On one hand, the topic of utterance is identified as an object. As a consequence, the respective sentence expresses a statement about an object. On the other hand, the property is assigned to a kind. As a consequence, the sentence expresses a statement about a kind. How can these two consequences be true at the same time?

To rescue the semantic composition one might want to assume that the copular has multiple meanings, the instance-of relation between an object and a kind being one of these. In (21a), one could then say, the copular verb actualises exactly this particular meaning (compare the discussion in Dölling 1997). However, such an analysis is not capable of dealing with sentences like (21b). Apart from that, I want to treat the predicate *is a dog* inclusive of the copular verb as a kind-level predicate. Only this treatment enables a unique semantic analysis for the formally identical predicates in (20c) and (21a) which is no doubt a desirable result from a theoretical point of view. Therefore, I commit myself to the strange assumption that a proper name is in some mysterious way related to a kind to which the kind-level predicate may assign its property.

I solve this problem by taking into account the fact that every reference to some object *automatically implies* the reference to some (often) underspecified kind. Whenever the speaker makes use of an object term to draw the hearer's attention to a certain object and is successful in doing so, both speaker and hearer must have categorised the object of their shared attention in one way or the other. It is impossible to draw attention to some object without adopting a certain perspective on the object, i.e. without categorising the object (Tomasello 1999:118). As Dölling (1993) puts it:

I suggest that a kind as a projection of conceptual structure is cognitively and hence ontologically prior to objects, which are specimens or instances of this kind. This follows from the fact that in order to identify an object instantiating a kind, we have to relate this object to the corresponding kind, while we can identify a kind independently of whether we refer to a special instance or not. (Dölling 1993:140)

Given this, the fact that proper names *when they are used* relate to kinds is in no way mysterious but rather obvious. That any object is necessarily related to a kind follows directly from the (nonlinguistic) assumptions I take for granted in this thesis (cf. chapter 2). Let me briefly repeat these assumptions.

We humans are (mobile) object individuals in a world full of other object individuals. To successfully navigate through this world of objects we have at our disposal the capacity to recognise objects. To recognise an object is, trivially, to identify an appearance *as an object*. This in turn is nothing else than to categorise an appearance as an object which, again trivially, presupposes the availability of the appropriate category: the object concept (or object schema).

Consequently, every identified object, whether or not we give a name to it, must be an instance of a sortal concept (kind), be it the most general of all kinds, the kind ‘object’. The kind ‘object’ can be thought of as the superkind of an enormous conceptual system of kinds (the kind domain) accumulated by every one of us during the course of our lifetime for the purpose of sorting the real world of objects in which we live. By means of language, it turns out, we may either refer to the particular objects of the real world (inclusive of ourselves) or to the abstract concepts, kinds, by means of which we sort the objects.

Adopting this (conceptual) view, we automatically subscribe to a strong linguistic assumption: since knowing a kind is prerequisite for knowing an object, linguistic reference to an object must always go together with kind reference (cf. Macnamara et al. 1994). I would like to speak of “dual reference” in this connection advocating a theoretical position called sortalism (cf. Blok, Newman & Rips 2005). A specific formalisation will be presented in chapter 7.

Against the dual reference of object terms, the contradiction with which my analysis of the sentences in (21) seems to be confronted disappears. The subject-NPs of these sentences refer to some real object (here always the same one: Fido). Because, due to the principle of dual reference, reference to an object always implies reference to a kind, the subject-NPs at the same time refer to some kind, namely to the kind as an instance of which the object (Fido) is being categorised, or perspectivised, in the relevant utterance situation. In the most general case this can be the kind ‘object’. It is, however, more likely that the “underlying kind” (Macnamara et al.) is further specified by world knowledge and/or contextual knowledge. Presumably, the actual kind underlying the referencing of Fido is the kind ‘dog’ or ‘neighbour dog’. Due to this dual reference, once to an object and once to an underlying kind, it is possible that the topic of the utterance is an object, namely Fido, but the predication nevertheless targets a kind, namely the kind underlying Fido.

How do we know which exactly kind it is that underlies the object referent in, for example, (21a). To come closer to an answer, let us recall what the sentence expresses, namely that Fido is an object instance of the kind ‘dog’. The pragmatic effect of the kind-level predication in (21a) is, I propose, a *specification* of the kind X underlying the object Fido to the kind ‘dog’. The communicative function of uttering the sentence is to make the kind as an instance of which the object Fido is known to the hearer prior to the utterance more specific. It is plausible to think of “making a kind more specific” as replacing the kind by some subkind of it. Given this, an utterance of (21a) is informative only if the hearer was, at the time of utterance, acquainted with Fido only as an instance of some superkind of ‘dog’ (e.g. as a living being, animal, pet or something like that), but not as an instance of ‘dog’ (let alone as a poodle or Australian poodle or something like that).

This analysis proposed for nominal predicates like *is a dog* easily extends to adjectival predicates like *is intelligent* or complex predicates like *has four legs*. Let us assume that the initial underlying kind is the kind ‘animal’. Upon hearing (21a) the state of knowledge of the hearer changes from knowing that Fido instantiates the kind ‘animal’ to knowing that Fido instantiates the kind ‘dog’. Note that the new knowledge state includes the old knowledge state. Virtually the same can be said with respect to (21b) und (21c). Also in these cases the kind initially underlying Fido is being further specified: (21b) leads to a change from knowing that Fido instantiates the kind ‘animal’ to knowing that Fido instantiates the kind ‘intelligent animal’; with respect to (21c) knowing that Fido instantiates the kind ‘animal’ changes to knowing that Fido instantiates the kind ‘animal having four legs’. Thus, also by uttering adjectival or complex predications like (21b) or (21c) the speaker invokes a replacement within the hearer’s knowledge system: the kind underlying a certain object is replaced by some subkind of it. After all, ‘intelligent animal’ and ‘animal having four legs’ are subkinds of the kind ‘animal’.

An interesting question that comes to mind is how to analyse quantificational subjects, as in the following example:

(22) *Each dog is intelligent.*

The kind-level interpretation (roughly: ‘each dog kind is an intelligent animal kind’) is self-explanatory: the quantifier ranges over kinds. The object-level interpretation (roughly: ‘each member of the dog kind is intelligent due to its dog kind membership’) is more interesting: Here the predicate assigns its kind property to the kind underlying the objects that the quantifier applies to. A precise investigation of such cases is desirable, but clearly beyond the scope of this work.

To sum up: It comes as no surprise that kind predicates can combine with kind terms. Besides that kind predicates may, more surprisingly, also combine with object terms. The reason is the nature of kinds being sortal concepts: because there can be no object individuation without categorisation, for every object to which I (linguistically or non-linguistically) refer, there must be some kind as an instance of which I perspectivise the object in the given situation. If an object term appears in the argument slot of a kind-level predicate, the kind-level property will be assigned to this underlying kind. In the case of categorising sentences the semantically underspecified (but pragmatically familiar) kind initially underlying the subject term is being identified with the kind referred to by the postcopular NP.

6.7 Carlsonian stage-level predicates

In contrast to stative predicates like *be a dog*, *be intelligent* or *have four legs*, predicates such as *smoke*, *bark* or *drink Whiskey* (which have been excluded from the discussion up to now) describe dynamic states of affairs. In English, most of them are verbs.

According to Carlson (1977), these predicates are special because their lexical property can be true or false of spatiotemporally bounded entities only. Because it is stages which serve as spatiotemporal entities within Carlson's ontology³, these predicates are dubbed "stage-level predicates". Having this special kind of meaning, Carlson argues, these predicates can be used to assign properties to stages.

On the basis of the distinction between spatiotemporal and non-spatiotemporal entities Carlson explains generic predications as being those predications involving property assignments to non-spatiotemporal entities (individual-level predications) and episodic predications as being those predications involving property assignments to spatiotemporal entities (stage-level predications). While I agree with Carlson in treating generic expressions as non-spatiotemporally denoting expressions, I disagree with respect to the ontological background assumptions. The main difference from Carlson's ontology is that, for me, objects are spatiotemporal entities – making stages dispensable. The only non-spatiotemporal entities are kinds. Accordingly, generic predications are always property assignments to kinds while episodic predications are property assignments to objects (compare the detailed discussion in chapter 4).

The predicates which have been identified by Carlson as stage-level predicates can all be used to form the grammatical predicate of a generic sentence. Pre-theoretically speaking, a generically used stage-level predicate expresses a habit. To explain the fact that lexically nongeneric stage-level predicates can have a generic meaning Carlson postulates a covert generalisation operator that transforms stage-level predicates into individual-level predicates. (23) shows simple examples of generically used stage-level predicates:

- (23) a. *Fido barks.*
b. *Mary smokes.*
c. *Petr drinks Whiskey.*

In contrast to Carlson I assume that, their special (dynamic) descriptive content notwithstanding, Carlsonian stage-level predicates are listed in the lexicon as property assignments to non-spatiotemporal entities, i.e. as kind-level predicates. In this respect, Carlsonian stage-level predicates pattern on a par

³Contrasting with the non-spatiotemporal objects and kinds, compare chapter 4.

with all other predicates. Arguably, the claim that every predicate is basically generic (kind-level) is supported by the simple fact that every (syntactically non-complex) predicate can be used to assign a property to a kind. If dynamic predicates are used in this way, this leads to sentences expressing rules such as the following ones:

- (24) a. *Dogs bark, but cats miaow.*
b. *A true cowboy smokes.*
c. *The Scotsman drinks Whiskey.*

Moreover, due to the principle of dual reference, dynamic predicates may be used as kind-level predicates in connection with object-referring expressions. This combination amounts to generic sentences like (23). Thus, so far, dynamic predicates can serve exactly the same linguistic purposes as stative predicates can. Both can be used to form generic (i.e. kind-level) predicates possibly true of either kinds or objects. The difference between the two sorts of predicates manifests itself only when they are being used non-generically. Used episodically, only dynamic predicates are capable of referring to spatiotemporally localised *events*. Stative predicates can also actualise spatiotemporally localised meanings but they cannot refer to events⁴.

Let us recapitulate: one important difference between Carlson's original genericity theory and my version of it is that Carlson assumes an ontological three-way distinction (stages versus objects and kinds), while I get along with a two-way distinction (objects versus kinds). The price I have to pay to do without stages is the commitment to the dual-reference-hypothesis: I have to assume that linguistic reference to objects always involves a kind. The price is not particularly high though. Indeed, the dual-reference-hypothesis comes for free if one considers kinds to be sortal concepts. A second difference relates to the fact that for Carlson the (English) lexicon contains predicates assigning properties to spatiotemporal entities. In contrast to that, I want the lexicon to contain kind-level predicates only. After all, every one of the predicates treated as stage-level by Carlson can be used generically. So why not treat them as kind-level in the first place? As a result, every lexical predicate would turn out to carry a generic meaning.

Nonverbal Carlsonian stage-level predicates pose a problem for my approach, at least at first sight:

⁴In fact, the distinction between predicates that can be used to refer to (Davisonian) events and those that cannot does not exactly correspond to the distinction between dynamic predicates and stative predicates. Certain statives like the position verbs *sit*, *stand*, *lie* and also *wait* or *sleep* can be used to refer to particular events (Maienborn 2000, 2003). Bach (1986) speaks of "dynamic states" as opposed to "stative states" in this respect.

- (25) a. *Grandpa is tired.*
 b. *Sharks are visible.* [Carlson 2000]
 c. *Peter is hungry.*
 d. *Firemen are available.*
 e. *Your hostel is over there.*
 f. *Here I am!*
 g. *A window is broken.*
 h. *The problem is solved.*

These predicates appear to be predestined to express spatiotemporally localised property assignments. For this reason, one might hesitate to accept the claim that they are listed in the lexicon as generic predicates.

In each of the examples (25e) and (25f) the nucleus of the predicate is formed by an indexical expression. That indexicals are to be interpreted spatiotemporally comes as no surprise because this is how they are defined: their interpretation builds on aspects of the real-world utterance situation. Since indexicals require a semantic treatment in their own right the respective examples do not question my proposal.

The reason why (25g) and (25h) are interpreted spatiotemporally rather than generically is probably that in these examples, the postcopular expressions represent morphological derivatives (participles) of spatiotemporally localised verb forms – in this case of *break* and *solve*. Given this, the predicates we see in (25g) and (25h) are not basic lexical predicates and, therefore, they are also compatible with my claim that lexical predicates are generally kind-level predicates.

It is true that also the predicates in (25a-d) rarely show up as generic predicates in actual usage. It is also true, however, that their generic use is not excluded:

- (26) a. *Old people are tired.*
 b. *Sharks are visible* (because they are not transparent.)
 c. [Experiences of an animal keeper:] *Young wolves are hungry.*
 d. *A bodyguard worth his money is available 24 hours a day.*

The fact that predicates like *tired* or *hungry* are frequently used episodically does not prove that they are listed in the mental lexicon as episodic predicates. How often one encounters a particular predicate in generic or episodic use, respectively, is a pragmatic question. Relevant for the linguistic system is not *how often* a certain form actualises a certain meaning, but only whether it actualises a certain meaning.

And so, instead of assuming that certain predicates are listed in the lexicon as episodic predicates that may be transformed into generic predicates by

grammatical means, I suggest that we look at things from the opposite direction. Carlsonian stage-level predicates are reinterpretable as (generic) kind-level predicates that may be transformed into episodic predicates given the respective speaker intention. In this respect, they are on a par with every other predicate with the exception of Carlsonian kind-level predicates (see below). Carlson's generalisation operation converting stage-level predicates into individual-level predicates is replaced by a localisation operation converting kind-level predicates into object-level predicates (cf. chapter 7). I furthermore noted that many of the predicates counted as stage-level by Carlson do in fact form a lexical class on their own: when used episodically, they refer to (Davidsonian) events.

To sum up: ignoring Carlsonian kind-level predicates for a moment, we note that every predicate is primarily a kind-level predicate that can accordingly be used to assign a property to some kind or to some object:

- (27) a. *Dogs are not very intelligent.*
 b. *The dolphin is a whale.*
 c. *Young chess players open with E2-E4.*
- (28) a. *The dog Rico is intelligent⁵.*
 b. *Moby Dick is a whale.*
 c. *Kasparov opens with E2-E4.*

6.8 Conceptual restrictions on spatiotemporal localisation

Every lexical predicate is basically a kind-level predicate and may, accordingly, be used generically. Moreover, setting aside Carlsonian kind-level predicates, every basic kind-level predicate may also be taken to episodically express a spatiotemporally localised meaning. As a matter of fact, however, predicates vary with respect to their “willingness” to spatiotemporally localise. The “willingness” to undergo spatiotemporal localisation is only loosely predictable from morphosyntactic form. Let me discuss at least some examples for illustration.

Predicates formed by the copular verb plus an article plus a noun heavily resist a non-generic usage (cf. Carlson 1977). To spatiotemporally localise such a predicate a strong context is usually called for, as in (29a), where it is the fairy tale context that makes the episodic interpretation plausible. It is not by chance that the examples under (29) describe exceptional circumstances – each time the described real-world state of affairs is limited to a particular time, i.e. relative to the spatiotemporal location at hand, different from what we expect it to be according to the conditions as they are determined in our conceptual

⁵Compare Kaminski, Call & Fischer (2004).

knowledge system (where a prince is a human being etc.). The sentences report on non-normal states of affairs:

- (29) a. *The prince is a frog at the moment.*
b. *Leipzig is a huge construction site at the moment.*
c. *This pawn is a queen now.*

(30) exemplifies the case of a sentence where, out of the blue, both of the two interpretations are more or less equally possible:

- (30) *Hinnerk is a taxi-driver.*

Under the generic, categorising reading Hinnerk is understood to be a professional taxi-driver (he is an instance of the kind ‘taxi-driver’). Being a taxi-driver is part of Hinnerk’s “identity”. Under the episodic, object-related reading Hinnerk is understood as working as a taxi-driver without this job necessarily making up his identity. Maybe Hinnerk is a private detective by profession but due to having no clients at the moment, he is forced to temporarily job as a taxi-driver to earn a living. In the first case, *is a taxi-driver* is understood as an “essential” or “defining” property of Hinnerk. In the second case it is understood as an “accidental” or “temporary” property. In many languages these two readings are formally disambiguated, for example in Russian past tense and future copular sentences (e.g. Geist 2004).

Predicates formed by the copular verb plus an adjective are likewise sometimes more, sometimes less easily localisable in space and time:

- (31) a. *John was upset on Wednesday.*
b. *John was intelligent on Wednesday.*⁶
c. *The sky is blue at the moment.*
d. *The bucket is blue at the moment.*

Whereas (31a) and (31c) are acceptable without reservation, (31b) and (31d) require quite strong contextual support. To arrive at an acceptable interpretation, we have to think of non-normal situations that run counter to our conceptual knowledge about intelligence and buckets. (31b) forces us to imagine that John’s intellectual capacities, usually on a low or at most average level, reach a high level during the course of one particular day. Such an short-period peak contradicts our expectations because we know that raising a person’s level of intelligence is a long term project – if possible at all. (31d) forces us to imagine that a bucket changes its colour from time to time. This likewise strains our common reasoning because we know that a bucket is a an artifact that primarily

⁶Adopted from Chierchia 1995.

serves a certain function, its outer appearance being of secondary relevance and, therefore, usually not important.

Verbal predicates do not resist spatiotemporal localisation. They like to denote particular events:

- (32) a. *Fido has been barking for hours now.*
b. *Mary is smoking.*
c. *Kasparov has once again opened with E2-E4.*

There is, however, a certain class of verbs, without exception stative verbs, which lack the lexical-conceptual prerequisites to denote events. With these, again, a spatiotemporal localised interpretation requires sometimes more, sometimes less contextual support. Padučeva (1996) calls them “relation verbs”, typical examples are *depend on*, *consist of*, *contain*, *include*, *involve*, etc.:

- (33) a. *Peter’s fate is depending on Mary’s decision now.*
b. *?The bucket is containing water.*

To sum up: All predicates are kind-level predicates that can be used to express (generic) properties directly true of some kind or indirectly true of a kind underlying some object. To express a (episodic) property true of some object such a basic kind-level predicate has to undergo spatiotemporal localisation. Spatiotemporal localisation is thus an operation transforming kind-level predicates into object-level predicates. By virtue of their lexical-conceptual structure, different predicates more or less readily undergo this operation. This is a pragmatic question in the end: for some predicates, situations in which they fit in under a spatiotemporal interpretation are frequent, for others, such situations are rare. Many stative predicates can actualise only non-eventive spatiotemporal interpretations. Apart from these “stative states” (Bach), the application of spatiotemporal localisation leads to event-denoting predicates.

The characterisation given accounts for all of the predicates that Carlson (1977) counted as stage-level and object-level predicates. Carlsonian kind-level predicates involve a certain lexical peculiarity that justifies a characterisation in their own terms. This will be given the next section.

6.9 What is special about Carlsonian kind-level predicates

I proposed to treat every lexical predicate as a kind-level predicate transformable into an object-level predicate given the appropriate communicative needs. I furthermore proposed that it is exactly those predicates that Carlson (1977) called “kind-level” which are non-typical kind-level predicates. The purpose of this

section is to highlight the peculiarity of Carlsonian kind-level predicates that distinguishes them from all other (“normal”) kind-level predicates.

Krifka et al. (1995) noticed that only kinds can be invented, be extinct, be exterminated etc., ruling out the combination of a Carlsonian kind-level predicate with an object term – unless under the representative object interpretation. From the perspective adopted in this dissertation this means that Carlsonian kind-level predicates, unlike all other (“normal”) predicates, cannot be spatiotemporally localised. We now ask why.

The lexical semantics of Carlsonian kind-level predicates has been discussed in some detail in chapter 5. Let me briefly repeat the results: In the case of a normal kind-level predicate, the predicate’s property is assigned to the kind delivered by the argument of the predicate and that is all there is to say. Normal predicates assign, so to speak, pure kind-level properties. The property assigned by a Carlsonian kind-level predicate, by contrast, is not a pure kind-level property: although the predication syntactically targets a kind, it semantically involves not only information about the kind denoted by the argument expression but also information about instances of the kind. This has consequences for the truth conditions.

The truth or falsity of a normal kind-level predication is fully determinable by evaluating the conditions holding in the kind domain⁷. If the conditions asserted by a normal kind-level predication map the conditions holding in the kind domain, the respective sentence will be true. The truth or falsity of a Carlsonian kind-level predication, by contrast, is determined not only relative to the kind domain, but also relative to the object domain. To judge whether a Carlsonian kind-level predication is true or false one has to take into consideration the relations holding among kinds as well as the relations holding among objects. Let us look at a classic example:

(34) *The dodo is a bird.*

What are the conditions for this sentence representing a normal kind-level predication to be true? It is true at time t_0 if the following holds at t_0 :

- (i) Within the kind domain, there is the kind ‘dodo’.
- (ii) Within the kind domain, there is the kind ‘bird’.
- (iii) The kind ‘dodo’ is a subkind of the kind ‘bird’⁸.

As can be seen, every one of the three truth conditions (i) to (iii) is concerned with a state holding at t_0 in the kind domain. For this reason, the truth of a

⁷Compare the discussion in connection with example (20a).

⁸This is equivalent with: ‘dodo’ is an element of the kind extension of the predicate *bird* and ‘dodo’ is not identical to the kind ‘bird’.

normal kind-level predication like (33) is independent of any state of affairs holding in the object domain (=in the real world) at t_0 . This is the reason why the sentence is true at the present moment of time even though at present, there is not a single dodo. Compare this with a Carlsonian kind-level predication:

(35) *The dodo is extinct.*

This sentence can be true at t_0 only if, at t_0 , the following conditions are fulfilled⁹:

- (i) Within the kind domain, there is the kind ‘dodo’.
- (ii) Within the object domain, there are no object instances of the kind ‘dodo’.
- (iii) Within the object domain, object instances of the kind ‘dodo’ existed at some time prior to t_0 .

As (ii) and (iii) show, this time the truth of the sentence is depending also on conditions holding in the object domain¹⁰.

So why is it that Carlsonian kind-level predicates cannot be spatiotemporally localised? The answer is a little surprising: the meaning of a Carlsonian kind-level predicate cannot be spatiotemporally localised because *it already is spatiotemporally localised*. The lexical-semantic structure of a Carlsonian kind-level predicate involves an assertion or presupposition, respectively, concerning the existence of object instances of the kind that is delivered by the expression filling the argument slot of the predicate.

It is helpful to contrast the present proposal with the classic Carlsonian analysis of kind-level predications. (36a) shows a Carlsonian semantics for *are extinct*, (36b) for *dodos* and (36c) for the combination of these two expressions¹¹:

- (36) a. $[[\textit{are extinct}]] = \lambda x_k. \text{EXTINCT}(x_k)$
 b. $[[\textit{dodos}]] = \text{DODO} \in D_k$
 c. $[[\textit{dodos are extinct}]] = \text{EXTINCT}(\text{DODO})$

⁹Note that I do not say that these three conditions were sufficient conditions for the truth of (33).

¹⁰Condition (iii) implies that, within the kind domain, there must have been the kind ‘dodo’ at some time prior to t_0 . This is the reason why it is impossible to use an indefinite NP to refer to the kind ‘dodo’ in (35), unless under the taxonomic reading (cf. chapter 5).

¹¹Subscript “k” marks that a variable is restricted to kinds. For objects, subscript “o” will be used.

In contrast to that, the present proposal amounts to the following semantic representations¹²:

- (37) a. $[[\textit{are extinct}]]$ = $\lambda x_k \exists x_o. \text{loc}(x_o, x_k) \& \text{VANISHED}(x_o)$
 b. $[[\textit{dodos}]]$ = $\text{DODO} \in D_k$
 c. $[[\textit{dodos are extinct}]]$ = $\exists x_o. \text{loc}(x_o, \text{DODO}) \& \text{VANISHED}(x_o)$

Saying that Carlsonian kind-level predicates initially have a spatiotemporally localised meaning I do not say that they would be camouflaged object-level predicates. (37a) identifies them as kind-level predicates with the special feature that the kind property entails the existence (or, in the case of *be extinct*, the former existence) of real objects instantiating the kind. In this respect they differ from normal or “pure” kind-level predicates, compare (38)¹³:

- (38) a. $[[\textit{is a bird}]]$ = $\lambda x_k. \text{BIRD}(x_k)$
 b. $[[\textit{the dodo}]]$ = $\text{DODO} \in D_k$
 c. $[[\textit{the dodo is a bird}]]$ = $\text{BIRD}(\text{DODO})$

Thus, I assume that every “normal” predicate is stored in the lexicon as a kind-level predicate and that an operator called *loc* (for “spatiotemporal localisation”) is made available by the language system to derive object-level predicates from kind-level predicates. I furthermore assume that the semantic effect of the application of *loc* to some kind-level predicate is that the property initially suited for kinds becomes suited for objects. Given these assumptions and the semantics for Carlsonian kind-level predicates as indicated in (37a), it is obvious why Carlsonian kind-level predicates cannot be spatiotemporally localised: the property of having object instances being rare, widespread, extinct, ... or invented can only be true of kinds. By definition, only kinds can have instances, instances cannot have instances. Therefore, the property of having object instances cannot be adapted to objects.

6.10 Carlsonian kind-level predicates as existence predicates

According to the semantic analysis outlined in the previous section a Carlsonian kind-level predicate is a predicate that ascribes to a kind the property of having instances (or having no longer instances). In this respect a kind-level predicate resembles existence predicates like the English *there be*. At least if we follow McNally (1997, 2005) assuming that the existence predicate corresponds to a

¹²The relation “*loc*” is a close relative of Carlson’s realisation relation “*R*”. The difference is only that *loc* does not relate stages to individuals, but objects to kinds.

¹³I abstract away from the semantics of the definite article in (38b).

function mapping entity correlates of properties onto a truth value such that a true proposition results in case the property-entity has instances.

McNally (2005) cites the philosopher Strawson who speculates that an existence sentence might be “a subject-predicate proposition in which the subject is a property or concept and in which the predicate declares, or denies, its instantiation”. If we, taking McNally’s approach further, consider “properties or concepts” to be kinds, Carlsonian kind-level predicates completely turn out to be specimens of the existence predicate, compare (39) with (37):

- (39) a. $[[\textit{there be}]] = \lambda x_k \exists x_o. \text{loc}(x_o, x_k)$
 b. $[[\textit{dodos}]] = \text{DODO} \in D_k$
 c. $[[\textit{there are dodos}]] = \exists x_o. \text{loc}(x_o, \text{DODO})$

The difference between *there be* and *be extinct* lies only in the additional lexical semantic condition that, in the latter case, is being imposed on the asserted instances and that I provisionally formulated as $\text{VANISHED}(x_o)$. Other Carlsonian kind-level predicates are to be analysed along virtually the same lines as indicated in the following semantic representations:

- (40) a. $[[\textit{be rare}]] = \lambda x_k \exists x_o. \text{loc}(x_o, x_k) \& \text{LOW-IN-NUMBER}(x_o)$
 b. $[[\textit{be numerous}]] = \lambda x_k \exists x_o. \text{loc}(x_o, x_k) \& \text{HIGH-IN-NUMBER}(x_o)$
 c. $[[\textit{be invented}]] = \lambda x_k \exists x_o \exists e. \text{loc}(x_o, x_k) \& \text{CREATE}(e) \& \text{UNDERGOER}(e, x_o)$
 d. $[[\textit{be widespread}]] = \lambda x_k \exists x_o. \text{loc}(x_o, x_k) \& \text{FAR-AWAY-FROM-EACH-OTHER}(x_o) \& \text{HIGH-IN-NUMBER}(x_o)$

The results arrived at are summarised in figure 6.2.

predicate	pure kind-level (pot. object-level)	“Carlsonian stage-level” “Carlsonian object-level”
	existence predicate	<i>there be</i> “Carlsonian kind-level”

Figure 6.2: revised predicate classification

On the one hand, there is the class of “pure” kind-level predicates from which, by means of the operation of spatiotemporal localisation, object-level predicates may be derived given the appropriate communicative needs. Normal kind-level predicates correspond to Carlsonian object-level and stage-level predicates. Note that, according to this analysis, *every* predicate, except for Carlsonian kind-level predicates, turns out to have a hybrid nature just like a “kind predicate” as identified by Krifka et al. (1995:10). On the other hand, there are the existence predicates which are likewise kind-level predicates but which

are not pure because their properties involve information about instances. The prototypical case is represented by the existence predicate *there be* which ascribes to a kind the property of having instances in the real world. Carlsonian kind-level predicates are likewise existence predicates. While *exist* is used to assert the “naked existence” of instances of some kind, a Carlsonian kind-level predicate is used to additionally inform about a certain mode of existence of instances of some kind.

In (40) I presented provisional semantic representations for some Carlsonian kind-level predicates. It is beyond the scope of this work to elaborate detailed lexical-semantic analyses of the various Carlsonian kind-level predicates. The semantics suggested for (40c), for example, is in fact too coarse to differentiate *be invented* from *be bred*. A more appropriate semantic analysis would have to additionally reflect the mode of creation of the new kind – that in the case of *breed*, unlike in the case of *invent*, the new kind results from crossing two already existing natural kinds etc.

In accordance with this remark, figure 6.3 is supposed to be the embryo of a lexical taxonomy of Carlsonian kind-level predicates to be further developed in future research. I take it to be self explaining (“CKP” is for “Carlsonian kind-level predicates”).

“CKP”	appearance	<i>invent, breed, create, discover, infest,...</i>
	disappearance	<i>exterminate, die out,...</i>
	spatial/numerical distribution	<i>be rare, be widespread, be numerous,...</i>
	gestalt variation	<i>come in different forms, grow bigger as one travels north,...</i>

Figure 6.3: cartography of Carlsonian kind-level predicates

6.11 Conclusions

Although every predicate is basically registered in the lexicon as a kind-level predicate, they not only allow for kind-referring arguments but also for object-referring arguments. For an object term to be the argument of a kind-level predicate, the denoted object must be interpretable as a representative of a kind so that the kind property of the predicate can be assigned to the kind represented by the object.

One has to distinguish between those cases where the object has gained the status of a kind's representative already prior to the utterance (41a) from those cases where an object is gaining the status of a kind's representative by virtue of the utterance. In the latter case one can further distinguish between the categorisation of an object as a representative of some kind already active in discourse (41b) and the categorisation of an object as a representative of some "reactivated" or even novel kind (41c). Arguably, in (41d), where the postcopular constituent is not formed by a noun phrase, the object is likewise categorised as a representative of a novel kind (here: the kind of blind individuals).

- (41) a. [newspaper headline:] *Moby Dick close to extinction*
 b. *Congratulations, you are the winner of the jackpot!*
 c. *"I am a Christian Buddhist"* (S. Deisler, German football star)
 d. *Mary is blind.*

The most surprising and most interesting discovery of this chapter is probably that the representative object interpretation of an object term is not a marginal phenomenon, but something central for the understanding of generic predications. This is explained by the nature of objects and kinds taken for granted in this work: by definition, objects are entities representing (=instantiating) kinds in the real world. From this it follows that we cannot refer to objects without perspectivising the object as a representative of one kind or the other. Semantically, this "underlying" (Macnamara) or "corresponding" (Dölling) kind is available to saturate the argument variable when the kind predicate syntactically combines with an object term. It is now understandable why the subject of a generic sentence, i.e. of a kind predication, may not only be formed by a kind term, but also by an object term.

To meet the objections that have been raised against the original Carlsonian predicate classification (e.g. by Croft 1986), we argued for a revised classification of natural language predicates where kind-level predicates represent the normal case. Normal kind-level predicates can be converted into object-level predicates by means of the spatiotemporal localisation operation (this will be the topic of the next chapter). It comes as no surprise that a spatiotemporally localised predicate is compatible with an object term.

Carlsonian kind-level predicates, by contrast, cannot be spatiotemporally localised. Together with the existence predicate they form a lexical predicate class in their own right. This class is characterised by that the described property involves existence conditions imposed on object instances of a kind. One can say that an object-level predication (which is a spatiotemporally localised kind-level predication) states about an object *how* it is existing in the real world of objects. An existence predication states about a kind *that* it is existing in the real world

of objects. A Carlsonian kind-level predication states about a kind *that and how* it is existing in the real world of objects.

Chapter 7

Overt and covert articles – a kind-based DRT approach to the semantics of noun phrases

7.1 Overview

The basic assumption taken for granted in this work is the assumption that the kind domain represents the fundamental ontological domain in relation to which nouns, verbs and adjectives determine their semantic content. If object-level reference is intended the meaning of a basic lexical predicate will have to undergo a certain grammatical operation. The aim of this chapter is to present one more important argument in support of this view. It concerns the semantic composition of NPs in general and of bare plural NPs in particular.

Section 2 determines my starting point for the discussion: the linguistic function of an article is to transform predicative expressions into referential expressions. I adopt the version of DRT recently presented by Farkas & de Swart (2003) to formalise this approach. However, before I turn to the semantics of bare nominals I insert three sections on predicate modification. After some general remarks on the formation of complex predicates (section 3), I point to the necessity of distinguishing between two modes of modification: internal or type-restricting modification, and external or token-restricting modification (section 4). Prenominal adjectives are assumed to always represent internal modifiers (section 5). In section 6, returning to the main topic of the chapter, I hypothesise that noun phrases (NPs) are always referential, never predicative. As a consequence, syntactic argument positions (NP-positions) are reserved for referential expressions; if a predicate appears in argument position, a referential expression must covertly derived from it (section 7). In section 8 I use DRT-representations to demonstrate that, once we accept all these assumptions, it is (in line with Carlson's 1977 reasoning) predicted that bare NPs are basically

kind-referring expressions. The fact that, besides the predicted generic reading, bare NPs actualise existential readings is addressed and explained (resorting to an idea put forward by Chierchia 1998) in section 9. Finally, in section 10, I demonstrate that even more support for the proposal I made comes from the deviant behaviour of bare plural NPs of the type *parts of that machine*. That in these cases the object-level interpretation of the bare plural NP is possible is due to the bare plural predicate being modified by a token-restricting modifier. Section 11 adds a brief discussion of how Chierchia and Krifka attempt to explain the bare plural facts. Section 12 is included to discuss proper names in the light of the present approach; unlike bare plural NPs, which are pure kind-level expressions, proper names by themselves always denote also at the object-level. Section 13 summarises the chapter.

7.2 The grammatical role of overt articles

When discussing the inner architecture of the kind domain in chapter 3, we were led to conclude that a kind-level noun is a janus-faced creature: on the one hand, it is a term naming a certain kind; on the other hand, at the same time, it is a predicate characterising a set of kinds. The noun *dog*, for example, names the kind ‘dog’ and denotes the set of kinds ‘dog’, ‘collie’, ‘dachshund’, ‘poodle’, ‘sledge dog’, ‘good dog’, ‘bad dog’, ‘straying dog’ and whatever more dog kinds one can think of. A closer look reveals, however, that there is a problem with treating a noun such as *dog* as a name of a kind. If *dog* was truly the name of a kind, we would expect the following sentence to be adequate, contrary to fact:

- (1) **Dog evolved from the wolf.*

To express the intended content, one has to choose one of the following options:

- (2) a. *The dog evolved from the wolf.*
 b. *A dog evolved from the wolf.*
 c. *The dogs evolved from the wolf.*
 d. *Dogs evolved from the wolf.*

Why is (1) excluded? Two answers are possible. If one wanted to keep with the idea that count nouns are listed in the lexicon as names of kinds, one would have to resort to a purely syntactic rule to explain the ungrammaticality of (1). Consequently, the derivational step from *dog* as in (1) to *the dog* as in (2) is viewed as not involving any change in semantic content. This kind of explanation is offered, for example, by Dölling (1992:60). Alternatively, one could assume that nouns are no *names* of kinds, but rather *symbols* of kinds (e.g. Foley

1998). Conceptually, a nominal like *dog* is a symbol of a kind, linguistically, however, it is a kind-level predicate. Under this assumption, the following frequently adopted explanation for (1) suggests itself: The subject of a declarative sentence must be formed by a referential expression¹. Lexical nouns are no referential expressions, however. Therefore, to be able to function as a referential expression, a noun like *dog* must undergo a certain semantic operation. In languages like English, this semantic operation manifests itself in the form of overt articles as in (2a)-(2c).

Padučeva (1979) argues that even in a language like Russian, that lacks article words, every nominal which is used as a referential expression necessarily consists of two meaning components. According to Padučeva (1979), the first semantic component is a certain property, while the second one is an instruction of how to determine a referent on the basis of this property within the context of an utterance. Lexically, a Russian noun is supposed to possess only the first semantic component:

Обščee imja ničego ne oboznačae: ono imeet smysl, t.e. vyražae nekotoroje svojstvo, no ne denotat. (Padučeva 1979:26)

A common noun does not refer at all: it has a sense, that is, it expresses a certain property, but it has no denotation. (own translation)

It is only if the the common noun combines with *a*, as she calls it, “actualiser” that the second meaning component comes into play, leading to a referential expression. Russian bare nominals functioning referentially are explained by the presence of some covert actualiser:

Aktualizator v sostave IG možeet byt’ nulevym. Denotativnyj status IG opredeljaetsja v étom slučae sintaksičeskim kontekstom, linejno-intonacionnoj strukturoj predloženiija i pročim. (Padučeva 1979:26)

The actualiser within a noun phrase may be a zero. In this case, the particular denotation of the noun phrase is determined by syntactic context, word order, prosody etc. (own translation)

I opt for the second possibility, thereby committing myself to two assumptions: First, a common noun is lexically a predicate. Secondly, the syntax (which is largely left unspecified in this work) provides argument positions, with the subject position being a case in point. Argument positions are reserved for referential (or quantificational, see footnote 1) expressions, syntactically categorised as NP (or DP, respectively). Those expressions that are NPs by themselves,

¹Throughout this work I abstract away from generalised quantifiers for ease of simplicity.

like proper names or certain article-noun-syntagms, are suited to fill such an argument position. However, expressions that are not NPs by themselves may occupy argument positions, too. A necessary condition for this to be possible is that the mismatch between the requirement of the syntactic position and the meaning format of the non-NP is dissolvable. In other words, the necessary condition is the availability of a type shifter which is able to turn the meaning format of the expression occupying the argument position into the meaning format associated with an NP.

Given that the lexical meaning of a content word is generally a kind-level meaning², we should consider an English noun to be listed in the lexicon as a kind-level predicate. Let D_k be the domain of kinds. For the nouns *dog* or *dogs*, respectively, this amounts to the semantic representation (3):³

- (3) a. $[[dog]] = \lambda X \in D_k. DOG(X)$
 b. $[[dogs]] = \lambda X \in D_k. DOG(X) \ \& \ PLURAL(X)$

Referential expressions (terms) result from the combination of such a kind-level predicate with a functor of which the semantic function it is to map the extension of the kind-level predicate onto one of its elements. This output element is the referent of the term. In theory, every subset of the kind extension is a possible referent of the derived term. In language-particular practice, however, the range of possible referents is morphosyntactically constrained. Two grammatical factors are relevant for the English noun: the category of number and the articles.

In the following presentation of the derivation of a referential NP, I use Farkas & de Swart's (2003) version of the discourse representation theory DRT as originally developed in Kamp & Reyle (1993). I modify it by installing a systematic distinction between kind-level arguments (symbolised by capital letters) and object-level arguments (symbolised by small letters).

The lexical meaning of the (kind-level) predicate *dog* corresponds to the discourse representation structure (DRS) in (4), the lexical meaning of the pluralised predicate *dogs* corresponds to the DRS in (5). Note that the DRSs in (4) and (5) do not contain any reference markers (=discourse referents) – the universe of each DRS is empty:

- (4) $[_N dog]$

²Note that proper names do not fall under this rule.

³I follow Farkas & de Swart's (2003) in modeling morphological pluralisation as an additional predication.

dog(X)

(5) [N *dogs*]

dog(X)
plural(X)

Only when these nouns are syntactically combined with articles to form NPs, their meanings are being – as Farkas & de Swart call it – “instantiated”. Technically speaking, instantiation is the replacement of a thematic argument X by a reference marker U. Note, by the way, that the elements bringing about “instantiation” are exactly the elements that Padučeva calls “actualisers”.

Farkas & de Swart consider the indefinite article to be the simplest of all possible instantiators because it semantically contributes nothing but the pure replacement of X (Y, Z, ...) by U (V, W, ...). Note, however, that the applicability of the indefinite article is lexically constrained: its semantic function is defined only for those nouns of which the extension includes exclusively atomic elements. In other words, the indefinite article is compatible only with nouns whose number value is singular.⁴

The semantic impact of the indefinite article leads us from the semantic representation in (4) to the one in (6), that corresponds to the indefinite NP *a dog*:

(6) [NP *a dog*]

U
dog(U)


How to precisely formulate the semantics of the *definite* article is left open by Farkas & de Swart, but Geurts (1999) made a useful proposal for analysis. Geurts adopts the view that the presence of a definite article in syntactic structure triggers a certain presupposition. The definite article is, simply speaking, the interpretive advice that the reference marker introduced by the definite

⁴This is the situation in English (and German). As is well known, other languages like French have a plural indefinite article at their disposal.

NP at hand must be bound to (co-indexed with) another reference marker already introduced into the discourse beforehand. If there is no appropriate pre-established reference marker to which the reference marker of the NP could be bound, such a reference marker must be accommodated. Geurts' proposal neatly captures the well-known observation that the definite article is used in those cases where the referent is "familiar".


Thus, the application of the definite article leads to the integration of the DRS induced by the NP to be computed *the dog* (little box) into the DRS previously built up induced by the pretext (big box). This is indicated in (7):

(7) [NP *the dog*]



Unlike the indefinite article, the English definite article freely combines with singular as well as with plural nouns:

(8) [NP *the dogs*]



What happens if reference to some non-familiar plurality is intended? This possible target of reference is not covered by the two lexically available instantiators (=articles). Hence, expressions capable of referring to non-familiar pluralities must be derived covertly. This raises the question of what the theoretical status of the English "null article" is, compare the subject of (2d). However, before I enter the difficult field of bare plurals, I would like to discuss some aspects of predicate modification that will be relevant for my line of argumentation. Predi-

cate modification is the syntactic combination of two or more predicates to form a complex predicate prior to the application of an (overt or covert) article.

7.3 Forming complex predicates

Following Dayal (2004), I assume that English articles are semantically underspecified with respect to the kind-/object-distinction. Speaking in formal semantic terms, the definite article is capable of mapping sets of objects as well as sets of kinds onto their respective maximal element, and the indefinite article is capable of mapping sets of objects or kinds onto an arbitrary element. Speaking in terms of DRT, the (in)definite article is capable of instantiating a thematic object argument as well as a thematic kind argument. For this reason, English NPs are systematically ambiguous between object reference and kind reference.

It frequently happens that the extension which is characterised by a particular lexical noun (be it kinds or objects) includes more entities than the speaker wants it to include, given her specific communicative goal in the specific context of utterance. To restrict the extension, the speaker has to impose an additional condition on the potential referent. This is achieved by syntactically combining the basic nominal predicate with a second predicate, in English prototypically with an adjective (see Croft 2001). On the semantic plane this additional predication leads from a less specific descriptive content of the basic predicate, that is, of the head noun, to the more specific description of the newly formed predicate-predicate combination, for example of an adjective-noun-syntagm.

Note that the well-known relationship between description (intension) and denotation (extension) holds not only for objects but also for kinds: the more specific the description, the less entities are likely to fall within the denotation. Consider the following example as an illustration:




Figure 7.1: nice dogs and bad dogs

Let us assume that this taxonomy represents a partial structure of the kind domain D_k of a given universe of discourse U . Thus, D_k includes among others the entities ‘animal’ (=A), ‘pig’ (=P), ‘dog’ (=D), ‘nice dog’ (=N) and ‘bad dog’

(=B). Let us furthermore assume that within the object domain D_o of U there are the three entities Piggy (=p), Fido (=f) and Bello (=b) and that the following relations hold between D_o and D_k : $\mathbf{R}(p,P)$, $\mathbf{R}(f,N)$ and $\mathbf{R}(b,B)$ ⁵: Piggy is an exemplar of the kind ‘pig’, Fido is an exemplar of the kind ‘good dog’ and Bello is an exemplar of the kind ‘bad dog’.

Given these assumptions, the noun *dog* denotes the set $\{f,b\}$ relative to D_o , while the modified nominal *good dog* denotes the smaller set $\{f\}$. Enriching the semantic content of *dog* with the content of *good* invokes a restriction of the extension. Bello fulfills the conditions to belong to the extension of the noun *dog* but not the conditions to belong to the extension of the modified noun *good dog*. How about kind-level predicates?

Relative to D_k , *dog* denotes the set $\{D, B, N\}$. The expression *bad dog*, by contrast, characterises the singleton set $\{B\}$. As can be seen, the interaction between intension and extension known from the domain of objects does also hold within the domain of kinds. Specifying a concept brings about the (potential) restriction of the set of entities falling under the concept. With respect to the kind domain, this interplay can be described as follows: if a N is modified to form a N^* everything will be “cut off” from the taxonomic tree of which the root is N except for the tree of which the root is N^* ⁶.

In theory, the modification of a basic predicate can be iterated arbitrarily often. The practical goals of communication, however, terminate subsequent modification in a natural way. After all, the speaker chooses to make use of a certain description in order to draw the hearer’s attention to a certain entity, be it a kind or an object. At some point the time has come to signal the hearer that she received enough information about the potential referent to successfully identify the actual referent within the context of utterance. The semantic step from a predicate characterising a set of entities to a term referring to a certain entity corresponds syntactically to the step from N or N^* , respectively, to NP ⁷. In terms of semantic types: when projecting from N to NP , an expression of the semantic type $\langle e,t \rangle$ is transformed into an expression of the semantic type $\langle e \rangle$.

As noted above, article languages like English or German have at their dis-

⁵“[L]et \mathbf{R} be the realization relation which relates kinds to their specimens [...]. Thus a formula $\mathbf{R}(x,k)$ states that the object x belongs to the kind k ” (Krifka et al. 1995:66). It follows from principles (3) and (4) in chapter 3 that the following relations must hold, too: $\mathbf{R}(f,D)$, $\mathbf{R}(f,A)$, $\mathbf{R}(p,A)$, $\mathbf{R}(b,D)$, $\mathbf{R}(b,A)$.

⁶I indicate modified expressions by *. If X is the syntactic category of a basic lexical predicate a , then X^* represents the expression b belonging to the same syntactic category as a , differing only in that it has undergone modification. Thus, the * is, strictly speaking, no syntactic diacritic.

⁷...or from NP to DP .

positional lexicalised functor elements which map predicates onto terms, namely the definite and the indefinite article⁸. From a functional perspective, we can think of the articles as linguistic symbols by means of which the speaker signals that the derivation of a possibly complex predicate on the basis of some head noun has come to an end, and that, from now on, it is the hearer's task to use the resulting description in order to find the intended referent within the context of utterance. The semantic contribution of the article terminates the nominal projection. The result is what is called a descriptive NP.

Let me summarise the basic facts that I take for granted in the discussion to follow. First, basic lexical expressions (content words) are predicates (semantic type $\langle e, t \rangle$). With respect to the grammar of English and many other languages, it is justified to give a special syntactic category label, N, to a special class of content words, namely those describing sortal concepts, i.e. nouns. Secondly, referential expressions result when a functor element (semantic type $\langle \langle e, t \rangle, e \rangle$) applies to a (possibly modified) lexical predicate. I take the syntactic category of such a referential expression to be NP (others prefer DP). Thirdly, English resembles German in that it possesses two lexicalised function words serving as “instantiators” (in the sense of Farkas & de Swart), the definite and the indefinite article. The application of an article terminates the derivation of a referential expression. Prior to the application of the article the nominal predicate may be modified by other predicates, e.g. by adjectives. At least, an NP may generally refer either to a kind (or a set of kinds), or to an object (or set of objects).

7.4 Internal and external modification

English definite and indefinite NPs are ambiguous – they allow for an object-related and for a kind-related interpretation (e.g. Krifka et al. 1995). Traditionally, it is often the article which is made responsible for the ambiguity; often, a special generic definite article is postulated which is supposed to be homomorphic to the “ordinary” definite article. Following Dölling (1992) and Dayal (2004), however, I assume that it is not the article which is the locus of the ambiguity but rather the nominal. This theoretical position has an impact on the interaction between complex predicate formation on the one hand, and on object-level versus kind-level denotation on the other hand. This is what I turn to now.


⁸Under the assumption that the definite article corresponds semantically to the iota-operator and that the indefinite article corresponds to a choice function, both articles are of semantic type $\langle \langle e, t \rangle, e \rangle$. This is, of course, not the only semantic analysis that has been suggested in the literature.

In this work I defend the thesis that basic lexical predicates are kind-level predicates which by default turn into kind terms as they compose with an article. Object terms are always the result of converting a basic kind-level predicate into an object-level predicate via spatiotemporal localisation prior to the application of the article. Consider (9) and (10) for illustration. The generic-referring subject in (9a) corresponds to the syntactic structure represented in (10a), and the particular-referring subject in (9b) corresponds to the syntactic structure in (10b):


(9) a. *The dog is man's best friend.*

b. *The dog is sitting on my place.*

(10) a.

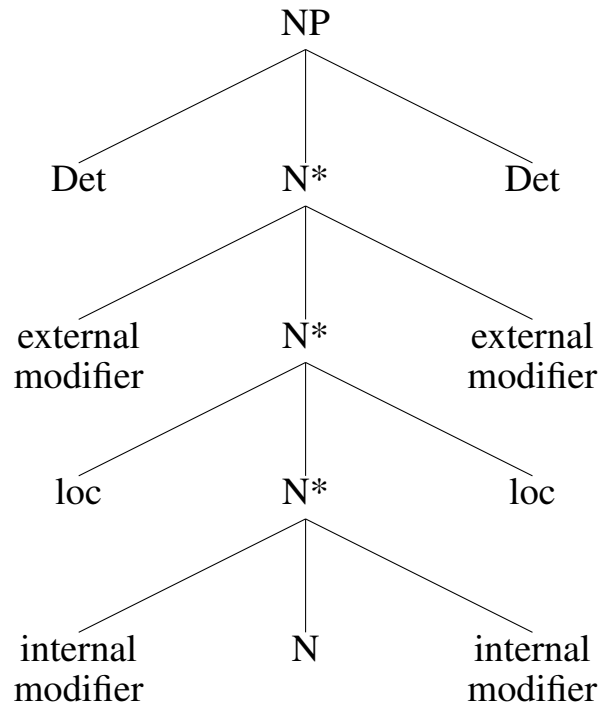


b.



These syntactic assumptions have important implications for the process of building up complex predicates. As has been noted in section 3, the modification of the descriptive content of a lexical head generally takes place prior to article application. Given this and given that there is an optional step *loc* prior to the stage when the article terminates the derivation, as indicated in (10), we expect that predicate modification can take place prior to *loc* or after *loc*. Modification prior to *loc* corresponds to semantic composition at the kind-level, whereas modification after *loc* corresponds to composition at the object-level. This is summarised in (11). Internal and external modification correspond to Bolinger's (1967) "reference modification" and "referent modification", respectively. Modification prior to spatiotemporal localisation can be called "internal modification"; modification afterward can be called "external modification".

(11)



The terminology “internal/external modification” is adopted from Maienborn (2001) who, investigating German, detects two different syntactic positions for local adverbials to modify verbal meanings, each position being associated with its own mode of modification⁹. For illustration, compare the following sentence:

- (12) *Paula schläft auf dem Bauch*
Paula is-sleeping on the belly
‘Paula is sleeping on her belly’


If understood episodically, the German sentence allows for two readings depending on the interpretation of the local adverbial *auf dem Bauch*. If the adverbial modifies event-internally, the local relation expressed by the prepositional phrase will be integrated into the description of the sleeping event; the relatum (landmark) belly will be understood as Paula’s belly, and the sentence reports on a situation where Paula is sleeping on her (own) belly. If the adverbial modifies event-externally, however, the whole sleeping event will be located within the auf-region of the relatum belly. As event participants belong to the event they participate in, the belly on which the event is located cannot be Paula’s. Therefore, if *auf dem Bauch* functions as an external modifier, (12) describes a situation in which Paula is sleeping on someone else’s belly.

Making use of the type/token-distinction, we can explain the availability of the two readings in (12) as follows. In the case of internal modification the local

⁹Maienborn still identifies a third syntactic position (not relevant for us here), namely the position for framesetting local adverbials.

adverbial functions as a type-level modifier – the event type ‘sleep’ is specified as to the event type ‘sleep on the belly’. The sentence reports on an instance (token) of the event type expressed by the complex VP *schläft auf dem Bauch*. By contrast, in the case of external modification, the sentence reports on an instance of the event type ‘sleep’. Whether the sleeper (the actor involved in the sleeping event) is sleeping on the belly or on the back is left open. Thus, the local adverbial localises the event token. According to Maienborn, this semantic function is correlated with a VP-external syntactic position.

With respect to the modification of a noun by an adjective, I would like to suggest that the prenominal position in English (and German) is generally restricted to internal modification. That is to say, prenominal adjectives modify at the type-level (which in the nominal realm is the kind-level):



Interestingly, German prepositional phrases that modify a nominal behave in analogy to prepositional phrases modifying a verb:

- (14) a. *Das Kind auf dem Bauch*
 the child on the belly
 ‘the child on her belly’
- b. *Der Kopffäger mit den zwei Köpfen*
 the headhunter with the two heads
 ‘the headhunter who has two heads’

The phrases in (14) can be read “literally” in a way that the child will be understood to be located in the auf-region of a certain belly and the headhunter will be understood to possess two trophies. Arguably, this literal reading will result when the prepositional phrase is semantically computed *after* the meaning of the head noun has become spatiotemporally localised. As a consequence, the modifier adds its property to a child-token or headhunter-token, respectively. However, the prepositional phrase can also be semantically computed *prior to* the application of loc. In this case, the property delivered by the PP adds to the kinds ‘child’ or ‘headhunter’, respectively, leading to more specific kinds. When the complex kind description is then spatiotemporally localised, the NPs

in (14) refer to instances of the kinds ‘child on her (the child’s) belly’ and ‘headhunter having two heads’; not to some child located on an alien belly, but to some child located on her own belly; not to some headhunter possessing two alien heads, but to some headhunter bizarrely possessing two own heads.

7.5 Attributive adjectives

Note that the ambiguous modifiers in (14) always appear in postnominal position. This comes as no surprise if, as I suspect, prenominal modification in German and English is restricted to kind-level modification. Consider the following examples from Bolinger (1967):

- (15) a. *the stolen jewels*
b. *the jewels stolen*
- (16) a. *the responsible man*
b. *the man responsible*
- (17) a. *the only navigable river*
b. *the only river navigable*

Bolinger (1967) observes that the property ascribed to the referent of a noun by an adjective is interpreted in different ways depending on the position of the adjective. According to Bolinger, a postnominal adjective modifies the “temporary” state of the referent of the noun. A prenominal adjective, by contrast, has a “characterising” impact on the nominal referent: together with the noun the adjective forms a complex category term. For this to be possible, Bolinger reasons, the category must be “customary” for the respective community (culture) of speakers. Dealing with the same semantic effect, Bhatt (1994) uses the terms “permanent property” to characterise the semantic contribution of prenominal adjectives and “temporary property” to characterise postnominal adjectives.

Bolinger’s and Bhatt’s observations easily blend in with the present framework. The permanence effect of prenominal adjectives results from the composition of the adjectival meaning and the nominal meaning at the kind-level. Recall that kinds (object types) are conceived of as mental categories (cf. chapter 3). The modification of a basic kind-level predicate leads to a syntactically complex kind-level predicate. The kind extension of the latter is a subset of the kind extension of the former.

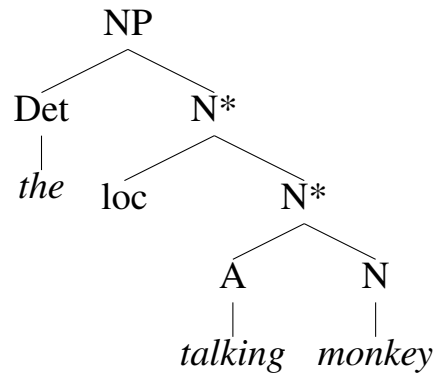
Let us look at an example. Whereas (18a) may well make sense, (18b) contains a contradiction:

- (18) a. *The talking monkey is not talking at the moment.*

- b. **The monkey talking is not talking at the moment.*

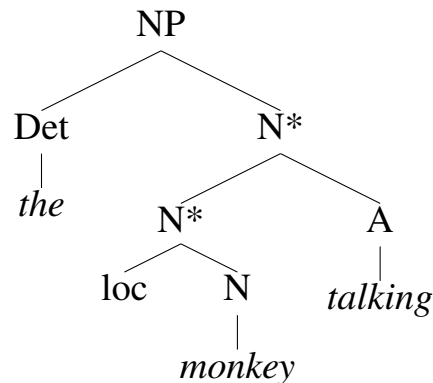
The reason is that *talking monkey* is a complex kind-level predicate formed in order to refer to a particular object in the given context of utterance. If this complex kind-level predicate is spatiotemporally localised in order to become an object-level predicate, it will characterise the set of all those objects populating the context that, first of all, are monkeys and that, secondly, are expected to be able to speak.

(19)



In *monkey talking*, by contrast, it is the noun alone that forms the kind-level predicate on the basis of which the referent is found in the utterance context. At first, the noun *monkey* is converted into an object-level predicate by means of spatiotemporal localisation. Then the adjective comes into play to modify this object-level predicate. This means that the adjective applies at the object-level. The noun and the deverbal adjective together form a complex object-level predicate which characterises the set of all objects in the context of the utterance which are, first, monkeys and, secondly, talking. As an object fulfilling these two conditions is not compatible with the predication in (18b), the sentence expresses a contradiction.

(20)



The assumption that prenominal attributive adjectives always modify at the kind-level might help to solve a problem that the truth conditional semantic

approach faces. In formal semantics, modification is modeled as set intersection. Accordingly, the expression *pink elephant* is analysed as characterising the set of all objects which at the same time belong to the set of pink objects and to the set of elephants (s. Heim & Kratzer 1998):

- (21) a. $[[pink]] = \lambda x \in D_o. PINK(x)$
 b. $[[elephant]] = \lambda x \in D_o. ELEPHANT(x)$
 c. $[[pink elephant]] = \lambda x \in D_o. PINK(x) \& ELEPHANT(x)$

The well-known problem is that, more often than not, the interpretation of the adjective turns out to be sensitive to the choice of the particular noun to which it attaches. This led Partee (1995) to formulate the “head primacy principle”:

In a modifier-head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted relative to the local context created from the former context by the interpretation of the head (Partee 1995:334)

With respect to (22a) the problem manifests itself as follows. Under the formal semantic intersection analysis we would also find the biggest mouse within the extension of *small mouse* because it belongs to the extension of *mouse* and to the extension of *small*. Even the biggest mouse is still a small animal after all.

- (22) a. *small mouse*
 b. *red hair*

The situation in (22b) is a little different. The intersection analysis predicts contra to fact that the extension of *red hair* will be empty, because red hair is not literally red by its colour but rather orange or brown (cf. Gärdenfors 2000)¹⁰.

If we assume, as proposed, that prenominal adjectives function as kind-level modifiers, the problem loses its efficacy. The adjective *small* no longer denotes the set of all *objects* characterisable as small, but rather the set of all kinds characterisable as small. Which kinds are small? The simple answer is that of an arbitrary kind ‘X’ two subkinds ‘small X’ and ‘big X’ can be formed; the first covers all those objects of X that are below average in size relative to X, the second covers all those that are above average. Accordingly, the extension of *small* relative to the taxonomic kind domain would be: $\{X \in D_k \mid SMALL(X)\}$. It includes ad hoc categories (cf. Barsalou 1983) like ‘small elephant’, ‘small mouse’, ‘small bucket’ etc., but also the more customary categories ‘small can’, ‘small car’, ‘small intestine’ etc¹¹.

¹⁰Let us ignore coloured hair for the sake of the argument.

¹¹Note that the extension of *small* must be non-finite: for every ‘small X’ the subkind ‘small small X’ can be coined, given communicative needs to do so.

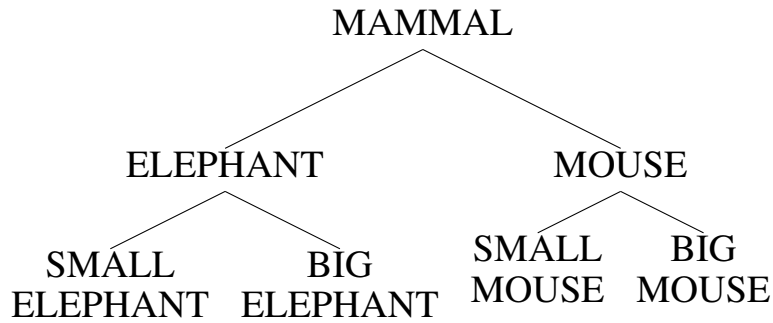


Figure 7.2: small elephant, big mouse

The problem related to the fact that even the biggest mouse is small dissolves if instead of the object domain the kind domain is chosen as the domain of evaluation for *small mouse*:

- (23) a. $[[\textit{small mouse}]] = \lambda X \in D_k. \text{SMALL}(X) \ \& \ \text{MOUSE}(X)$

Relative to the taxonomy given in figure 7.2, the kind-level predicate *mouse* denotes the set {‘mouse’, ‘small mouse’, ‘big mouse’}. The kind-level predicate *small* denotes the set {‘small elephant’, ‘small mouse’}. The complex kind-level predicate *small mouse* denotes the intersection of these two sets: {‘small mouse’}.

The complex predicate *red hair* can be dealt with by the same token:

- (24) a. $[[\textit{red}]] = \lambda X \in D_k. \text{RED}(X)$
 b. $[[\textit{hair}]] = \lambda X \in D_k. \text{HAIR}(X)$
 c. $[[\textit{red hair}]] = \lambda X \in D_k. \text{RED}(X) \ \& \ \text{HAIR}(X)$

The fact that red hair is “in reality” not red and that small elephants are “in reality” big (compared to mice) is irrelevant if adjectival modification computes at the kind-level. The “reality” talked about here is the world of spatiotemporally localised entities, i.e. the object domain. The conditions holding within this “reality” are crucial only for determining which objects truly instantiate the (complex) kinds ‘small elephant’ and ‘red hair’, which have been composed with respect to the kind domain.

7.6 On the semantic function of NPs

At the end of section 2, we stopped at the discussion of bare plurals in English (and German), which has gained much attraction since Carlson (1977). I return to it now. Let me briefly recapitulate the essence of section 2.

As English (like German) has a grammaticalised and hence obligatory singular-plural distinction, English nominal predicates come in two forms: those characterising sets of singularities, and those characterising sets of pluralities¹². Two lexicalised functors (“instantiators”) are available to derive a referential expression from a nominal predicate. The descriptive content of the nominal is being connected to some familiar reference marker (i.e. to a reference marker that has been established in the discourse model already prior to the utterance) by means of the definite article. The indefinite article signals that the nominal description is connected to a novel reference marker, i.e., to a reference marker that is being introduced through the utterance. Because, unlike the definite article, the indefinite article is lexically restricted to singular nouns, the following question arises: what if reference to some novel plurality is intended?

My answer to this question presupposes quite a strong assumption. I take it to be desirable from a theoretical point of view that *every* NP has the grammatical status of an argument expression. This position contrasts with many modern semantic theories in which it is assumed that NPs may also be used as predicates, i.e., as property-denoting expressions of the semantic type $\langle e,t \rangle$ ¹³.

For instance, the postcopular NP in a sentence like *John is a doctor* is usually viewed as a predicative NP (e.g. Partee 1987, Heim & Kratzer 1998). Van Geenhoven (1998) analyses the direct object of a sentence like *John ate an apple* under its so-called weak reading as a “semantically incorporated” predicate expression. McNally (1998) suggested that the pivot NP of an existential sentence as in *There is a fly in my soup* is formed by a predicative NP¹⁴. Kallulli (2005) takes existentially interpreted bare plural subjects as in *Dogs are barking* to be of semantic type $\langle e,t \rangle$. In the following, I will defend my hypothesis, facing all these proposals:

NP-HYPOTHESIS An NP (or DP if you prefer) is *always* an argument. It refers either to some object (or set of objects), or to some kind (or set of kinds).

In terms of DRT the hypothesis takes on the following form:

NP-HYPOTHESIS (DRT-version) An NP (DP) *always* introduces some reference marker into the discourse – either one standing for an object or one standing for a kind. The reference marker can have its own (newly introduced) index or be co-indexed with some pre-established reference marker.

¹²Again, I ignore mass nouns to reduce the complexity of the discussion.

¹³Or $\langle s, \langle e,t \rangle \rangle$ in the intensional version.

¹⁴But see McNally (1997) and (2005) for a different view.

Now remember that, as discussed in chapter 3, properties and kinds are informationally equivalent and related toward each other in a systematic way. Every property corresponds to a kind:

To any natural property, like the property of being a dog, there corresponds a kind, viz. the dog-kind. Conversely, any natural kind will have a corresponding property (the property of belonging to that kind) (Chierchia 1998:348)

Taking this fact into account, my proposal to consider apparently non-referential, predicative, NPs to be indeed referential, namely kind-referential, loses the revolutionary flavour it might have at first sight. To be an instance of the kind ‘doctor’ is to have the property of being a doctor. Thus, if we accept that “predicative” NPs are not of semantic type $\langle e,t \rangle$, but of type $\langle e_k \rangle$, nothing will change with respect to the informational content of the NP. What will change is the relationship between form and content (syntax-semantics-map). It will turn out to be as simple as can be, namely an isomorphism: the semantic function of an NP is to introduce a reference marker into the discourse.

syntactic category	semantic type
NP	$\langle e \rangle$

Figure 7.3: NP, syntax-semantics-map

7.7 The grammatical null article

When discussing bare nominals in English or elsewhere, one has to distinguish carefully between two different syntactic levels. Being plural nouns, bare plurals are predicate expressions functionally characterising sets. Being noun phrases, bare plurals are argument expressions (according to the NP-hypothesis put forward in the previous section). This view implies that a covert functor must be responsible for choosing a certain element from the set characterised by the bare plural noun to become the referent of the bare plural NP.

Does this mean that I commit myself to the existence of a null article in the grammar of English? Yes and no. Yes, inasmuch as a covert operator transforms the N-predicate into an NP-argument on the semantic plane. No, because this covert operator is unlike the overt articles in that it is no lexical element. Instead, the covert article is conceived of as a type shifter (cf. Partee 1987, Dölling 1992, Chierchia 1998) made available by the grammatical system. This type shifter is

of the same semantic type as the articles: $\langle\langle e,t\rangle,e\rangle$. What precisely are the conditions for it to apply?

As mentioned above, I assume that the syntax provides well-defined argument positions. Such are the subject and direct object position of a sentence formed on the basis of an (in)transitive verb or the position following a preposition, for instance. Argument positions must be filled by referential expressions, that is, expressions of the semantic type $\langle e\rangle$. When an expression of the semantic type $\langle e,t\rangle$ appears in syntactic argument position, its semantic type has to be coerced to suit the semantic requirements of the syntactic position. In other words, a covert type shift from $\langle e,t\rangle$ to $\langle e\rangle$ is triggered. This type shift applies automatically whenever the described constellation occurs. In the following examples, “ \emptyset ” symbolises the taking place of covert instantiation:

- (25) a. $[S[NP \emptyset [N \textit{dogs}]] [VP \textit{bark}]]$
b. $[S[NP \emptyset [A \textit{red}]] [VP \textit{is} [NP \textit{a} [N \textit{colour}]]]]$
c. $[NP \textit{the} [N \textit{lady} [PP \textit{in} [NP \emptyset [A \textit{red}]]]]]$

(25a) illustrates the case of a syntactic configuration in which a bare plural noun is shifted to a bare plural NP. In (25b) and (25c) the adjective *red*, which is lexically a predicate denoting the property of being red, is typeshifted to the NP *red*. The latter NP denotes an entity that must be contextually inferred as fulfilling the condition of being red. Due to the postcopular NP *a colour*, this amounts to the abstract entity ‘the colour red’ in (25b). The context in (25c) suggests the referent of the NP *red* be determined as a red dress.

Syntactic argument positions must be filled by referential (or quantificational) expressions¹⁵. When a predicate expression appears in a syntactic argument position, it will *automatically* be converted into a referential expression. However, this automatism is constrained by the principle of lexical blocking (Chierchia 1998, Krifka 2004):

Coercion is constrained by a blocking principle that says that if a language has overt means to express a type shift, then they have to be used. (Krifka 2004:7)

I now collected all the pieces necessary to describe a system that predicts exactly the types of (count noun based) NPs that are actually met in English. Let me repeat the building blocks:

¹⁵In general (cf. Krifka 2004) it holds that syntactic argument positions must be filled by expressions able to compose with the predicate via functional application. This may happen by the filler saturating an argument slot of the predicate (referential expressions) or by the filler itself selecting for a predicate (generalised quantifier). As already said, I restrict myself to the discussion of referential expressions.

- NPs are always of the semantic type $\langle e \rangle$ whereby “e” can mean “ e_o ” (for objects) or “ e_k ” (for kinds)
- syntactic argument positions are reserved for NPs; $\langle e, t \rangle$ -expressions appearing in syntactic argument position will be covertly typeshifted to $\langle e \rangle$
- the lexical articles (*a*, *the*) are function words combining with $\langle e, t \rangle$ -expressions to form $\langle e \rangle$ -expressions; their availability blocks the application of the covert typeshift

7.8 The requirement of earliest application

The difference in nature between overt articles being lexical creatures and covert articles being grammatical creatures has important implications for their functioning. As is well known, bare plural NPs differ empirically from NPs flagged by an overt article in a number of respects. The probably most important observation is that bare plurals generally have narrow scope (Carlson 1977, Chierchia 1998). Let us recall a classic example:

- (26) a. *Miles wants to meet policemen*
 b. *Miles wants to meet a policeman*

The bare plural NP in (26a) actualises only the “nonspecific” or “opaque” interpretation according to which it does not matter for Miles which particular policemen he meets, any would do. The indefinite NP in (26b), by contrast, allows also for the “specific” or “transparent” interpretation according to which there is a particular policeman and Miles wants to meet exactly him.

One might want to think that the structure of the bare plural NP contains a *lexical* null article that like the overt articles converts kind-level and object-level predicates into kind terms and object terms. However, this would leave the different linguistic behavior as observable in (26) totally unexplained (cf. Krifka 2004).

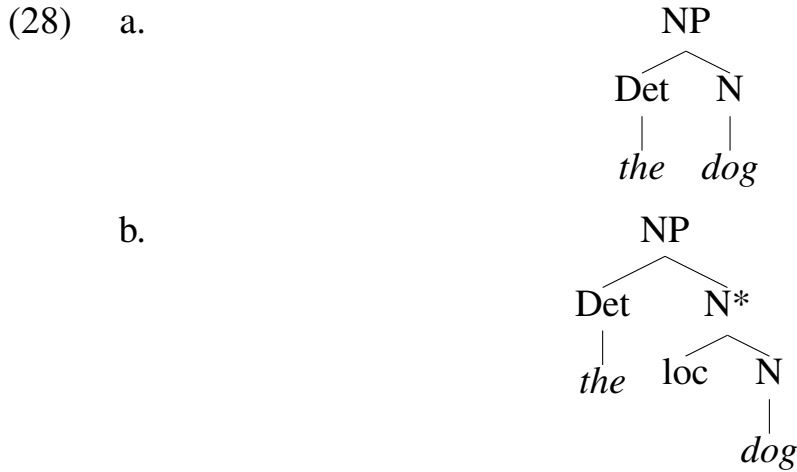
I will now present how the grammatical architecture proposed in this work can help to explain the empirical differences between overt (lexical) articles and covert (grammatical) articles. What is meant by “proposed grammatical architecture” is the assumption that the kind domain is the fundamental domain which is relevant to the formation of linguistic structures which means that object-level expressions are necessarily derived from kind-level expressions. With respect to the formation of noun phrases, this amounts to the following:

A count noun is lexically a kind-level predicate. A nominal expression can denote objects only if sentence-internal and/or sentence-external

information shifts the domain of evaluation from the kind-level to the object-level (spatiotemporal localisation).

To understand the derivational difference between a kind-referring NP and an object-referring NP compare again the following examples. The kind-referring subject in (27a) is derived in accordance with (28a), while the object-referring subject in (27b) is derived in accordance with (28b):


- (27) a. *The dog is man's best friend.* [= (9a)]
 b. *The dog is sitting on my place.* [= (9b)]



(28) illustrates how we may think of the overt article (in this case *the*) applying either prior to or after spatiotemporal localisation to terminate the derivation of the NP. An overt lexical article may operate on sets of kinds or on sets of objects. Which application is triggered depends on whether or not the co(n)textual elements that flank the noun phrase call for a spatiotemporal interpretation. If the noun phrase shows up in the argument position of an episodic (object-level) predicate like *is sitting on my place*, for instance, it is clear that an object must be delivered and that, accordingly, structure (28b) is motivated. One might want to call this mechanism “semantic agreement” (see Šmelev 1996).


The syntactic structure in (28a) corresponds to the semantic representation in (7). The structure in (28b) involves an additional syntactic step associated with an additional semantic operation, compare (29b):

- (29) a. $[_N \textit{dog}]$




↓ spatiotemporal localisation

- b. $[_{N^*} [_{loc} \emptyset] [_N \textit{dog}]]$



↓ overt instantiation

c. [NP [Det *the*] [N* [loc \emptyset] [N *dog*]]]



Thus, an overt article (in this case *the*) can instantiate thematic kind arguments as well as thematic object arguments. Note that whenever a thematic object argument is instantiated, so is a thematic kind argument. Technically, the semantic computation of an article corresponds to a construction rule saying that all *x*'s (or *X*'s) in the DRS must be replaced by *u*'s (or *U*'s).

The situation is different in the case of covert instantiation. I assume that the type shifter (that I also referred to as the grammatical null article) must apply in the derivation as soon as possible. I cannot exactly tell why this is so. Presumably, the requirement for earliest possible application is an economy constraint. The grammar should make use of covert type shifters only if the application is unavoidable to repair a semantic mismatch, and the type shifter should do as little covert work as necessary to make uninterpretable linguistic structure interpretable. The requirement for earliest application guarantees that only a minimum of invisible changes in meaning takes place.

The requirement for earliest application has the consequence that the grammatical null article, unlike its overt counterparts, applies always *before* any co(n)textual elements flanking the bare plural could switch the domain of evaluation from the kind-level to the object-level, i.e., prior to spatiotemporal localisation. This has the effect that as a rule the grammatical null article operates on kinds; bare plural NPs are, in accordance with Carlson (1977), always kind terms.

(30) shows the derivation of a kind-referring bare plural NP as in *Dogs evolved from the wolf*:

(30) a. [N *dogs*]

dog(X) plural(X)

↓ covert instantiation

b. [NP [Det \emptyset][N *dogs*]]

U
dog(U) plural(U)

Thus, the prediction is that bare NPs in general and English bare plural NPs in particular are basically kind-referring NPs.

7.9 Existentially interpreted bare plural NPs

Of course, not every bare plural NP is interpreted generically. Often, the bare plural NP is felt as not relating to the kind, but to manifestations of it.

According to Carlson, the existence of stages of the kind is asserted under this reading, which is called the existential reading. Working without stages and taking objects to be spatiotemporal entities instead, as I do (cf. chapter 4), this means that existentially interpreted bare plural NPs come with an existence assertion concerning *objects*. Object-referring bare plural NPs are found whenever a bare plural kind term occurs as the syntactic argument of an object-level predicate. The semantic mismatch resulting from this constellation – required is an expression of the semantic type $\langle e_o \rangle$, given is an expression of the semantic type $\langle e_k \rangle$ – can be repaired by an interpretive mechanism to which Chierchia (1998) refers to as “derived kind predication”:

[W]henver an object-level argument slot in a predicate is filled by a kind (in an episodic frame), the type of the predicate is automatically adjusted by introducing a (local) existential quantification over instances of the kind [...] Derived Kind Predication (DKP): If P applies to objects and k denotes a kind, then $P(k) = \exists x [{}^U k(x) \& P(x)]$ (Chierchia 1998:364)

In example (31a) the kind-referring use of the subject-NP is lexically explicated. According to Chierchia’s rule of derived kind predication, the sentence expresses the statement that there are instances (objects) of the kind described

as *that kind of animal* which at the moment of speech are ruining the speaker’s garden.

- (31) a. *That kind of animal is ruining my garden.* [Chierchia 1998]
 b. *Rats reached Australia in 1770.* [Krifka et al. 1995]

In (31b) the subject-NP is realised as a bare plural. Given the Carlsonian view (which I subscribe to, see above) that bare plural NPs are always kind-denoting is true and given the DKP-rule, the sentence expresses that there are instances of the kind ‘rat’ that reached Australia in the year 1770.

What is going on can be described as follows: syntactically, an episodic verbal predicate combines with a kind term; semantically, however, it assigns, by virtue of its episodicity, a property to an object. This is the reason why we have to infer the existence of object instances of the kind denoted by the kind term in order to interpret the sentence; the inferred objects are needed to take over the role of the bearer of the object-level property denoted by the verb.

The semantic representation in (32) is the result of computing the subject-NP *rats* and the VP *reached Australia in 1770* each in its own right. Above the dotted line we see the conditions contributed by the kind term *rats*. Below we see the conditions contributed by the VP. Among them there is also the semantic role relation “theme(e,x)”, which requires that an object token argument x participates in the event token e (“E” is a reference marker standing for an event type).

- (32) [_S *Rats reached Australia in 1770*]

U e E
rat(U)
plural(U)

reached Australia(E)
in 1770(e)
loc(e,E)
theme(e,x)

Being a kind term, the subject expression *rats* cannot deliver the object-level argument required by the episodic predicate *reached Australia in 1770*. After all, the semantic contribution of *rats* consists only in an object type referent (=kind referent) which is represented by the reference marker U. Note, however, that this is exactly the kind of semantic mismatch that activates the rule of derived kind predication (DKP-rule) in Chierchia’s system. Translated into the present framework, the DKP-rule brings about the post hoc spatiotemporal localisation

of the kind reference marker U . This corresponds to additionally introducing the relation $\text{loc}(u,U)$ and a reference marker u :

(33) [_S *Rats reached Australia in 1770*]

u U e E
rat(U)
plural(U)
loc (u,U)
reached Australia(E)
in 1770(e)
loc(e,E)
theme(e,u)

Be it noted that, of course, object terms too can fill the argument slot of such an episodic predicate, which shows that the episodic predicates are ordinary ones:

(34) *James Cook reached Australia in 1770.*

7.10 The semantic contribution of a token-restricting modifier

Unlike the lexicalised overt article the covert grammatical “article” applies as a rule at the kind-level. Therefore, its contribution is usually a reference marker standing for a kind. There are, however, important exceptions to this rule. These exceptional cases come about whenever the bare plural noun is modified by an expression of which the interpretation presupposes object reference (e.g. by some indexical). The reason is that in the presence of a modifier, the earliest possible stage for the grammatical null article to apply is after the modification. Given that the modifier is a token-restricting modifier¹⁶, modification will switch the domain of evaluation from the default kind-level to the object-level. As a consequence, this yields bare plural NPs with the potential to actualise so-called wide scope readings (Carlson 1977, Chierchia 1998):

(35) *Betty is looking for books she lost yesterday.* [Carlson 1977:317]

If the grammatical null article must always apply as soon as possible, and if the kind domain represents the basic ontological domain, it will be predicted that covert instantiation automatically leads to kind term formation – as long as no external modifier causes the switch to the object domain before the null article can apply. It seems to me that this is a neat explanation for the existence of examples like (35):

¹⁶The term “token-restricting modifier” is adopted from Yury Lander and tantamount to the term “external modifier” that I used in section 4.

(36) a. [_N *books*]

book(X) plural(X)

↓ external modification

b. [_{N*} *books she lost yesterday*]

book(X) plural(X) she lost yesterday(x) loc(x,X)

↓ covert instantiation

c. [_{NP} *books she lost yesterday*]

u,U
book(U) plural(U) she lost yesterday(u) loc(u,U)

So why do certain bare plural NPs like, for instance, *parts of that machine*, *boys sitting here* or *books she lost yesterday* allow for the transparent (“wide scope”) interpretation? The answer is, because the semantic computation of each of these complex syntagms involves an indexical element triggering spatiotemporal localisation. This in turn has the consequence that an object reference marker will be introduced into the discourse when the predicate undergoes instantiation. Now, if the semantic representation contains an object reference marker, in every situation which makes the sentence true there will have to be an appropriate object; for example, a plurality of books that Betty lost yesterday. If the semantics of the sentence entails this meaning component, the sentence will have to be interpreted transparently: there are books such that Betty lost them yesterday and it is exactly these books that Betty is looking for today.

So far, so good. This is not the only interpretation of (35) though. The sentence also allows for an opaque reading. Accordingly, Betty is not looking for the specific book tokens she lost yesterday but for any tokens of the type

‘book that Betty lost yesterday’. This reading is pragmatically marked but not impossible. Imagine Betty for some reason suspects that she lost books yesterday, neither knowing for sure that this really was the case nor how many books she apparently lost. Now Betty starts looking around for these potentially lost books. In this case it is possible to refer to this looking-for-event by means of (35).

Doesn’t this run counter to the proposal I made before? How can the opaque reading come about when the relative clause *she lost yesterday* triggers spatiotemporal localisation, as I argued above? Admittedly, I have no good explanation for that. I believe that under the opaque reading the crucial constituent *books she lost yesterday* undergoes some sort of “re-intensionalisation” taking us from the literal spatiotemporally localised meaning (as it is predicted by my theory) to some derived meaning paraphrasable as: ‘books of which Betty *believes* that she lost them yesterday’. By virtue of this reconceptualisation, the semantic impact of *loc* is annulled (the spatiotemporal meaning is relativised to possible worlds). Consequently, the relative clause no longer has a spatiotemporally localising effect. The transparent interpretation loses its ground and what remains is the opaque interpretation.

This kind of reasoning is in line with Carlson (1977:320) who likewise considers it to be a conceptual problem whether a nominal syntagm characterising a set of objects can also be used generically. Compare also Chierchia’s approach to which I will turn in the next section.

Let us finally compare the DRSs of (33) and (36) for clarification. In each case we ended up with a DRS which includes the condition $loc(u,U)$. There is an important difference in the respective history of this condition, however. In (33), *loc* is due the DKP-rule – a well-defined type shifter that applies to maintain compositionality. Note that it applies *after the covert article has turned the bare plural noun into a kind-referring bare plural noun phrase*. In (36), by contrast, it is the covert article itself that introduces *loc* thereby instantiating an object-level property, i.e. a property that entails the condition $loc(x,X)$. So here the ultimate source of *loc* is the indexical component of the modifier. What is responsible for *loc* is not a repair strategy, as in (33), but the ordinary semantics of the articles, as conceived of in this work.

To repeat, articles serve to instantiate (in the sense of Farkas & de Swart 2003) thematic arguments. As long as they meet kind-level thematic arguments only, their semantic impact will result in kind-referring NPs. To form an object-referring NP, an article must attach to a structure that involves also an object-level thematic argument. This holds for overt as well as for covert articles. The difference is, so I proposed, that the latter must apply “as soon as possible”. This brings it about that a bare NP by itself is never ambiguous (as is an NP with an overt article) but, as a rule, always kind-referring. The bare plural in

(33) is no exception to this rule. After all it is the very fact of kind reference of the bare NP that calls for DKP. The bare plural in (36) is an exception, but it is the exception that supports the rule. This is because in this special case the trigger of $\text{loc}(x,X)$ lies within the NP so that instantiation results in $\text{loc}(u,U)$ even though the covert article applies as soon as it can.

7.11 Chierchia's and Krifka's explanations

I am of course not the first to propose a solution to the problem of bare plurals. Let me discuss two prominent approaches.

For Chierchia (1998) kind terms are derived from predicates by means of the down-operator \ulcorner (see chapter 3). Chierchia reserves this operator for predicates that characterise sets of objects showing a “sufficiently regular behavior”:

[K]inds are generally seen as regularities that occur in nature. [...] What counts as kind is not set by grammar, but by the shared knowledge of a community of speakers. It thus varies, to a certain degree, with the context, and remains somewhat vague. Lexical nouns identify kinds. Complex nouns may or may not. (Chierchia 1998:348)

For Chierchia, “complex nouns” like *parts of that machine* or *books she lost yesterday* denote properties that usually do not correspond to kinds, unlike “lexical nouns” like, for instance, *policemen*. Therefore, \ulcorner is not defined in cases of complex nouns, and the only type shifter available to turn the respective predicate into an argument is \exists .

Chierchia's approach and the approach presented in this work are similar in that they both draw a distinction between two sorts of predicates. Chierchia distinguishes predicates characterising sets of objects showing a sufficiently regular behaviour from predicates characterising sets of objects showing no sufficiently regular behaviour. My own approach is based on a distinction between non-spatiotemporal (i.e. kind-level) predicates and spatiotemporal (i.e. object-level) predicates. In both approaches, it is assumed that predicates occurring in argument positions undergo type shifts. In Chierchia's system, “regular predicates” are shifted to kind terms, and “irregular predicates” are shifted to generalised existential quantifiers. In my approach, non-spatiotemporal predicates are shifted to kind terms, and spatiotemporal predicates are shifted to object terms.

As far as I can see, the crucial difference is what is taken to be the default mode of reference in each case. In Chierchia's system, properties must fulfill a special condition (namely the condition of being regular) to be mappable onto kinds. Thus, object-level denotation is the default. In my system, properties

must fulfill a special condition (namely, the condition of being spatiotemporally localised) to be mappable onto objects. So this time kind-level denotation is the default.

In other words, for Chierchia, kind-referring bare plural NPs represent the special case for which syntactically simple nominals are suited by their very nature and for which syntactically complex nominals are suited under certain circumstances. The normal case, however, is the existentially interpreted bare plural. Only in this case there is an (existential) quantifier that may interact with other quantifiers to yield the two readings in (35), i.e., the wide scope and the narrow scope interpretations.

To demonstrate the scopal ambiguity of sentences based on “irregular predicates”, Chierchia uses not (35) but (37a):

- (37) a. *John is looking for parts of that machine.*
b. *John is looking for boys sitting here.*

No doubt, (37a) has both of the two relevant interpretations. Whether (35) really has both interpretations remains an open question, however. Now compare (37b) in this respect. Does this sentence have the opaque reading at all? What we can conclude from these observations is that not any bare plural modification by means of some indexical can realise the opaque reading. This supports the view suggested above that the opaque reading is the result of a conceptual reinterpretation.

An (admittedly weak) argument against Chierchia relates to the following sentence:

- (38) *Peter is looking for single fathers with twins.*

Out of the blue, this sentence is more likely to actualise the opaque reading. A typical context to embed it is the following one: Peter who works at the psychology department of some scientific institute is planning to exercise an experiment on twin research. For this purpose, he is looking for test persons. If Chierchia was right, the transparent reading of (38) should be preferred. After all, single fathers with twins are not known to show a sufficiently regular behaviour. The preferred interpretation is the opaque reading, however.

My proposal predicts that as long as the modifying constituent of a complex predicate does not contain any indexical elements, complex bare plurals, too, can be used to refer to kinds. Even in cases where we can hardly say that the predicate shows a regular behaviour:

- (39) *The detective is looking for shortsighted Marlboro smokers with shoe size 46 and extraordinary computer skills.*

Krifka (2004) offers a different explanation for the deviant linguistic behaviour of bare plural NPs like *parts of that machine* etc. His explanation relies on the principle of lexical blocking, according to which a semantic operation cannot take place covertly if there are overt linguistic means to serve the same purpose. For Krifka it is the availability of *some* in English which is responsible for the fact that bare plural NPs like *policemen, boys, books* etc. cannot actualise transparent readings in contexts of the type [*Betty is looking for ___*] – in contrast to bare plural NPs like *parts of that machine, boys sitting here, books she lost yesterday* etc., that can do so.

Krifka describes the difference between the two classes of nominals as follows: a bare plural like *parts of that machine* denotes a finite set of objects, whereas a bare plural like *policemen* denotes an open set. Krifka observes that when *some* combines with a nominal denoting an open set, the *some*-phrase forms a referential expression; Krifka suggests to treat *some* semantically as a choice function. In contrast to that, when *some* combines with a nominal denoting a finite set, the *some*-phrase is a partitively interpreted quantificational structure. (40b), for example, implies that Betty is not looking for all of the books that she lost yesterday.

- (40) a. *Betty is looking for some books.*
 b. *Betty is looking for some books she lost yesterday.*

Krifka argues that if there was a covert type shifter with the same meaning as *some*, i.e. the meaning of a choice function, the application of this covert operator would only in those cases be lexically blocked by *some* where a bare plural denotes an open set. In contrast to that, bare plural predicates characterising finite sets may well undergo type shifting by means of the covert choice function, because with these, *some* gives rise not to the choice function reading, but to the partitive reading.

- (41) a. *Betty is looking for {books/policemen/boys/...}.*
 b. *Betty is looking for {books she lost yesterday/policemen in this building/boys sitting here/...}.*

Therefore, only (41b) has a referential (transparent) reading brought about by the covert choice function operator. To get the same reading in (41a) *some* must be used.

There is a problem in connection with this approach. First note that it presupposes that there are two homomorphous expressions, *some*₁ and *some*₂. If there was only one *some*, this expression would necessarily always block the semantically equivalent covert type shifter. To see this, let us take on a more abstract perspective: The attested empirical asymmetry is explained by the fact that one and the same covert operator is in some cases (context 1) lexically blocked,

while it is free to apply in other cases (context 2). The reason is that only in context 1 a semantically equivalent overt operator is available, but not in context 2. This is because only in context 1 the appropriate grammatical conditions are met for the application of the semantically equivalent overt operator; context 2 is just not suited for it. Therefore, *some* in context (40a) must be different from *some* in context (40b). Otherwise, the argument would not work.

The problem is: if the covert operator can apply in (41b), why cannot the semantically equivalent overt operator *some* apply, too? Possibly, there is some kind of syntactic restriction which excludes the semantically equivalent *some* from contexts like (41b). Note, however, that such a syntactic restriction would undermine the principle of lexical blocking, which says that when a language has an overt operator, it must be used. If there are no syntactic reasons to rule out the semantically equivalent *some* in (41b), there will have to be semantic reasons. However, if there are semantic reasons to bar the semantically equivalent *some* from (41b), this will of course also bar the covert operator from (41b). They are semantically equivalent, after all.

7.12 Object-level arguments in kind-level positions

Above we saw how kind-level arguments compose with object-level predicates. How about the opposite situation? How can the present framework cope with object terms that appear in the argument position of kind-level predicates?

What does sentence (42) express? Presumably, it expresses the proposition according to which the kind ‘whale’ belongs to the set of subkinds of the kind ‘mammal’.

(42) *A whale is a mammal*

Given this, (42) would be true relative to any situation in which the tree in figure 7.4 constituted a partial structure of the taxonomic kind domain.

What does sentence (43) express? Intuitively speaking, it expresses that the object individual Moby Dick is an instance of the kind ‘whale’.

(43) *Moby Dick is a whale.*

It is tempting to identify the semantic contribution of the copular in (42) with the subkind-of relation **T** and the semantic contribution of the copular in (43) with the instance-of relation **R**.¹⁷ However, this would lead to a not very appealing ambiguity of the copular verb¹⁸. Can we find a compositional analysis

¹⁷See Krifka et al. (1995) on **T** and **R**; see also chapter 3.

¹⁸Compare the discussion in Dölling (1997), who presents arguments against this solution.




Figure 7.4: mammal and whale

for (42) and (43) that maintains a uniform semantics for the copular, that can derive the subkind-of interpretation of (42), and that can derive the instance-of interpretation of (43)?

(44) *A blue whale is a whale.*


For ease of comparison, I will discuss (44) instead of (42). By hypothesis, the two instances of *is a whale* in (43) and (44) are not only formally, but also semantically equivalent. What could be the common meaning? One syntactically combines with an object term, the other combines with a kind term. Given this, how can they be semantically equivalent?

The basic assumption is that a singular noun like *whale* is listed in the lexicon as a kind-level predicate. Its extension includes, among other kinds, the elements ‘whale’, ‘blue whale’, ‘dolphin’, ‘sword whale’, ‘beluga’. (45) shows the semantic representation of *whale* in terms of predicate logic. Capitals signal that a variable is kind-level, small letters signal that a variable is object-level (thus, the information “ $\in D_k$ ” is, strictly speaking, redundant):

(45) $[[whale]] = \lambda X \in D_k. WHALE(X)$

As outlined above, I adopt Farkas & de Swart’s (2003) revision of Kamp & Reyle’s (1993) original DRT version which I extend by a distinction between kind-arguments and object-arguments. Given such a framework, the nominal predicate *whale* induces the representation (46a), and the combination with the indefinite article yields (46b):

(46) a. $[_N whale]$



↓ D-instantiation

b. [NP [Det *a*][N *whale*]]

U
whale(U)

Please note that the condition “whale(U)” does not imply that the discourse referent U stands for *the* kind ‘whale’. Rather, U can stand for any kind which satisfies the description associated with *whale*, be it the superkind ‘whale’, or a subkind like ‘blue whale’, ‘dolphin’ etc.

Thus, by means of the NP *a whale* the speaker refers to a whale kind. What happens if this NP is following the copular verb? Here is a proposal:

(47) a. [NP *a whale*]

U
whale(U)

↓ application of the copular

b. [VP [cop *is*][NP *a whale*]]

U
whale(U)
=(X,U)

Accordingly, the copular verb selects for an entity-referring expression to form a predicate characterising the set of all entities identical to the input entity. This way a VP is formed which semantically corresponds to a predicate, the extension of which is the set of all kinds identical to whatever whale kind is represented by U. Now let us assume that this predicate appears in (43) as well as in (44). How does the meaning of this predicate compose with the meanings of the respective subject expressions?

Farkas & de Swart (2003) speak of “A-instantiation” when the meaning of an argument expression saturates the argument slot of a predicate. Technically, A-instantiation is the substitution of the predicate’s thematic argument by the argument’s discourse referent. Let us look at (44) first. In analogy to (46), the subject NP introduces a kind referent:

(48) [NP *a blue whale*]

V
blue whale(V)

The NP introduces a discourse referent standing for a kind, which is required to satisfy the description associated with *blue whale*. I assume that, unless there is information to the contrary, the discourse referent will be interpreted as standing for the (maximum) kind ‘blue whale’ named by the head noun. Due to this default mechanism, the condition “blue whale(V)” is specified to “blue whale=V”. If the resulting NP syntactically combines with the predicate *is a whale*, the discourse referent V of the NP will a-instantiate the thematic argument X of the predicate:

(49) a. [VP *is a whale*]

U
whale(U)
=(X,U)

↓ A-instantiation

b. [S [NP *a blue whale*][VP *is a whale*]]

V U
whale(U)
=(V,U)
blue whale=V

The sentence (44) is true in all situations (models) into which the DRS (49b) is embeddable. These are all those situations in which there is the kind ‘blue whale’ and in which this kind is identical to a kind satisfying the description associated with *whale*. In other words, the sentence is true in all situations in which the taxonomic domain is structured in such a way that the kind ‘blue whale’ forms a subkind of the kind ‘whale’, as exemplified in figure 7.5.

Let us then turn to (43). The problem here is that the subject expression *Moby Dick* is object-referring. How can its meaning compose with the meaning of the kind-level predicate (47b)? My solution to this problem rests on assumptions that I argued for in chapter 6. Here, I repeat only the gist of it.

Formal semanticists adopt a very abstract perspective onto the empirical facts to be accounted for. This is necessary to make mathematical tools (in particular




Figure 7.5: whale and blue whale

set theory) available for the description of linguistic meanings. Only at a very abstract level of representation, entities of very different ontological sorts can be treated on a par as elements describable in set theoretical terms. Unfortunately, in the case of kinds and objects, essential differences tend to become blurred by abstracting away from the real nature of things. This is because the difference between kinds (types) and objects (tokens) lies at the heart of cognition.

The domain of objects is the domain (real world) we live in; the domain of kinds, by contrast, is the domain that we mentally represent (or construct) in order to sort the real world we live in. This difference has strong implications. Whereas objects *exist* in the common sense of existence, namely in the sense of existence in space and time, kinds “only” *exist* in some metaphorical sense of existence, namely in the sense of existence in the mind. Representing linguistic meanings as sets made up of entities of various ontological sorts and taking the overall set of these entities to be the “world” (model), we are in danger of tracing over the different modes of existence of objects and kinds. Specifically, we are in danger of tracing over the fact that, cognitively, objects can only be accessed via kinds.

The purpose of sorting the world of objects into kinds is not to make linguistic life easier, but rather to make it possible at all. By means of categorisation, a conceptual level of symbolic representation is established without which language (and not only language) would be unthinkable. Language is the symbolic system that enables us to communicate with each other about the real world (of objects) we live in. Therefore, to talk about real objects, we have to start off from the conceptual kind domain.

This way of looking at things holds the key to solve the Moby Dick problem because it implies the dual-reference-hypothesis (see chapter 6). Here it is again in the DRT version:

DUAL-REFERENCE-HYPOTHESIS (DRT version) Every expression introducing an object discourse referent must at the same time introduce a kind discourse referent, and a condition which determines

that the respective object must be an instance of the respective kind.

That reference to objects always involves a kind is pointed out by sortalist like, for instance, Macnamara et al. (1994), Xu (1997) or Carey & Xu (1999); compare also chapter 2. Semanticists like Krifka (1995:399) and Dölling (1992, 1993) make the same observation, which leads them to conclude that the kind domain is ontologically prior to the object domain.

Note that this view implies a quite unusual semantics for proper names. From a sortalist perspective, reference to an object by means of a proper name involves a kind, too (in particular: Macnamara et al. 1994). Accordingly, a proper name like *Moby Dick* induces a DRS containing an object discourse referent, a kind discourse referent and a realisation relation holding among them. This amounts to the following DRS:

(50) $[\text{NP } Moby Dick]$

v V
$Moby Dick = v$
$\text{loc}(v, V)$

Thus, I assume that a proper name does not only introduce an object discourse referent v into the discourse, but also a kind discourse referent V . For which kind exactly V is standing is semantically underspecified, and must be determined by pragmatics. What is known about this kind is only that the bearer of the name must be an instance of it; the semantics allows for a wide range of values for V .

The *Moby Dick* problem can now easily be solved. Given that (50) is the semantic representation of *Moby Dick*, the kind-level predicate *is a whale* can assign its property to the kind symbolised by the kind discourse referent V . In terms of Farkas & de Swart's extended DRT-version, we can say that the kind discourse referent V a-instantiates the thematic kind argument X of the predicate:

(51) a. $[\text{VP } is a whale]$

U
$whale(U)$
$=(X, U)$

↓ A-instantiation

b. $[\text{S } [\text{NP } Moby Dick][\text{VP } is a whale]]$

v V U
whale(U)
=(V,U)
Moby Dick=v
loc(v,V)

Given the dual-reference-hypothesis is true, a uniform semantic analysis of the predicate *is a whale* in (43) and (44) is possible. The generic character of both of these sentences (the feeling that the assigned property says something essential or defining about the respective subject referent) is due to the fact that, in both cases, the predicate is kind-level.

7.13 Conclusions

Generally, content words can function as type-level or token-level predicates. According to the view advocated in this work, the type-level function is the basic one because this is just the lexical format of content words¹⁹. To function as a token-level predicate, a content word must be supported by the co(n)text.

To derive a referential expression on the basis of a lexical noun, the noun has to undergo a certain semantic operation which Farkas & de Swart (2003) call “instantiation”. Languages can have at their disposal lexicalised function words or morphemes specialised in the task of instantiation. If a language has no such overt instantiators, referential expressions must be derived covertly²⁰.

English has two overt articles. They are semantically underspecified inasmuch as they are capable of turning kind-level predicates into kind terms, and object-level predicates into object terms. Other languages possess articles which are lexically restricted to either kind-level or object-level instantiations. For example, Maori possesses two overt indefinite articles (Bauer 1993, Chung & Ladusaw 2004). As I argued elsewhere (Mueller-Reichau 2005), one of them is reserved for the creation of indefinite kind terms (*he*), and the other one for the creation of indefinite object terms (*teetahi*).

As the English indefinite article is lexically restricted to singular nominals, the English article system has a gap: there is no overt instantiator available for indefinite reference to pluralities. In such cases, the principle of lexical blocking is suspended and the grammatical null article can apply. The main point of this chapter was that, due to its type shifter nature, the null article always applies, unlike overt articles, as soon as possible. This prevents the nominal

¹⁹This is also the reason why lexical word formation rules usually operate on types.

²⁰Compare Padučeva (1979) on Russian.

predicate with which the null article “combines” from undergoing spatiotemporal localisation – unless a token-restricting modifier attaches to the nominal predicate.

In the tradition of Carlson (1977) it was argued that bare NPs are generally kind-referring expressions. This was related to the assumption that nouns are stored in the lexicon as kind-level predicates. The idea was that, as a rule, object-level denotation presupposes some sort of semantic process operating on a basic lexical unit. Several grammatical operations can qualify a lexical noun for object-level denotation.

First, the articles. If a noun is accompanied by an article, it may denote at the kind-level or at the object-level:

- (52) a. {*the whale/ a whale*} *is a mammal.*
b. {*the whale/ a whale*} *was killed.*

Secondly, prosodic means. In the German example (53b), the object-level (existential) bare plural subject carries a special accent. Arguably, this marked prosodic structure reflects the marked semantic structure; the default kind-level interpretation (53a) gets a neutral prosodic realisation:

- (53) a. *Hunde BELLen.*
dogs bark
'Dogs bark'
b. *HUNDe bellen.*
dogs bark
'Dogs are barking'

Third, token-restricting modifiers. If a noun is modified by an expression that links reference to objects populating the context of the utterance, the bare NP will be object-level. The following example from Mehlig (1983) is accordingly ambiguous. The bare nominal subject may refer to the kind ‘student of the Department of Slavonic Studies of the University of Kiel’. In that case the Genitive-NP does not refer to the closed class of students that at the moment belong to the department in Kiel, but rather to the open class of potential students. Alternatively, the bare nominal may refer to the concrete group of actual students of the department at the moment of speech.

- (54) *Studenten unseres Instituts sprechen gut Russisch.*
students our.GEN department.GEN speak well Russian
'Students of our department speak Russian well'

At least, the context within which a bare nominal appears can motivate its object-level interpretation. The trigger may be the sentence-internal context,

e.g. the episodic predicate in the English translation of (53b), but also the sentence-external context. An illustrative example is from Sheila Glasbey (cf. Maienborn 2003:38). Unlike the subject of (55a), which may well be understood generically, the subject of (55b) can hardly be understood generically:

- (55) a. *Drinkers were under-age.*
b. *John was shocked by his visit to the Red Lion. Drinkers were under-age, drugs were on sale, and a number of fights broke out while he was there.*

Jäger (2001) and Maienborn (2003:37) observe that, in out of the blue contexts, the generic reading is the preferred interpretation of a bare plural subject, even if it is combined with a Carlsonian stage-level predicate.

- (56) a. *Feuerwehrleute sind hungrig.* [?]
firemen are hungry
'Firemen are hungry'
b. *Emperors are naked.* [?]

This observation is in line with the predictions following from my analysis. After all, I assume that nominal predicates begin their linguistic life at the kind-level, and that they enter the object-level only if they are supported by the context in which they appear. As in out of the blue contexts there is, by definition, no context that might support the bare nominal, it is predicted that the existential reading requires additional effort from the side of the interpreter.

References

- Abbott, B. (2008): Issues in the Semantics and Pragmatics of Definite Descriptions in English. In: J. Gundel & N. Hedberg (Eds.), *Reference: Interdisciplinary perspectives*. Oxford: Oxford University Press, 61–72.
- Altmann, H. (1981): *Formen der “Herausstellung” im Deutschen*. Tübingen: Max Niemeyer Verlag.
- Bach, E. (1986): The Algebra of Events. *Linguistics & Philosophy* 9, 5–16.
- Bach, E. (1994): The Semantics of Syntactic Categories: A Cross-Linguistic Perspective. In: J. Macnamara & G. E. Reyes (eds.), *The Logical Foundations of Cognition*. New York: OUP, 264–281.
- Bache, C. & N. Davidsen-Nielsen (1997): *Mastering English. An Advanced Grammar for Non-native and Native Speakers*. Berlin, New York: Mouton de Gruyter.
- Baker, M. (2003): *Lexical Categories: Verbs, Nouns and Adjectives*. Cambridge: CUP.
- Barsalou, L. (1983): Ad hoc Categories. *Memory & Cognition* 11, 211–227.
- Barsalou, L. (2000): Concepts: Structure. In: A. Kazdin (ed.), *Encyclopedia of Psychology*, Vol. 2. Oxford: OUP, 245–248.
- Bartsch, R. (1998): *Dynamic Conceptual Semantics: A Logico-Philosophical Investigation into Concept Formation and Understanding*. Stanford: CSLI Publications.
- Bauer, W. (1993): *Maori*. London and New York: Routledge.
- Bhatt, D. (1994): *The Adjectival Category: Criteria for Differentiation and Identification*. Amsterdam/Philadelphia: John Benjamins.
- Bierwisch, M. & R. Schreuder (1992): From Concepts to Lexical Items. *Cognition* 42, 23–60.

- Blok, S., G. Newman & L. J. Rips (2005): Individuals and their concepts. In: W. k. Ahn et al. (eds.), *Categorization inside and outside the lab*. Washington, D.C.: American Psychological Association, 127–149.
- Bolinger, D. (1967): Adjectives in English: Attribution and Predication. *Lingua* 18, 1–34.
- Bulygina, T. (1982): K postroeniju tipologii predikatov v ruskom jazyke. In: O. Seliverstova (ed.), *Semantičeskie tipy predikatov*. Moskva: Nauka, 7–85.
- Carey, S. (1995): Continuity and Discontinuity in Cognitive Development. In: E. Smith (ed.), *Invitation to Cognitive Science, Part III: Thinking*. Cambridge: MIT Press, 101–129.
- Carey, S. & F. Xu (1999). Sortals and Kinds: An Appreciation of John Macnamara. In: R. Jackendoff, P. Bloom & K. Wynn (eds.), *Language, Logic, and Concepts*. Cambridge: MIT Press, 311–335.
- Carlson, G. (1977): *Reference to Kinds in English*. Amherst: University of Massachusetts.
- Carlson, G. (1992): Natural kinds and common nouns. In: A. von Stechow & D. Wunderlich (eds.), *Handbook of Semantics*. Berlin: de Gruyter, 370–398.
- Carlson, G. (1995): Truth Conditions of Generic Sentences: Two Contrasting Views. In: G. Carlson & F. Pelletier (1995), 224–237.
- Carlson, G. & F. Pelletier (1995). *The Generic Book*. Chicago, London: University of Chicago Press.
- Chierchia, G. (1984): *Topics in the Syntax and Semantics of Infinitives and Gerunds*. Amherst: University of Massachusetts.
- Chierchia, G. (1989): Introduction. In: G. Chierchia, B. Partee & R. Turner (eds.), *Properties, types and meaning. Volume II: Semantic issues*. Dordrecht, Boston, London: Kluwer Academic Publishers, 1–20.
- Chierchia, G. (1995): Individual-Level Predicates as Inherent Generics. In: G. Carlson & F. Pelletier (1995), 176–223.
- Chierchia, G. (1998): Reference to Kinds Across Languages. *Natural Language Semantics* 6, 339–405.
- Chierchia, G. & R. Turner (1988): Semantics and Property Theory. *Linguistics & Philosophy* 11, 261–302.

- Chung, S. & W. Ladusaw (2004): *Restriction and Saturation*. Cambridge: MIT Press.
- Cohen, A. (1999): *Think Generic! The Meaning and Use of Generic Sentences*. Stanford: CSLI Publications.
- Cohen, A. (2001): On the generic use of indefinite singulars. *Journal of Semantics* 18, 183–209.
- Cohen, A. (2002): Genericity. *Linguistische Berichte* 10, 58–89.
- Corbett, G. & N. Fraser (1993): Network Morphology. A DATR account of Russian Nominal Inflection. *Journal of Linguistics* 29, 113–142.
- Corbett, G. & N. Fraser (1999): Default Genders. In: B. Unterbeck (ed.), *Gender in Grammar and Cognition. Approaches to Gender*. Berlin, New York: Mouton de Gruyter, 55–97.
- Croft, W. (1986): *Universal Quantifiers and Generic Expressions*. Ms.
- Croft, W. (2001): *Radical Construction Grammar. Syntactic Theory in Typological Perspective*. Oxford: OUP.
- Dahl, Ö. (1995): The Marking of the Episodic/Generic Distinction in Tense-Aspect Systems. In: G. Carlson & F. Pelletier (1995), 412–425.
- Davidson, D. (1967): The Logical Form of Action Sentences. In: D. Davidson (ed.), *Essays on Actions and Events*. Oxford: Clarendon Press, 105–122.
- Dayal, V. (2004): Number Marking and (In)definiteness in Kind Terms. *Linguistics & Philosophy* 27, 393–450.
- de Swart, H. (1998). *Introduction to Natural Language Semantics*. Stanford: CSLI Publications.
- Dölling, J. (1992): Flexible Interpretation durch Sortenverschiebung. In: I. Zimmermann & A. Strigin (eds.), *Fügungspotenzen*. Berlin: Akademie Verlag, 23–62.
- Dölling, J. (1993): Commonsense Ontology and Semantics of Natural Language. *Zeitschrift für Sprachtypologie und Universalienforschung* 46, 133–141.
- Dölling, J. (1997): Ist die Kopula mehrdeutig? Anmerkungen zu einem Vorurteil. In: U. Scheffler & K. Wuttich (eds.), *Termingebrauch und Folgebeziehung*. Berlin: Logos, 5–24.
- Farkas, D. & H. de Swart (2003): *The Semantics of Incorporation*. Stanford: CSLI.

- Finch, S. (2001): Copular Elements in Bengali and the Stage-/Individual-Level Distinction. Ms.
- Foley, W. A. (1998): Symmetrical voice system and precategoriality in Philippine languages. Ms.
- Gatt, A. (2004): Regular and Generic Possessives in Maltese. In: J.-Y. Kim, Y. Lander & B. Partee (eds.), *Possessives and Beyond: Semantics and Syntax*. Amherst: GLSA, 199–215.
- Geist, L. (1999): Russisch *byt'* ('sein') als funktionale und/oder lexikalische Kategorie. *ZAS Papers in Linguistics* 14, 1–40.
- Geist, L. (2004): *Kopulasätze an der Syntax-Semantik Schnittstelle. Die russische Variante*. Berlin: Humboldt Universität zu Berlin.
- Gendler Szabó, Z. (2000): Descriptions and Uniqueness. *Philosophical Studies* 101, 29–57.
- Geurts, B. (1997): Good News about the Description Theory of Names. *Journal of Semantics* 14, 319–348.
- Geurts, B. (1999): *Presuppositions and Pronouns*. Oxford: Elsevier.
- Geurts, B. (2004): *The Pragmatics of Rigidity*. Paper presented at the ZAS Mini Workshop on Proper Names, 10. September 2004.
- Gärdenfors, P. (2000): *Conceptual Spaces*. Cambridge: MIT Press.
- Hawkins, J. (1991): On (in)definite articles: implicatures and (un)grammaticality prediction. *Journal of Linguistics* 27, 405–442.
- Heim, I. (1983): File Change Semantics and the Familiarity Theory of Definiteness. In: R. Bäuerle, C. Schwarze & A. von Stechow (eds.), *Meaning, Use and the Interpretation of Language*. Berlin: de Gruyter, 164–190.
- Heim, I. (1988): *The Semantics of Definite and Indefinite Noun Phrases*. New York: Garland.
- Heim, I. (1991): Artikel und Definitheit. In: A. von Stechow and D. Wunderlich (eds.), *Semantik*. Berlin, New York: de Gruyter, 487–535.
- Heim, I. & A. Kratzer (1998): *Semantics in Generative Grammar*. Oxford: Blackwell.
- Jackendoff, R. (1983): *Semantics and Cognition*. Cambridge, MA: MIT Press.
- Jäger, G. (2001). Topic-comment Structure and the Contrast between Stage-Level and Individual-Level Predicates. *Journal of Semantics* 18, 83–126.

- Kallulli, D. (2005): On existential bare plural “subjects”: they don’t exist! Ms.
- Kaminski, J., J. Call & J. Fischer (2004). Word learning in a domestic dog: evidence for fast mapping. *Science* 304, 1682–1683.
- Kamp, H. & U. Reyle (1993). *From Discourse to Logic*. Dordrecht: Kluwer.
- Katz, G. & R. Zamparelli (2005). *Genericity in Natural Language*. Lecture notes: European Summer School in Language, Logic and Information, Edinburgh 2005.
- Kratzer, A. (1995): Stage-Level and Individual-Level Predicates. In: G. Carlson & F. Pelletier (eds.), *The Generic Book*. Chicago, London: University of Chicago Press, 125–175.
- Krifka, M. (1995): Common Nouns: A Contrastive Analysis of Chinese and English. In: G. Carlson & F. Pelletier (eds.), *The Generic Book*. Chicago, London: University of Chicago Press, 398–411.
- Krifka, M. (2001): Kinds of Kind Reference. Extended handout of a talk given at the Workshop on Genericity, University of Cologne, Nov 10-11, 2001.
- Krifka, M. (2003): Compositionality. In R. A. Wilson & F. Keil (eds.), *The MIT Encyclopedia of Cognitive Sciences*. Cambridge, MA: MIT Press.
- Krifka, M. (2004): Bare NPs: Kind-referring, Indefinites, Both, or Neither? In: R. B. Young & Y. Zhou (eds.), *Proceedings of Semantics and Linguistic Theory (SALT) XIII*, University of Washington, Seattle. Cornell: CLC Publications.
- Krifka, M. et al. (1995): Genericity: An Introduction. In: G. Carlson & F. Pelletier (eds.), *The Generic Book*. Chicago, London: University of Chicago Press, 1–124.
- Löbner, S. (1985): Definites. *Journal of Semantics* 4, 279–326.
- Löbner, S. (2002): *Understanding semantics*. London: Arnold.
- Longobardi, G. (1994): Reference and Proper Names: A Theory of N-Movement in Syntax and Logical Form. *Linguistic Inquiry* 25, 609–665.
- Macnamara, J., M. La Palme Reyes & G. E. Reyes (1994). Reference, Kinds, and Predicates. In: J. Macnamara & G. E. Reyes (eds.), *The Logical Foundations of Cognition*. New York: OUP, 91–143.
- Maienborn, C. (2000): Zustände - Stadien - stative Ausdrücke: Zur Semantik und Pragmatik von Kopula-Prädikativ-Konstruktionen. *Linguistische Berichte* 183, 271–307.

- Maienborn, C. (2001): On the Position and Interpretation of Locative Modifiers. *Natural Language Semantics* 9, 191–240.
- Maienborn, C. (2003). *Die logische Form von Kopula-Sätzen*. Berlin: Akademie Verlag.
- McNally, L. (1997): *A semantics for the English existential construction*. New York, London: Garland.
- McNally, L. (1998): Existential Sentences without Existential Quantification. *Linguistics & Philosophy* 21, 353–392.
- Mehlig, H. R. (1983): Nominale Referenz, Zeitreferenz und Prädikatssemantik. In H. R. Mehlig (ed.), *Slavistische Linguistik 1982. Referate des VIII.Konstanzer Slavistischen Arbeitstreffens*. München: Otto Sagner, 48–75.
- Mueller-Reichau, O. (2004). From concepts to objects and eventualities. In: R. Asatiani et al. (eds.), *Proceedings of the Fifth Tbilisi Symposium on Language, Logic and Computation, October 6-10, 2003, Tbilisi State University*. Amsterdam: ILLC Publications.
- Mueller-Reichau, O. (2005): A kinds-based DRT approach to the semantics of articles across languages. Paper presented at ESSLLI 2005 Workshop “Formal semantics and cross-linguistic data”, Edinburgh 8.-19. August 2005.
- Padučeva, E.V. (1979): Denotativnyj status imennoj gruppy i ego otaženie v semantičeskom predstavlenii predloženiya. *Naučno-Techničeskaja Informacija* 2/9, 25–31.
- Padučeva, E.V. (1996): *Semantičeskie issledovanija*. Moskva: Jazyki Russkoj Kul'tury.
- Partee, B. (1987): Noun Phrase Interpretation and Type-Shifting Principles. In: *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*. Foris.
- Partee, B. (1995): Lexical Semantics and Compositionality. In: L. Gleitman & M. Liberman (eds.), *Invitation to Cognitive Science, Part I: Language*. Cambridge: MIT Press, 311–360.
- Rosch, E., C. Mervis, W. Gray, D. Johnson & P. Boyes-Braem (1976): Basic Objects in Natural Categories. *Cognitive Psychology* 8, 382–439.
- Šmelev, A. (1996): *Referencial'nye Mechanizmy Russkogo Jazyka*. Tampere: Tampereen yliopisto Jäljennepalvelu.

- Spelke, E. (1990): Principles of Object Perception. *Cognitive Science* 14, 29–56.
- Spelke, E., K. Breinlinger, J. Macomber & K. Jacobson (1992): Origins of Knowledge. *Psychological Review* 99/4, 605–632.
- Tomasello, M. (1999): *The Cultural Origins of Human Cognition*. Harvard:Harvard University Press.
- Van Geenhoven, V. (1998): *Semantic Incorporation and Indefinite Descriptions*. Stanford: CSLI.
- Vergnaud, J.-R. & M. Zubizarreta (1992). The Definite Determiner and the Inalienable Construction in French and in English. *Linguistic Inquiry* 23, 595–652.
- Weiss, D. (1983). Indefinite, definite und generische Referenz in artikellosen slavischen Sprachen. In H. R. Mehlig (ed.), *Slavistische Linguistik 1982. Referate des VIII.Konstanzer Slavistischen Arbeitstreffens*. München: Otto Sagner, 229–261.
- Wiese, H. (1999). Die Verknüpfung sprachlichen und konzeptuellen Wissens: Eine Diskussion mentaler Module. In: I. Wachsmuth & B. Jung (eds.), *Kog-Wis99. Proceedings der 4. Fachtagung der Gesellschaft für Kognitionswissenschaft Bielefeld, 28. September - 1. Oktober 1999*, 92–97. St. Augustin:Infix-Verlag.
- Xu, F. (1997): From Lot's Wife to a Pillar of Salt: Evidence that Physical Object is a Sortal Concept. *Mind & Language* 12, 365–392.
- Xu, F. (2005): Categories, Kinds, and Object Individuation in Infancy. In: C. Gershkoff-Stowe & D. Rakison (eds.), *Building object categories in developmental time. Papers from the 32nd Carnegie Symposium on Cognition*. New Jersey: Lawrence Erlbaum, 63–89.
- Xu, F. & S. Carey (1996): Infants' metaphysics: the case of numerical identity. *Cognitive Psychology* 30, 111–153.
- Xu, F., S. Carey & N. Quint (2004): The emergence of kind-based object individuation in infancy. *Cognitive Psychology* 49, 155–190.
- Zamparelli, R. (1998): A Theory of Kinds, Partitives and of/z Possessives. In: A. Alexiadou & C. Wilder (eds.), *Possessors, Predicates and Movement in the Determiner Phrase*. Amsterdam: John Benjamins, 259–301.
- Zamparelli, R. (2002): Definite and Bare Kind-denoting Noun Phrases. In: C. Beyssade et al. (eds.), *Romance Languages and Linguistic Theory 2000*. Amsterdam: John Benjamins, 305–342.