In: P. Ludewig & B. Geurts (eds.), Lexikalische Semantik aus Kognitiver Sicht. Tübingen, 141–171. Gunter Narr Verlag

Reinhard Blutner

Lexical Underspecification and Pragmatics[1]

Abstract

This paper is part of a theoretical attempt to investigate interactions between the (mental) lexicon and pragmatics. It aims at giving a systematic and explanatory account of pragmatic phenomena that are connected with the semantic underspecification of lexical items. Cases in point are the pragmatics of adjectives, systematic polysemy, the distribution of lexical and productive causatives, blocking phenomena, the interpretation of compounds and many phenomena presently discussed within the framework of Cognitive Semantics.

After emphasizing some important consequences of the traditional view of lexical semantics (= contrastive analysis of lexemes within the Katz-Fodor tradition of semantics), several phenomena are collected that seem to conflict with the theoretical assumptions made by it. These phenomena are taken as arguments in favor of a particular account of the division of labor between lexical semantics and pragmatics. This account combines the idea of lexical underspecification with a theory of pragmatic strengthening. The basic pragmatic mechanism rests on conditions of updating the common ground and allows us to give a precise explication of notions such as generalized conversational implicature and pragmatic anomaly. It is suggested that this approach may provide a principled account of several of the lexical-pragmatic phenomena that are currently discussed.

1 Introduction

This paper is part of a theoretical attempt to investigate interactions between the (mental) lexicon and pragmatics. Already Katz & Fodor (1963) have stressed the point that a full account of lexical meaning has to include more information than that which allows one to discriminate the meanings of different words. In one of their examples they argue that *take back* is

[1] Thanks to Petra Ludewig and Bart Geurts for the inspired idea and excellent organization of this workshop. Thanks also for their helpful comments and interactions. Moreover, I am greatly indebted to Jutta Romberg, Hans-Martin Gärtner and Jörn. As always, responsibility for errors and matters of judgement rest entirely with me.

used in very different ways in the sentences (1a,b), although the relevant lexical entries are semantically unambiguous.

(1) a. Should we take the lion back to the zoo?
b. Should we take the bus back to the zoo?

An obvious difference between these sentences is that the lion is the *object* taken back to the zoo in (1a), but the bus is the *instrument* that takes us back to the zoo in (1b). The problem for the pragmatic component of utterance interpretation is to explain the difference in terms of different conceptual settings („world knowledge"), starting from a lexicon that doesn't discriminate the two occurrences of *take back* semantically and from a syntax that is completely parallel for the two sentences.

As another introductory example let's consider perception verbs in English (cf. Sweetser 1990). If Saussure is right, there is an essential arbitrary component in the association of words or morphemes with what they mean. Consequently, the feature of arbitrariness could be taken at least as a sufficient condition for the presence of semantic information. It is certainly an arbitrary fact of English that *see* (rather than, say, *buy* or *smell*) refers to visual perception when it is part of the utterance (2a). Given this arbitrary association between a phonological word and its meaning, however, it is by no means arbitrary that *see* can also have an epistemic reading as in (2b).

(2) a. I see the tree.
b. I see what you're getting at.

Moreover, it is not random that other sensory verbs such as *smell* or *taste* are not used to express an epistemic reading. Sweetser (1990) tries to sketch an explanation for such facts and insists that they have to do with conceptual organization. It is our knowledge about the inner world that implicates that vision and knowledge are highly related, in contrast to, say, smell and knowledge or taste and knowledge, which are only weakly related for normal human beings. If this claim is correct, then the information that *see* may have an epistemic reading but *smell* and *taste* do not need no longer be stipulated semantically. Instead, this information is pragmatic in nature, having to do with the utterance of words within a conceptual setting, and can be derived by means of some general mechanism of conceptual interpretation.

Considerations of this kind raise a standard puzzle for lexical semantics when we ask how to separate the (mental) lexicon from the (mental) encyclopedia. How should we separate information about the meaning of words from information about the (supposed) reality associated with these words? Admittedly, it may be rather difficult to discriminate these two kinds of information. Reliable, theory-independent empirical tests simply don't exist. There are two principal possibilities of dealing with this situation. First, the distinction between the lexicon and the encyclopedia is said to be illusory (as it has sometimes been suggested by representatives of Cognitive Semantics, e.g. Lakoff 1987). In this case, all the relevant information has to be put into the lexicon (leading to a highly non-compositional account of meaning projection). The second possibility is to take the distinction as an important one. As a consequence, we are concerned with two different types of mechanisms:

- a mechanism that deals with the combinatorial aspects of meaning
- a pragmatic mechanism that deals with conceptual interpretation.

Once we have adopted such theoretical mechanisms, the problem of discriminating lexical semantic information from encyclopedic information needs no longer look so hopeless, and we really may profit from a division of labor between semantics and pragmatics. It is the position of this paper to argue in favour of the second option.

From a Gricean perspective, two different ideas of how to overcome the divergences between (formal) meaning and natural language interpretation come into mind. The first one uses conventional implicatures as an extension of the classical information entries. The second idea uses conversational implicatures as a method to overcome the divergences. Whereas I believe that modern semantic theories (which usually are characterized as dynamic, epistemic, nonmonotonic) make the conception of conventional implicature superfluous as an addendum to the semantic component, I do not think that the same holds for conversational implicature. In fact, in this paper I will argue that the proper use of conversational implicature will resolve some of the problems of lexical interpretation that remain otherwise unsolved.

The conceptual core of the theory I want to propose demands a straight formulation of conversational implicature. Paired with the idea of (radical) semantic underspecification in the lexicon and an appropriate representation of contextual and encyclopedic knowledge this conception

avoids both unmotivated lexical ambiguities as well as expansive reinterpretation and coercion mechanisms.

There are two basic aims of this paper. First, I want to demonstrate some general problems we are confronted with when trying to analyze the utterance of words within concrete conceptual and contextual settings and to go beyond the aspects of meaning typically investigated by a contrastive analysis of lexemes within the Katz-Fodor tradition of semantics. This may help to develop a feeling of what kind of problems may be approached by means of the division of labor between lexical semantics and pragmatics. Second, I would like to argue in favor of a particular account of the interaction between lexical semantics and pragmatics, one that combines the idea of (radical) semantic underspecification in the lexicon with a theory of pragmatic strengthening (based on conversational implicatures).

The organization of this paper is as follows. In the next three sections I will concentrate on some general features that can be ascribed to the *traditional view of (lexical) semantics*: The principle of compositionality (section 2), the monotonicity of the lexical system (section 3), and the persistence of anomaly (section 4). These features are not intended to characterize the family of theories called the *traditional view* in any sense completely. Rather, their selection is intended to emphasize several properties that may become problematic when a broader view of utterance meaning is taken. In section 5, I will use these features for marking out the borderline between semantics and pragmatics. Furthermore, I introduce a particular way of combining (radical) semantic underspecification with a theory of pragmatic strengthening. Finally, section 6 outlines my present theory of pragmatic strengthening. The basic mechanism rests on conditions of updating the common ground and allows to give a precise explication of notions as generalized conversational implicature and pragmatic anomaly. It is suggested that this approach may provide a principled account of several of the lexical-pragmatic phenomena that are recently discussed.

2 Systematicity and the principle of compositionality

One nearly uncontroversial feature of our linguistic system is the *systematicity of linguistic competence*. According to Fodor & Pylyshyn (1988: 41-42) this feature refers to the fact that the ability to understand and produce some expressions is intrinsically connected to the speaker's ability to

produce and understand other expressions that are semantically related. The classical solution to account for the systematicity of linguistic competence crucially makes use of the *principle of compositionality*. In its general form this principle states the following:

(3) The meaning of a complex expression is a function of the meanings of its parts and their syntactic mode of combination.

In an approximation that is sufficient for present purposes, the principle of compositionality states that „a lexical item must make approximately the same semantic contribution to each expression in which it occurs“ (Fodor & Pylyshyn 1988). As a simple example, consider adjective-noun combinations such as *brown cow* and *black horse*. Let's take "absolute" adjectives (such as *brown* and *black*) as one-place predicates. Moreover, non-relational nouns are considered as one-place predicates as well. Let's assume further that the combinatorial semantic operation that corresponds to adjectival modification is the intersection operation. Fodor & Pylyshyn (1988) conclude that these assumptions may explain the feature of systematicity in the case of adjectival modification. For example, when somebody is able to understand the expressions *brown cow* and *black horse*, he should understand the expressions *brown horse* and *black cow* as well. Note that it is the use of the intersection operation that is involved in explaining the phenomenon, not compositionality per se. Nevertheless, the principle of compositionality is an important guide that helps us to find specific solutions to the puzzle of systematicity.

Lexical semantics are concerned with the meanings of the smallest parts of linguistic expressions that are assumed to bear meaning. Assumptions about the meanings of lexical units are justified empirically only in as far as they make correct predictions about the meanings of larger constituents. Consequently, though the principle of compositionality clearly goes beyond the scope of lexical semantics, it is indispensable as a methodological instrument for lexical semantics. I state the principle of compositionality as the first feature characterizing the standard view of (lexical) semantics.

Until now, we have taken adjectives like *red, interesting* or *straight* as intersective adjectives, and I have illustrated how this pretty simple analysis brings together systematicity and compositionality. Unfortunately, the view that a large range of adjectives behaves intersectively proved to be questionable. For example, Quine (1960) notes the contrast between *red apple* (red on the outside) and *pink grapefruit* (pink on the

inside), and between the different colors denoted by *red* in *red apple* and *red hair*. In a similar vein, Lahav (1989, 1993) argues that an adjective such as *brown* doesn't make a simple and fixed contribution to any composite expression in which it appears.

> In order for a cow to be brown most of its body's surface should be brown, though not its udders, eyes, or internal organs. A brown crystal, on the other hand, needs to be brown both inside and outside. A brown book is brown if its cover, but not necessarily its inner pages, are mostly brown, while a newspaper is brown only if all its pages are brown. For a potato to be brown it needs to be brown only outside, ... Furthermore, in order for a cow or a bird to be brown the brown color should be the animal's natural color, since it is regarded as being 'really' brown even if it is painted with all over. A table, on the other hand, is brown even if it is only painted brown and its 'natural' color underneath the paint is, say, yellow. But while a table or a bird are not brown if covered with brown sugar, a cookie is. In short, what is to be brown is different for different types of objects. To be sure, brown objects do have something in common: A salient part that is wholly brownish. But this hardly suffices for an object to count as brown. A significant component of the applicability condition of the predicate 'brown' varies from one linguistic context to another. (Lahav 1993: 76)

Some authors - for example, Keenan (1974), Partee (1984), Lahav (1989, 1993) - conclude from facts of this kind that the *simplistic view* mentioned above must be abolished. As suggested by Montague (1970), Keenan (1974), Kamp (1975) and others, there is a simple solution that addresses such facts in a descriptive way and obeys the principle of compositionality. This solution considers adjectives essentially to be adnominal functors. Such functors, for example, turn the properties expressed by *apple* into those expressed by *red apple*. Of course, such functors have to be defined disjunctively in the manner illustrated in (4):

(4) RED(X) means roughly the property
- a. of having a red inner volume if X denotes fruits only the inside of which is edible
- b. of having a red surface if X denotes fruits with edible outside
- c. of having a functional part that is red if X denotes tools

Let us call this view the *functional view*. It should be stressed that the functional view describes the facts mentioned above only by enumeration. Consequently, it doesn't account for any kind of systematicity concerning

our competence to deal with adjective-noun combinations in an interesting way. Another (notorious) problem of this view has to do with the treatment of predicatively used adjectives. In that case the adjectives must at least implicitly be supplemented by a noun. Various artificial assumptions are necessary which make such theory inappropriate (cf. Bierwisch 1989 for more discussion of this point). We may conclude that compositionality doesn't necessarily lead to systematicity.

There is a third view about treating the meanings of adjectives, which I call the *free variable view*. In a certain sense, this view can be seen as preserving the advantages of both the simplistic as well as the functional view, but as overcoming their shortcomings. The free variable view has been developed in considerable details in case of gradable adjectives (see, for example, Bierwisch 1989 and the references given therein). It is well-known that the applicability conditions of restricting adjectives that denote gradable properties, such as *tall, high, long, short, quick, intelligent* vary depending on the type of object to which they apply. What is high for a chair is not high for a tower and what is clever for a young child is not clever for an adult. Oversimplifying, I can state the free variable view as follows. Similar to the first view, the meanings of adjectives are taken to be one-place predicates. But now we assume that these predicates are complex expressions that contain a free variable. Using an extensional language allowing λ-abstraction, we can represent the adjective *long* (in its contrastive interpretation), for example, as λx LONG(x,X), denoting the class of objects that are long with regard to a comparison class, which is indicated by the free variable X. At least on the representational level the predicative and the attributive use of adjectives can be treated as in the first view: *The train is long* translates to (after λ-conversion) LONG(**t**,X) and *long train* translates to λx [LONG(x,X) ∧ T(x)]. In these formulas **t** is a term denoting a specific train and T refers to the predicate of being a train.

Free variables are the main instrument for forming underspecified lexical representations. To be sure, free variables simply have the status of place holders for more elaborated subpatterns and expressions containing free variables should be explained as representational schemes. Free variables do not only stand as place holders for a comparison class X as just indicated. The view can be generalized to include other types of free variables as well, for example a type of variable connected with the specification of the dimension of evaluation in cases of adjectives such as *good* and *bad* or a type of variable connected with the determination of the object-dependent spatial dimensions in cases of spatial adjectives such as *wide*

and *deep*. In what follows, a variety of other kinds of variables will be considered, leading to rather complex types of lexical underspecification.

Of course, it is not sufficient to postulate underspecified lexical representations and to indicate what the sets of semantically possible specifications of the variables are. In order to grasp natural language interpretation ("conceptual interpretation"), it is also required to provide a proper account of contextual enrichment, explaining how the free variables are instantiated in the appropriate way. Obviously, such mechanism has to take into consideration various aspects of world and discourse knowledge. We are presented here with a kind of *selection task*: How to select from a set of possibilities an appropriate one where (weak) restrictions are given in the form of world and discourse knowledge.

In some particular cases, the instantiation of free variables may be done by using ordinary (monotonic) unification. If that works fine, it may be concluded that the mechanism of contextual enrichment has the feature of compositionality. In other words, the principle of compositionality stated for semantic representations can be transferred to the level of contextually enriched forms. In Blutner (1997), I have considered some examples that demonstrate that monotonic unification doesn't suffice for contextual enrichment.

There is a variety of other examples that demonstrate that our comprehension capacities have salient, non-compositional aspects. The most prominent class of examples may be found within the area of *systematic polysemy*. This term refers to the phenomenon that one lexical unit may be associated with a whole range of senses which are related to each other in a systematic way.[2] The phenomenon has traditionally been thought intractable, and in fact it *is* intractable when considered as a problem of lexical semantics in the traditional sense.

---

[2] Unfortunately, the term *systematic polysemy* covers a whole family of empirically different subphenomena which no unified terminology is available for. Expressions such as *open and closed polysemy* (Deane 1988), *conceptual specification* and *conceptual shift* (Bierwisch 1983), *sense modulation* and *sense change* (Cruse 1986), *constructional polysemy* and *sense extension* (Copstake and Briscoe 1995) may be convenient to indicate a rough outline of the classification.

3 The monotonicity of the lexical system

Another general characteristic of the standard view is connected with the idea of analyzing the meanings of lexical items as a complex of more primitive elements. The main motivation for a componential (decompositional) analysis is connected with the explanation of certain semantic relations such as antonymy, synonymy, and semantic entailment. If the meaning of a lexical item was not analyzed into components, the lexical system of grammar would have to simply enumerate the actually realized relations as independent facts. This procedure would not be descriptively very economic. More important, it would miss the point that these facts are *not* independent from each other. The componential approach can be found both in theories of meaning in generative semantics (cf. Fodor 1977) and in model-theoretic based (especially Montagovian) semantic work (cf. Dowty 1979).

Defining the meaning of lexical items in terms of a repertoire of more primitive elements leads to a second order property which I will call the *monotonicity* of the lexical system. In short, the monotonicity restriction refers to the fact that we can incrementally extend the lexical system (by adding some definitions for new lexical material) without influencing the content of elements already defined.

At first glance, the monotonicity of the lexical system looks quite natural as a constraint within formal semantics. Of course, it would be very surprising if the content of *...is a bachelor* would change if the system learns what a spinster is (by acquiring the corresponding definition). Similarly, the meaning of *prime, even, odd (number)* should be independent of whether the system knows the meaning of *rational number* or *perfect number*. [3]

It should be stressed that it is not the idea of decomposition (definition) per se that leads to the monotonicity feature of the lexical system. Instead, it is its classical treatment within a formal metalanguage that exhibits all features of a deductive system in the sense of Tarski.[4]

---

[3] A perfect number is a natural number that is identical to the sum of its true divisors; e.g. 6 = 1+2+3 or 28 = 1+2+4+7+14.

[4] Within a deductive system a consequence relation |= is defined. |= explicates the notion of a logical consequence: The formula (of a particular formal language *L*) is a logical consequence of the set of premises Γ (of *L*). For the present purpose it is unessential to consider the details of constructing the consequence relation. What is essential, however, is to remember what Tarski stated quite generally as some minimal requirements which a deductive consequence relation must fulfill if it is truly to be a logical notion. A logical consequence relation has to

In the simplest case, definitions are explicit and can be represented as $Q(x) \leftrightarrow C(x)$, where Q is the definiendum and C the definiens (an expression constructed in terms of a given system of lexical "primes"). In other cases, for example when we have to define disposition-like expressions as *soluble*, Carnap's (1936) reduction pairs may be used. An interesting case are bilateral reduction sentences. They have the form $F(x) \rightarrow (Q(x) \leftrightarrow C(x))$, with definiendum Q and definiens C (under condition F). In both cases, the system of (explicit or implicit) definitions bears the feature of monotonicity.

Figure 1 illustrates the difference between monotonic systems and non-monotonic ones in a schematic way. The picture simplifies matters by identifying meanings with extensions (represented by Venn-diagrams). In the case of a monotonic system, the addition of a new predicate R doesn't change the extensions of the old predicates P and Q. However, the same doesn't hold in the case of a non-monotonic system. In this case, we have "field"-effects: There seem to be attracting and repelling "forces" that shift the extensions of old predicates in a particular way when new lexical material comes into play. The non-monotonic system I have in mind corresponds to the so-called *Voronoi tesselation* defining a partitioning of some (abstract) space in terms of a given set of prototypes. The construction stipulates that the element x belongs to the same category as the closed prototype of the given set of prototypes. It is evident that previously defined categories may change when we add new prototypes. (For more details and for the cognitive significance of this construction, see Gärdenfors 1996).

A general problem that lexical semantics has to address is the phenomenon of (partial) lexical blocking. The existence of blocking effects may raise serious doubts about the validity of monotonicity as a property of the lexical system. The phenomenon has been demonstrated by a number of examples where the appropriate use of a given expression formed by a relatively productive process is restricted by the existence of a more "lexicalized" alternative to this expression. One case in point is due to Householder (1971). The adjective *pale* can be combined with a great many color

---

satisfy the following principles (here Γ and Γ' range over sets of formulas and ϕ over isolated formulas of *L*):

a. REFLEXIVITY: $\Gamma \models \Gamma$

b. CUT: if $\Gamma \models \Gamma'$ and $\Gamma \cup \Gamma' \models \phi$, then $\Gamma \models \phi$

c. MONOTONICITY: if $\Gamma \models \phi$, then $\Gamma \cup \Gamma' \models \phi$

The most important characteristic is MONOTONICITY. Informally, this principle states that the old theorems remain valid when the system of axioms (definitions, meaning postulates, factual knowledge) has been augmented by adding some new axioms.

words: *Pale green, pale blue, pale yellow*. However, the combination *pale red* is limited in a way that the other combinations are not. For some speakers *pale red* is simply anomalous and for others it picks up whatever part of the pale domain of *red pink* has not preempted. This suggests that the combinability of *pale* is fully or partially blocked by the lexical alternative *pink*.

Another standard example is the phenomenon of blocking in the context of derivational and inflectional morphological processes. Aronoff (1976) has shown that the existence of a simple lexical item can block the formation of an otherwise expected affixally derived form synonymous with it.

In particular, the existence of a simple abstract nominal underlying a given *-ous* adjective blocks its nominalization with *-ity*:

(5) a. curious - curiosity
tenacious – tenacity
b. furious - *furiosity - fury
fallacious - *fallacity – fallacy

While Aronoff's formulation of blocking has been limited to derivational processes, Kiparsky (1982) notes that blocking may also extend to inflectional processes and he suggests a reformulation of Aronoff's blocking as a subcase of the *Elsewhere Condition* (special rules block general rules in their shared domain). However, Kiparsky cites examples of *partial blocking* in order to show that this formulation is too strong. According to Kiparsky, partial blocking corresponds to the phenomenon that the special (less productive) affix occurs in some restricted meaning and the general (more productive) affix picks up the remaining meaning (consider examples like *refrigerant - refrigerator, informant - informer, contestant - contester*).

McCawley (1978) collects a number of further examples demonstrating the phenomenon of partial blocking outside the domain of derivational and inflectional processes. For example, he observes that the distribution of productive causatives (in English, Japanese, German, and other languages) is restricted by the existence of a corresponding lexical causative.

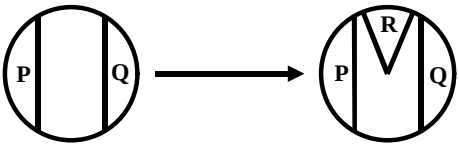

monotonic extension

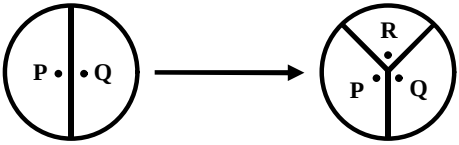

non-monotonic extension

Figure 1:
Monotonic and non-monotonic extensions of a (lexicalized) system of concepts

Whereas lexical causatives (e.g. (6a)) tend to be restricted in their distribution to the stereotypic causative situation (direct, unmediated causation through physical action), productive (periphrastic) causatives tend to pick up more marked situations of mediated, indirect causation. For example, (6b) could be used appropriately when Black Bart caused the sheriff's gun to backfire by stuffing it with cotton.

(6) a. Black Bart killed the sheriff
b. Black Bart caused the sheriff to die

The phenomenon of blocking can be taken as evidence demonstrating the apparent non-monotonicity of the lexical system. This becomes pretty clear when we take an ontogenetic perspective on the development of the lexical system. Children overgeneralize at some stage while developing their lexical system. For example, they acquire the productive rule of deriving adjectives with *-able* and apply this rule to produce *washable*, *breakable*, *readable*, but also *seeable* and *hearable*. Only later, after forms like *seeable* and *visible*, *hearable* and *audible* have coexisted for a while, the meanings of the specialized items block the regularly derived forms. Examples of this kind suggest that the development of word meanings cannot be described as a process of accumulating more and more denotational knowledge in a monotonic way. Instead, there are highly non-monotonic stages in lexical development. At the moment, it is not clear whether this

ontogenetic feature must be reflected in the logical structure of the mental lexicon. Rather, it is possible that pragmatic factors (such as Gricean rules of conversation) play an important role in determining which possible words are actual and what they really denote (McCawley, 1978, Horn 1984, Dowty 1979; see also section 6).

4 The persistence of anomaly

Lexical semantics has to account for semantic contradictions as **married spinster*, **female bachelor*, **reddish green* and for other types of semantic anomalies as exemplified by the famous **colorless green ideas sleep furiously*. Usually, semantic anomaly of an expression is defined as logical incompatibility of (some part of) the formal translation of the expression taken in union with a given system Γ of definitions and/or meaning postulates (e.g. McCawley 1971). Explicating incompatibility in terms of inconsistency and inconsistency in terms of contradictory entailments makes it possible to derive a second order property which I call the *persistence of anomaly*.

The persistence of anomaly comes in two variants: (i) if we add some new axioms to Γ, then any former anomaly persists; and (ii) if a (propositional) formula is anomalous, then every other formula that implies it is anomalous as well.[5] Both varieties seem to be satisfied empirically. It would be very surprising if the anomaly of **married bachelor* would be cancelled by learning the meaning of several new words. Once an anomaly is established, it seems to persist when the system is extended. In a similar sense it would be perplexing if the anomaly of the expression **the idea sleeps* would not persist if the expression is made more specific, e.g. **the new idea sleeps*.

It is straightforward that the notion of semantic anomaly can be converted in a notion of pragmatic anomaly if the system Γ of axioms is assumed to include other sources of knowledge, such as conceptual and ontological knowledge. Not surprising, the persistence of anomaly persists in this case.

Next, I will present some examples which may challenge the persistence of (pragmatic) anomaly as a fundamental conceptual principle. Take the well-known phenomenon of "conceptual grinding", whereby or-

---

[5] Again, it is the classical, deductive character of the entailment relation that leads to this conclusion.

dinary count nouns acquire a mass noun reading denoting the stuff the individual objects are made of, as in *There is fish on the table* or *There is Dog all over the street*. There are several factors that determine whether "grinding" may apply, and, more specifically, what kind of "grinding" (meat grinding, fur grinding, universe grinding, ...) may apply. Some of these factors have to do with the conceptual system, while others are language-dependent (cf. Nunberg & Zaenen 1992; Copestake and Briscoe 1995; Leßmöllmann 1996).

One of the language-dependent factors affecting the grinding mechanism is lexical blocking. For example, in English the specialized mass terms *pork, beef, wood* usually block the grinding mechanism in connection with the count nouns *pig, cow, tree.* This explains the contrasts given in (7).

(7) a. I ate pork/?pig
b. Some persons are forbidden to eat beef/?cow
c. The table is made of wood/?tree

The important point is the observation that blocking is not absolute, but may be cancelled under special contextual conditions. That is, we find cases of *deblocking*. Nunberg & Zaenen (1992) consider the following example:

(8) Hindus are forbidden to eat cow/?beef

They argue that "what makes *beef* odd here is that the interdiction concerns the status of the animal as a whole, and not simply its meat. That is, Hindus are forbidden to eat beef only because it is cow-stuff." (Nunberg & Zaenen 1992: 391). Examples of this kind strongly suggest that the blocking phenomenon is pragmatic in nature. Furthermore, these examples suggest that (pragmatic) anomaly does not necessarily persist when specific contextual information is added. Copestake & Briscoe (1995) provide further examples substantiating this claim.

Previously, I introduced a second variant of the notion of persistent anomaly. It concerns the specificity of linguistic information and less that of contextual information. There is a variety of examples showing that this variant of the persistence of (pragmatic) anomaly likewise must fail (cf. Nunberg & Zaenen 1992):

(9) a. This wine is particularly good with ?mammal/lamb
b. ?Mammal/canine is healthy food
c. She likes to wear ?mammal/?sheep/angora

5 The division of labor between (lexical) semantics and pragmatics

In the following, I take the distinction between the lexicon and the encyclopedia as an important one. With the same attitude I look at the general distinction between semantics and pragmatics. I will use the features reported in the last three sections for marking out the borderline between lexical semantics and pragmatics. This opens the possibility to make the discrimination between lexical information and encyclopedic information empirically more comprehensible.

The most notable feature in the present discussion is compositionality. Many researchers agree in seeing compositionality as a principle satisfied at the semantic level of representation but violated at the level of utterance interpretation. Respecting the non-compositionality of utterance interpretation, several of these researchers seem to consider it virtuous and advantageous to deviate from compositionality in a minimalist way. A typical approach following this path of virtue is the so-called *coercion view* (e.g. Pustejovsky 1989, 1991, 1993, 1995). This approach can be contrasted with another view that is more radically founded on underspecified representations. I will call it the *radical underspecification view* (e.g. Hobbs 1983, Pinkal 1995, Alshawi 1990, Poesio 1991, and others). Generally, the two approaches can be characterized and contrasted as shown in the following table:

The coercion view	The radical underspecification view
Every lexical unit determines a primary conceptual variant which can be grasped as its (literal) meaning.	Every lexical unit determines an underspecified representation (i.e. a representation that may contain, for example, place holders and restrictions for individual and relational concepts)
The combinatorial system of language determines how the lexical units are combined into larger units (phrases, sentences).	The combinatorial system of language determines how lexical units are combined into larger units (phrases, sentences).
There is a system of type and sortal restrictions which determines whether the resulting structures are well-formed.	There is a system of type and sortal restrictions which determines whether structures of a certain degree of (under)specification are well-formed.

There is a generative device (called type/sort coercion) that tries to overcome type or sortal conflicts that may arise by strict application of the combinatorial system of language. The coercion device is triggered (only) by type or sort violations.	There is a mechanism of contextual enrichment (pragmatic strengthening based on contextual and encyclopedic knowledge). This inferential mechanism is controlled by cost factors and doesn't need triggering by type or sort violations.

In Blutner (1997) several problems about the coercion view in general and Pustejovsky's account in particular are discussed. Some cues are: combinatorial explosion puzzle, restrictiveness of the coercion mechanism, problem of defeasibility, inflation of shifting operation. Taken together, these problems suggest that it is more promising to favour the radical underspecification view.

It should be added that the radical underspecification view shares some ideas with Bierwisch's (1983) two-level semantics: (i) the distinction between lexicon and encyclopedia, i.e. between semantics and pragmatics is taken as an important one, (ii) the features of compositionality, monotonicity, and (perhaps) persistence of anomaly, are taken as crucial characteristics marking out the domain of semantics.

However, in contrast to Bierwisch's two-level semantics (and Pustejovsky's generative lexicon), the radical underspecification view disregards monotonic unification and type/sort coercion as mechanisms of contextual enrichment. Instead, it explores alternative proposals stressing open-ended default inference on real world knowledge. Here is a collection of candidates that may provide a suitable mechanism for the contextual enrichment of underspecified representations:

- Defaults as rules for filling in information gaps (see various papers in van Deemter & Peters, 1995)
- Discourse interpretation based on a default conditional logic (e.g. Lascarides & Asher 1993)
- Persistent Default-Unification (Lascarides, Briscoe, Asher & Copestake 1995, Copestake & Briscoe 1995)
- Weighted abduction (Hobbs et al. 1993)
- Conversational implicature and lexical pragmatics (Blutner, Leßmöllmann & van der Sandt 1996, Blutner 1997)

In the following, I will refer to the last-mentioned account as an explicit variant of the radical underspecification view. This view starts with certain underspecified representations. In the simplest case these are representa-

tions containing a special type of individual expressions (indicated by $\underline{x}$, $\underline{y}$, $\underline{z}$,...) which can be seen as *placeholders* for individuals (or individual concepts). As a general rule, the semantic value of an underspecified expression is a set of elements, where each element is the value of a corresponding ordinary expression.[6] For the sake of illustration, I will consider the example (10a). Roughly simplifying, I assume that the formula (10b) is the corresponding underspecified representation. [7]

(10) a. red apple
b. $\lambda y\,[\mathrm{APPLE}(y) \wedge \mathrm{PART}(y, \underline{x}) \wedge \mathrm{COLOR}(\underline{x}, \mathrm{RED})]$

Here, $\underline{x}$ indicates the placeholder for the part of the apple y that is red. This λ-expression may be applied to a term. The semantic rules for the composition of two expressions are the ordinary ones lifted to the level of sets of ordinary values. For the operation of application, as an example, we have the following:

(11) Let $\underline{P}$ be a predicate expression and $\underline{t}$ a term (both possibly underspecified):

$||\underline{P}(\underline{t})||_M = \{\Pi(o): \Pi \in ||\underline{P}||_M \text{ and } o \in ||\underline{t}||_M\}$
(The index M refers to the model)

It is obvious that the semantic value of *underspecified sentences* comes out as one of four possible truth values: {1} true, {0} false, {1,0} underspecified, {} uninterpretable.

A sentence like (12a) translates as (12b). In accordance with the application rule, it comes out as true if all the parts of the apple are red. On the other hand, it is false if no part of the apple is red.

(12) a. This is a red apple
b. $\mathrm{APPLE}(\mathrm{THIS}) \wedge \mathrm{PART}(\mathrm{THIS}, \underline{x}) \wedge \mathrm{COLOR}(\underline{x}, \mathrm{RED})$

---

6 At first glance, the use of *placeholders* $\underline{x}$, $\underline{y}$, $\underline{z}$ shares some features with the *free variable view* mentioned in section 2. However, the semantic value of *placeholders* is of a completely different kind than the value of free variables (according to the standard treatment).

7 For the moment, it doesn't matter in which way the underspecified representation (10b) is projected from the lexicon. One possible choice would be to include the PART-feature into the lexical representation of the color-predicate. Another choice leaves this feature completely outside the lexicon and includes it in the rules for adjectival modification.

In case some parts of the apple are red and others are not, the sentence is neither true nor false. Instead, its truth value is {0,1} (underspecified). An interesting consequence of this view is that we can say that color predicates are *divisive* (with regard to the mereological structure of objects like apples).

The example illustrates in which way a very simple *compositional* semantics can be defined. The price we have to pay for the simplicity is that the truth conditions defined so far give only a very weak and incomplete characterization of the states of affairs where sentences like (12a) are *intuitively* true. Obviously, it's not necessary that all parts of our apple must be red in order to call (12a) true. It is sufficient (under normal circumstances) that it has a red surface. On the other hand, already from the observation that (the most part of) the surface is green, we can conclude that (12a) is false. (Imagine an apple with a green peel which has red pulp).

The other two principles (monotonicity and persistence of anomaly) are satisfied as well when the semantic component is construed in terms of underspecification as illustrated before. The important point is the definition of the entailment relation in case of underspecified expressions. Let φ and ψ be formulas representing underspecified sentences. I consider the following two possibilities to define a (monotonic) entailment relation[8]:

(13) a. $\varphi \models^{s} \phi$ iff $\|\varphi\|_M = \{1\}$ implies $\|\phi\|_M = \{1\}$, for each M
b. $\varphi \models^{w} \phi$ iff $\|\varphi\|_M \neq \{0\}$ implies $\|\phi\|_M \neq \{0\}$, for each M

According to the first notion, a sentence such as *this is a red apple* would entail that *each part of the apple is red*. Intuitively, this notion of entailment seems to be far too strong. The second (weak) notion corresponds better to our intuitions: *This is a red apple* entails only that *a part of the apple is red.* Let's take definition (13b) as the appropriate one. This choice makes one point clear again: The semantic part of the machinery gives only very weak determination of the kind of knowledge we intuitively connect with the meaning of contentful expressions. A full account of (lexical)

[8] Van Deemter (1996, 1997), among others, makes some tentative efforts in investigating the whole spectrum of possible entailment relations that can be defined in the case of underspecification.

meaning has to include more information than that which establishes a truth-conditional, compositional semantics.

What is needed is a suitable mechanism for the contextual enrichment of underspecified representations. Such a mechanism is intended to resolve the task of discourse interpretation, and it carries the main burden in explaining restrictions on interpretation. Due to its inferential character, this mechanism is structured non-compositionally. In the next section, I will describe a mechanism of pragmatic strengthening which is based upon a reformulation of the Gricean idea of conversational implicature.

6 Conversational implicature: Outline of a reformulation

For Griceans, conversational implicatures are those non-truth-functional aspects of utterance interpretation which are conveyed by virtue of the assumption that the speaker and the hearer are obeying the *cooperative principle of conversation*, and, more specifically, various *conversational maxims*: *maxims of quantity, quality, relation* and *manner*.

While the notion of conversational implicature doesn't seem hard to grasp intuitively, it has proven difficult to define precisely. The generality of the cooperative principle and the conversational maxims makes it difficult to specify just which maxims are involved in particular implicatures. Essential concepts mentioned in the maxims are left undefined (what is *relevance, adequate evidence*, etc.). However, before we can start to "flesh out" something like the maxims, Grice's view of implicature raises even more basic questions. Are there just the maxims Grice mentioned, or might others be needed (as he suggested himself)? Or could the number of maxims be reduced? Sperber & Wilson (1986) are an extreme case in suggesting one only, the maxim of relevance. And what is the rational behind the cooperative principle and the maxims? Are they norms which speakers and hearers must *know* in order to communicate adequately (as Grice and most followers suggest)? Or are they generalizations about certain forms of inferential behavior which speakers and hearers need no more to know to communicate than they need to know the principles of digestion to digest (Sperber and Wilson's position, a position which is also adopted in the present account). An important step in reducing and explicating the Gricean framework has been made by Atlas and Levinson (1981) and Horn (1984). Taking quantity as starting point they distinguish between two

principles, the Q-principle and the I-principle (termed R-principle by Horn 1984). Simple but informal formulations of these principles are as follows:

Q-principle:

- Say as much as you can (given I) (Horn 1984: 13).
- Make your contribution as informative (strong) as possible (Matsumoto 1995: 23).
- Do not provide a statement that is informationally weaker than your knowledge of the world allows, unless providing a stronger statement would contravene the I-principle (Levinson 1987: 401).

I-principle:

- Don't say more than you must (given Q) (Horn 1984: 13).
- Say as little as necessary, i.e. produce the minimum linguistic information sufficient to achieve your communicational ends (bearing the Q-principle in mind) (Levinson 1987: 402).
- Read as much into an utterance as is consistent with what you know about the world (Levinson 1983: 146-147).

Obviously, the Q-principle corresponds to the first part of Grice's quantity maxim (*make your contribution as informative as required*), while it can be argued that the countervailing I-principle collects the second part of the quantity maxim (*do not make your contribution more informative than required*), the maxim of relation and possibly all the manner maxims. As Horn (1984) seeks to demonstrate, the two principles can be seen as representing two competing forces, one force of *unification* minimizing the Speaker's effort (I-principle), and one force of *diversification* minimizing the Auditor's effort (Q-principle).

Conversational implicatures which are derivable essentially by appeal to the Q-principle are called Q-based implicatures. Standard examples are given in (14):

(14) a. Some of the boys are at the party
> Not all of the boys are at the party

b. I'm happy
> I'm happy but not ecstatic
(*Scalar implicatures*, Gazdar 1979)

c. Rick is a philosopher or a poet
> Rick is not both a philosopher and a poet
(*Scalar implicatures*, Grice 1968; Atlas & Levinson 1981)

d. Rick is a philosopher or a poet
> Rick may (not) be a philosopher
> Rick may (not) be a poet
(*Clausal implicatures*, Gazdar 1979; Atlas & Levinson 1981)

e. Black Bart caused the sheriff to die
> Black Bart didn't kill the sheriff in a direct way (McCawley 1978)

Conversational implicatures which are derivable essentially by appeal to the I-principle are called I-based implicatures. Some standard examples are given in (15):

(15) a. John turned the key and the engine started
>John turned the key and then the engine started (temporal sequence)
> John turned the key and therefore the engine started (causal connectedness)
(*Conjunction buttressing*, Atlas & Levinson, 1981)

b. If you mow the lawn, I'll give you $5
> If and only if you mow the lawn, will I give you $5 .
(*Conditional perfection*, Geis & Zwicky 1971)

c. John unpacked the picnic. The beer was warm.
> The beer was part of the picnic.
(*Bridging*, Clark & Haviland, 1977)

d. John said 'Hello' to the secretary and then he smiled.
>John said 'Hello' to the female secretary and then he smiled.
(*Inference to stereotype*, Atlas & Levinson 1981)

e. Harry and Sue bought a piano.
> They bought it together, not one each.
(*Mirror maxim*, Harnish, 1976)

f. John came in and he sat down
> $John_1$ came in and he_1 sat down

(Preferred Co-reference)

I believe that the proper treatment of conversational implicature crucially depends on the proper formulation of the Q- and the I-principle. Furthermore, such formulation also has to account for the interplay between these two principles and their interaction with the quality maxim. The explication developed in Blutner, Leßmöllmann & van der Sandt (1996) and Blutner (1997) rests on the assumption that the underspecified semantic description sem(α) of an utterance α determines a whole range of possible enrichments m, one of which covers the intended content. In order to develop a fairly explicit notion of possible enrichments, the idea of abductive specification (cf. Hobbs, Stickel, Appelt, & Martin 1993) has been used. Loosely speaking, <α, m> is called a *possible enrichment pair* if m is an abductive specification of sem(α) that can be generated by means of general world knowledge and discourse knowledge.[9] Weighted abduction gives for each possible enrichment pair a cost value $\underline{c}(\alpha, m)$ that reflects the "proof" cost for deriving m from sem(α). Roughly, this cost is correlated with the *surprise* the particular enrichment has for an agent confronted with the underspecified representation sem(α).

The Q- and the I-principle can be seen as conditions constraining possible enrichment pairs <sem(α), m>:

(16) Constraining *possible enrichment pairs*:

a. <α, m> satisfies the Q-principle iff <α, m> is a possible enrichment pair and there is no other <α', m> [satisfying the I-principle] such that $\underline{c}(\alpha', m) < \underline{c}(\alpha, m)$.

b. <α, m> satisfies the I-principle iff <α, m> is a possible enrichment pair and there is no other <α, m'> [satisfying the Q-principle] such that $\underline{c}(\alpha, m') < \underline{c}(\alpha, m)$.

---

[9] With regard to the model-theoretic semantics of underspecification, we must be careful to distinguish between specification and precisification (cf. Pinkal 1995). Specifying a formula reduces the set of models where the formula is true (but typically extends the set of models where it is false). Making a formula more precise, on the other hand, reduces the set of models where the formula has a gap (= underspecified truth-value {0,1}). It is not difficult to see that we get precisification if we take the mechanism of abductive specification and treat the *placeholders* as Skolem-functions within that formalism. For example, within the domain of fruits we can apply the "rule" PART(y,$\underline{x}$) ← PEEL(y,$\underline{x}$) abductively to the expression PART(a,$\underline{x}$). The result is the formula PEEL(a,$\underline{x}$), which is *more precise* than PART(y, $\underline{x}$).

In this (rather symmetrical) formulation, the Q- and the I-principle constrain possible enrichments in two different ways. The I-principle constrains them by selecting the *minimal surprising* enrichments (provided Q has been satisfied), and the Q-principle constrains them by blocking those enrichments which can be grasped more economically by an alternative linguistic input α' (provided I has been satisfied).

It is not difficult to see that the Q-principle carries the main burden in explaining the blocking effects discussed in section 3. The additions put in brackets were introduced to explain the "division of pragmatic labor" (Horn 1984): The use of marked expressions, when a corresponding unmarked expression is available, tends to be interpreted as conveying a marked message. (Recall, for example, the case of productive causatives, as illustrated in (6); for a detailed explanation the reader is referred to Blutner 1997).

Now, I informally introduce the notion of *common ground*, an information state containing all the propositions shared by several participants, including general world knowledge and discourse knowledge. For simplicity, I take common grounds as sets of formulas, and I will assume that the following condition is satisfied for common grounds cg:

(17) $\varphi \in cg \Leftrightarrow \mathbf{K}_S(\varphi) \in cg$ and $\mathbf{K}_H(\varphi) \in cg$ (for each formula φ)

$\mathbf{K}_i(\varphi)$ should be read as "agent i believes that φ". For the addition of some (new) formula ϕ to the common ground I write cg[ϕ]. If cg is a common ground, then cg[ϕ] doesn't necessarily satisfy the condition (17). I assume an operation * which closes off a set of formulas as a common ground. The formula cg[ϕ]* then expresses the updating of the common ground with ϕ (such that the resulting set again counts as a common ground). For the task of giving a precise definition of the closure operation, I refer the reader to Gerbrandy (1997) and Zeevat (1997). Informally, we can see the operation * as providing the deductive closure of a set of formulas σ, where at the same time all expressions of the form $\mathbf{K}_S(\varphi)$ and $\mathbf{K}_H(\varphi)$ are added to the set provided φ is in the set. The closure operation * sometimes yields the whole language. I call this the *impossible information state* and indicate it by *nil*. This is the case, for example, for inconsistent sets such as $\{p, \neg p\}$,

or for *epistemically indefensible* sets[10] such as {p, $\neg \mathbf{K}_S(p)$}. Now the following definitions can be stated:

(18) Definition of basic pragmatic notions

a. A possible enrichment pair <α,m> is called *pragmatically licensed* (in a common ground cg) iff <α,m> satisfies the Q- and the I-principle and cg[m]* is different from *nil*.[11]

b. An utterance α is called *pragmatically sound/anomalous* (in cg) iff there is a/no pragmatically licensed enrichment pair <α,m>.

b. A propositional formula ϕ is called a *conversational implicature* of α (in cg) if ϕ is an element of cg[m]* for each m of a pragmatically licensed enrichment pair <α,m>.

Figure 2 shows the basic idea underlying these definitions in a schematic form. In order to check whether an utterance α (with regard to a given common ground cg) is pragmatically anomalous or not, we have to proceed in two steps. In the first step all the enrichments are selected which satisfy both the I- and the Q-principle. This is the step where Horn's division of pragmatic labor takes place. For simplicity, I will assume for the moment that there is only one selected enrichment m_i. In the second step, this enrichment is added to cg and the whole set is closed off as a common ground. In case the resulting common ground is different from *nil*, the utterance α is proven to be pragmatically sound; in the other case it is proven to be pragmatically anomalous. Conversational implicatures then simply can be taken as elements of cg $[m_i]$*.

[10] Hintikka (1962) calls a proposition φ *epistemically indefensible* just in case $\mathbf{K}_S(\varphi)$ is inconsistent (with regard to his epistemic logic system).

[11] I have explained elsewhere (Blutner 1997) that the last clause in this definition reflects Grice's *maxim of quality*.

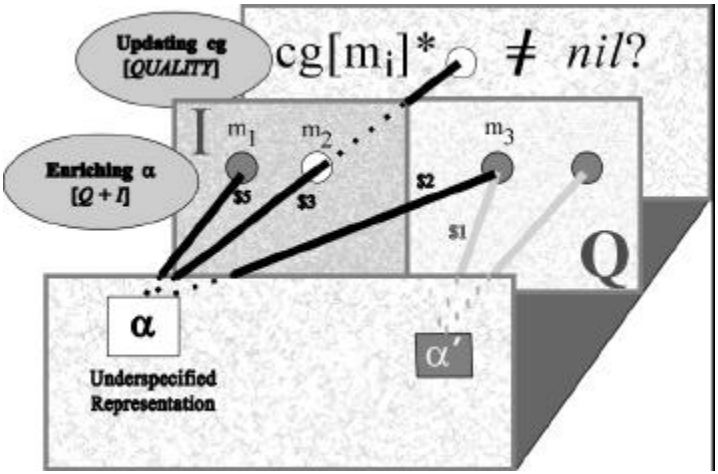

Figure 2:
Enriching underspecified representations and updating the common ground

It is not difficult to see in which way the general mechanism of conversational implicature reflects (a) the non-compositional aspect of utterance interpretation, (b) the phenomena of blocking and deblocking, and (c) the general fact that pragmatic anomalies usually don't persist.

First, let's consider compositionality. Almost everything in the formulation of conversational implicature has non-compositional character: Both the formulation of the Q-principle as well as the I-principle are "holistic" in addressing a whole range of alternative expressions; the conceptions of informativeness, surprise (measured in terms of conditional probability) and linguistic complexity are non-combinatorial and cannot be reduced to the corresponding properties of the parts of an expression; the mechanism of information enrichment as based on abductive inference is non-compositional (see Blutner 1997).

Next, our system deals with blocking and deblocking. The crucial mechanism involved is due to the Q-principle. In the same way, the present system captures the "field" effects, which are very important if the extensions of lexical concepts are considered.

The third point concerns the persistence of anomalies. The general definition of pragmatic anomaly doesn't simply define this notion as some kind of inconsistency. Instead, non-representational parameters (such as

surprise, cue validity, relevance, frequency of use, etc.) are crucially involved in controlling the selection and suppression of possible enrichments. Within this setting, typically some kind of garden path effect may arise . This constitutes the pragmatic anomaly.

What follows is a brief illustration of how this framework can be used to solve some of the lexical-pragmatic phenomena considered in the first part of the paper. I consider two of Quine's puzzles concerning the pragmatics of adjectives (see section 2). The first one concerns the observation that the (preferred) interpretation of adjective noun combinations seems to affect different parts of the subject term in cases like (19a,b). The second puzzle has to do with the explanation of pragmatic anomalies in examples like (19c), where it is very difficult to get interpretation (19d).

(19) a. The apple is red [interpretation: its peel is red]
b. The apple is sweet [interpretation: its pulp is sweet]
c. ?The tractor is pumped up
d. The tires of the tractor are pumped up

In order to sketch how the mechanism solves the first puzzle, let us concentrate on example (19a).[12] Input of the analysis is an underspecified representation expressing that a certain part of the apple is red (roughly: APPLE(d) $\wedge$ PART(d,$\underline{x}$) $\wedge$ COLOR($\underline{x}$,$\underline{u}$) $\wedge$ $\underline{u}$=RED). The precisification of the relevant part(s) is guided by parameters of subjective probability (cue validity, diagnostic value). For example, it is plausible to assume that the color of the peel is more diagnostic for classifying apples than the color of other apple parts (such as the color of the pulp). From this assumption it can be derived that the *red peel*-enrichment is the cost minimum one. Consequently, the I-principle selects the *red peel*-enrichment (and blocks the *red pulp*-enrichment). It follows that the proposition expressing that the peel of the apple is red is a conversational implicature of (19a) (but not the proposition expressing that the pulp of the apple is red). In the case of (19b), analogous considerations give the *sweet pulp*-enrichment as the preferred interpretation.

Next, what about the pragmatic anomaly in cases like (19c), which contrast with examples like (19d) which are acceptable? Surely, the underspecified semantics of (19c) (saying that some part of the tractor is pumped up) isn't inconsistent with usual background knowledge. If it was, the

[12] The technical details of the analysis are given elsewhere (Blutner 1997).

sentence (19d) should be deviant in the same way. Consequently, the pragmatic anomaly of (19c) must be explained in another way. I think it follows from the fact that those parts of tractors that may be pumped up (the tires) are only marginally diagnostic for classifying tractors. If this is correct, the *pumped up tires*-enrichment is blocked by enrichments that refer to more salient parts (such as the motor or the coachwork). However, the latter enrichments suffer from sort conflicts and therefore come out as not pragmatically licensed. In summary, a kind of garden path effect brings about that (19c) is pragmatically anomalous.

This example nicely helps to see that the present notion of anomaly isn't persistent in general. The anomaly can be canceled under special contextual conditions. For example, imagine the situation in a garage where we find tractors whose tires are pumped up and tractors whose tires are not. In this situation sentence (19c) sounds fine. The explanation is that the pressure state of the tires in this situation should be highly diagnostic for classifying tractors.

The approach just sketched can be extended to analyze corresponding effects in case of systematic polysemy (see Nunberg 1979, Bierwisch 1984, Blutner 1997). For example, we can apply the analysis for explaining the contrast between (20a) and (20b).

(20) a. The school has a flat roof.
b. ?The government has a flat roof.

In a first approximation, the underspecified semantic representation of these sentences may look like (21):

(21) a. $\mathrm{SCHOOL}(e) \wedge \mathrm{REALIZE}(e,\underline{x}) \wedge \mathrm{HFR}(\underline{x})$
b. $\mathrm{GOV}(e) \wedge \mathrm{REALIZE}(e,\underline{x}) \wedge \mathrm{HFR}(\underline{x})$

In these examples, *e* refers to a certain entity which can be understood as a conceptual frame and can be classified according to the variety of *institute*-type (*school* in (21a), *government* in (21b)). It is further assumed that this entity can be considered under different *perspectives.* These perspectives are assumed to provide more concrete *realizations* of the rather abstract concept of a certain *institute*-type e, perhaps realized as building, process or institution property. The role of the placeholder $\underline{x}$ is to guarantee that the particular perspective adopted remains underspecified, and, consequently, the concrete realization of the intended *institute*-type remains

semantically open. The precisification of the relevant realization(s) is guided again by parameters of subjective probability. For the example of *school*, let's consider two possible precisifications: One corresponding to the institution realization and the other to the building realization. It is plausible to assume that the saliences of the two realizations are of comparable strength. Evidence for this hypothesis comes from prototype semantics. The building and the institution reading of *school* can be seen as realizing basic level concepts of both buildings and institutions. However, salience is not the only factors that effects the costs of deriving the two variants. Another factor is what can be saved in assumption costs, supposed the selection restrictions of the predicate are satisfied. Since the predicate *has a flat roof* selects for the category of buildings, it results that the building-variant is less expensive than the institution variant. Consequently, the I-principle selects the building-reading (and blocks the institution-reading).

What about the pragmatic anomaly in cases like (20b), which contrast with examples like (22) which are acceptable?

(22) The government building has a flat roof.

In this case, we have a strong imbalance between the salience of the institution realization and that of the building realization. This follows from the assumption that governments as institutions are certainly basic level institutions, whereas government buildings are not basic level buildings (at least to my "cultural background" which is not that of Munich or Saarbrücken). Consequently, it is plausible to assume that the (conceptually inconsistent) institution variant is less expensive than the (conceptually consistent) building-variant. In this case, the I-principle selects the inconsistent variant. Again we have a kind of garden path effect: Updating the common ground leads to the impossible information state. And this explains why (20b) is pragmatically anomalous (for details, see Blutner 1997).

Finally, I want to suggest that the present exercise of working out the interplay between lexical underspecification and pragmatic strengthening may be useful and applicable when it comes to consider word formation in general (e.g. Aronoff 1976, Bauer 1983) and the interpretation of compounds in particular (e.g. Meyer 1993, Wu 1990). Moreover, the investigation of other kinds of polysemy than those sketched here may be helpful in order to see the ubiquity of these problems (cf., for instance,

Lakoff's (1987) study on English prepositions and Sweetser's (1990) investigation of English perception verbs).

As a matter of priority, methods are needed allowing to measure the values of the probabilistic parameters that control and organize conceptual knowledge. It is these non-representational parameters that crucially determine and control the computational system of pragmatic strengthening.

References

Alshawi, H. (1990): "Resolving quasi logical forms". Computational Linguistics 16, 133-144.

Aronoff, M. (1976): Word Formation in Generative Grammar. Cambridge, Mass.: MIT Press.

Atlas, J. & Levinson, S. (1981), "*It*-clefts, informativeness and logical form". In: P. Cole (ed): Radical Pragmatics. New York: Academic Press, 1-61.

Bauer, L. (1983): English Word-Formation. Cambridge: Cambridge University Press.

Bierwisch, M. (1983): "Semantische und konzeptuelle Repräsentation lexikalischer Einheiten". In: W. Motsch & R. Ruzicka (Hrsg.): Untersuchungen zur Semantik. Berlin: Akademie Verlag, 61-99.

Bierwisch, M. (1989), "The semantics of gradation". In: M. Bierwisch & E. Lang (eds.): Dimensional Adjectives. Berlin: Springer-Verlag, 71-261.

Blutner, R. (1997): "Lexical pragmatics". In: R. van der Sandt, R. Blutner, and M. Bierwisch (eds.): From Underspecification to Interpretation. IBM Working Papers of the Institute for Logic and Linguistics, Nr. 29. Heidelberg: IBM Germany, 55-89.

Blutner, R., Leßmöllmann, A., and van der Sandt, R. (1996): "Conversational implicature and lexical pragmatics". In: Proceedings of the AAAI Spring Symposium on Conversational Implicature. Stanford, 1-9.

Carnap, R. (1936): "Testability and meaning". Philosophy of Science 3, 419-471.

Clark, H. & Haviland, S. (1977): "Comprehension and the given-new contract". In: R. Freedle (ed): Discourse Production and Comprehension. Norwood, NJ: Ablex.

Copestake, A. & Briscoe, T. (1995): "Semi-productive polysemy and sense extension". Journal of Semantics 12, 15-67.

Cruse, D.A. (1986): Lexical Semantics. Cambridge: Cambridge University Press.

Deane, P.D. (1988): "Polysemy and cognition". Lingua 75, 325-361.

Dowty, D. (1979): Word Meaning and Montague Grammar. Dordrecht: Kluwer.

Fodor, J.A. & Pylyshyn Z.W. (1988): "Connectionism and cognitive architecture: a critical analysis". Cognition 28, 3-71.

Fodor, J.D. (1977): Semantics: Theories of Meaning in Generative Grammar. New York: Crowell.

Gärdenfors, P. (1996): "Conceptual spaces and cognitive semantics". Course material presented on the eighth European Summer School in Logic, Language, and Information. Prague.

Gazdar, G. (1979): Pragmatics. New York: Academic Press.

Geis, M. & Zwicky, A.M. (1971): "On invited inference". Linguistic Inquiry 2, 561-566.

Gerbrandy, J. (1997): "Changing the common ground". In: A. Benz & G. Jäger (eds.): Proceedings of the Munich Workshop on Formal Semantics and Pragmatics of Dialogue. University of Munich, 40-58.

Grice, P. (1968): Logic and Conversation. Text of Grice's William James lectures at Harvard University. Published in: Grice (1989): Studies in the Way of Words, Cambridge, MA: Harvard University Press.

Harnish, R.M. (1976): "Logical form and implicature". In T. Bever, J. Katz & T. Langendoen (eds.): An Integrated Theory of Linguistic Ability. New York: Crowell, 464-479.

Hobbs, J.R. (1983): "An improper treatment of quantification in ordinary English". In: Proceedings of ACL-83. Cambridge, MA, 57-63

Hobbs, J.R., Stickel, M.E., Appelt, D.E., & Martin, P. (1993): "Interpretation as abduction". Artificial Intelligence 63, 69-142.

Horn, L.R. (1984): "owards a new taxonomy for pragmatic inference: Q-based and R-based implicatures". In: D. Schiffrin (ed.): Meaning, Form, and Use in Context. Washington: Georgetown University Press, 11-42.

Householder, F.W. (1971): Linguistic Speculations. London and New York: Cambridge University Press.

Kamp, H. (1975), "Two theories about adjectives". In: E.L. Keenan (ed.): Formal Semantics for Natural Language. Cambridge: Cambridge University Press, 123-155.

Katz, J.J. & Fodor, J.A. (1963): "The structure of semantic theory". Language 39, 170-210.

Keenan, E.L. (1974): "The functional principle: Generalizing the notion of Subject of". In: Papers from the Tenth Regional Meeting of the Chicago Linguistic Society. Chicago: Illinois, 298-310.

Kiparsky, P. (1982), "Word-formation and the lexicon". In F. Ingeman (ed.): Proceedings of the 1982 Mid-America Linguistic Conference.

Lahav, R. (1989): "Against compositionality: the case of adjectives". Philosophical Studies 55, 111-129.

Lahav, R. (1993): "The combinatorial-connectionist debate and the pragmatics of adjectives". Pragmatics and Cognition 1, 71-88.

Lakoff, R. (1987): Women, Fire, and Dangerous Things: What Categories Reveal About the Mind. Chicago: University of Chicago Press.

Lascarides, A. & Asher, N. (1993): "Temporal interpretation, discourse relation, and common sense entailment". Linguistics and Philosophy 16, 437-494.

Lascarides, A., Briscoe, T., Asher, N., & Copestake, N. (1995): "Order independent and persistent typed default unification". Linguistics and Philosophy 19, 1-90.

Leßmöllmann, A. (1996): Das Problem der Modifikation polysemer Ausdrücke. Magister Artium, Humboldt University Berlin.

Levinson, S. (1983): Pragmatics: Cambridge, England: Cambridge University Press.

Levinson, S. (1987): "Pragmatics and the grammar of anaphora". Journal of Linguistics 23, 379-434.

Matsumoto, Y. (1995), "The conversational condition on Horn scales". Linguistics and Philosophy 18, 21-60.

McCawley, J.D. (1971): "Interpretive semantics meets Frankenstein". Foundations of Language 7, 285-296.

McCawley, J.D. (1978): "Conversational implicature and the lexicon". In: P. Cole (ed.): Syntax and Semantics 9: Pragmatics. New York: Academic Press, 245-259.

Meyer, R. (1993): "Compound Comprehension in Isolation and in Context". Tübingen: Max Niemeyer Verlag.

Montague, R. (1970): "Universal Grammar". Theoria 36, 373-398.

Nunberg, G. (1979): "The non-uniqueness of semantic solutions: Polysemy". Linguistics and Philosophy 3, 143-184.

Nunberg, G. & Zaenen, A. (1992): "Systematic polysemy in lexicology and lexicography". In: K. Varantola, H. Tommola, T. Salmi-Tolonen and J. Schopp (eds.): Euralex II. Tampere: Finland.

Partee, B. (1984): "Compositionality". In: F. Landman and F. Veltman (eds.): Varieties of Formal Semantics. Dordrecht: Foris, 281-311.

Pinkal, M. (1995): Logic and Lexicon. Dordrecht: D. Reidel.

Poesio, M. (1991): "Relational semantics and scope ambiguity". In: J. Barwise, J.M. Gawron, G. Plotkin, and S. Tutiya (eds.): Situation semantics and its Applications, Vol. 2. Stanford, CA: CSLI, 469-497.

Pustejovsky, J. (1989): "Type coercion and selection". Paper presented at WCCFL VIII, April 1989, Vancouver. B.C.

Pustejovsky, J. (1991): "The generative lexicon". Computational Linguistics 17, 409-441.

Pustejovsky, J. (1993): "Type coercion and lexical selection". In: J. Pustejovsky (ed.): Semantics and the Lexicon. Dordrecht: Kluwer, 73-96.

Pustejovsky, J. (1995): The Generative Lexicon. Cambridge, Mass.: The MIT Press.

Quine, W.V.O. (1960): Word and Object. Cambridge, Mass.: MIT Press.

Sperber, D. & Wilson, D. (1986): Relevance. Oxford: Blackwell.

Sweetser, E.E. (1990): From Etymology to Pragmatics. Cambridge: Cambridge University Press.

van Deemter, K. (1996): "Towards a logic of ambiguous expressions". In: K. van Deemter and S. Peters (eds.): Semantic Ambiguity and Underspecification. Stanford, California: CSLI Publications.

van Deemter, K. (1997): "Ambiguity and the principle of idiosyncratic interpretation". In: R.van der Sandt, R. Blutner, and M. Bierwisch (eds.): From Underspecification to Interpretation. Working Papers of the Institute for Logic and Linguistics, Nr. 29. Heidelberg: IBM Germany, 91-111.

van Deemter, K. & Peters, S. (1996): Semantic Ambiguity and Underspecification. Stanford, California: CSLI Publications.

Wu, D. (1990): "Probabilistic unification-based integration of syntactic and semantic preferences for nominal compounds". Proceedings of the 13th International Conference on Computational Linguistics (COLING 90), Helsinki, 413-418.

Zeevat, H. (1997): "The common ground as a dialog parameter". In: A. Benz & G. Jäger (eds.): Proceedings of the Munich Workshop on Formal Semantics and Pragmtics of Dialogue. University of Munich, 195-214