Conversational Implicature and Lexical Pragmatics

Reinhard Blutner and Annette Leßmöllmann

Max-Planck-Research Group
at the Humboldt-University
Jägerstraße 10/11, 10117 Berlin, Germany

Rob van der Sandt

University of Nijmegen
Dept. of Philosophy
P.O. Box 9103, 6500 HD Nijmegen, The Netherlands

Abstract

Lexical Pragmatics is a research field that tries to give a systematic and explanatory account of a number of pragmatic phenomena that are connected with the semantic underspecification of lexical items. The approach combines a constraint-based semantics with a Gricean mechanism of pragmatics. The basic pragmatic mechanism rests on conditions of updating the common ground and allows to give a precise explication of notions as (generalized) conversational implicatures and pragmatic anomaly. The integrating and unifying character of the basic account is demonstrated by solving some puzzles with regard to Atlas & Levinson's Q- and I-principles, Horn's division of pragmatic labor, etc. The basic mechanism is then extended by an abductive reasoning system (guided by subjective probability) and it is illustrated how this framework can be used to solve Quine's riddles concerning the pragmatics of adjectives.

1 Introduction

Lexical Pragmatics is a research field that tries to give a systematic and explanatory account of pragmatic phenomena that are intimately connected with the semantic underspecification of lexical items. Cases in point are the interpretation of compounds, systematic polysemy, the distribution of lexical and productive causatives (McCawley 1978), blocking phenomena (Horn 1984), the pragmatics of adjectives (Lahav 1993) and many phenomena presently discussed within the framework of Cognitive Semantics (Lakoff 1987). The approach combines a compositional semantics with a Gricean mechanism of pragmatics. Starting off from an underspecified semantic representation an abductive mechanism is invoked to yield the appropriate specification with regard to the common ground.

The setup of this paper is as follows. We first give an account of conversational implicature that rests on pragmatic conditions of updating the common ground. Then we discuss some consequences of the basic mechanism demonstrating the integrating and unifying character of the basic approach. In section 4 we extend the basic mechanism by including an abductive reasoning system guided by subjective probability. Finally, we present an example illustrating how this framework can be used to solve a specific riddle of the pragmatics of adjectives.

2 Toward the Proper Treatment of Conversational Implicature

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, the various *conversational maxims*: *maxims of quantity, quality, relation* and *manner*. While the notion of conversational implicature seems not hard to grasp intuitively, it has proven difficult to define precisely. 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). A simple but informal formulations of these principles is as follows:

> <u>Q-principle</u>: Say as much as you can (given I). (Horn 1984: 13)
> 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: Say no more than you must (given Q). (Horn 1984: 13)
Say as little as necessary, i.e. produce the minimal linguistic information sufficient to achieve your communicational ends (bearing the Q-principle in mind). (Levinson 1987: 402)

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 is 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).

We guess that the proper treatment of conversational implicatures crucially depends on the proper formulation of the Q- and the I-principle. As we will demonstrate subsequently, such a formulation has to account for the interplay of these two principles and their interaction with the two quality maxims. The present explication rests on the assumption that the semantic description sem(α) of an utterance α is an underspecified representation determining a whole range of possible specifications or refinements m_i one of which covers the intended content $m_{intended}$. This idea is expressed by assuming a general constraint ℂ defining the set of the possible pairs [sem(α),m]. In the simplest case sem(α) is a first-order formula, m a (partial) state description and ℂ(α) is realized as the set of (partial) state descriptions such that sem(α) holds in them ($\mathbb{C}(\alpha) =_{def} \{m: m \models sem(\alpha)\}$). A more refined conception of the general constraint ℂ can be defined in terms of abduction. In this case, ℂ(α) is realized as the set of abductive variants that can be generated from sem(α) by means of general world and discourse knowledge (see section 4).

The original formulation of the maxims seems to suggest that some concern primary what is said—so to speak, they concern sem(α) (e.g. the maxims under "Manner"), while others concern primary what is meant (e.g. "be relevant"). So it is certainly an appropriate picture to explicate the maxims as constraining [sem(α), m]-pairs. In order to capture notions like linguistic complexity and informativeness let us assume a global cost-function $\underline{c}(\alpha,m)$ defined as follows.

(1) $\underline{c}(\alpha,m) = compl(\alpha)\cdot c(sem(\alpha),m)$,
where compl(α) is a positive real number and $c(sem(\alpha),m) = -\log_2 pr(m|sem(\alpha))$

The global cost-function multiplies the complexity compl(α) of the linguistic aspects of α with a cost function c(sem(α),m) expressing the cost to correlate the linguistic meaning sem(α) with the (partial) state description m. We will assume that the latter costs are the negative logarithms of the (subjective) probabilities to associate m with sem(α). That means, the more probable the realization of a certain (partial) state description m—given the range of possible models provided by sem(α)—the less surprising this m should be and the less it should cost to assume.

The informal formulation of these principles shows a kind of partial circularity: in expressing the Q-principle regress to the I-principle has been made and vice versa. We think we must live with this kind of partial circularity, but we must give a precise formulation for it in order to see its consequences. The following is a step in this direction:

(2) (a) [α, m] satisfies the Q-principle iff there is no [α', m]∈ℂ satisfying the I-principle such that $\underline{c}(\alpha',m) < \underline{c}(\alpha,m)$
(b) [α, m] satisfies the I-principle iff there is no [α, m']∈ℂ satisfying the Q-principle such that $\underline{c}(\alpha,m') < \underline{c}(\alpha,m)$.

In this (rather symmetrical) formulation, the Q- and the I-principle constrain the set ℂ of possible [sem(α),m]-pairs in two different ways. The I-principle constraints the set by selecting the *minimal surprising* state descriptions with respect to a given semantic content sem(α) (provided Q has been satisfied), and the Q-principle constraints the set by blocking those state descriptions which can be grasped more economically by an alternative linguistic input α' (provided I has been satisfied).

Before we come to a more close inspection of the formulation (2) we have to introduce the notion of *common ground* and we have to investigate the effects of the *maxims of quality*. According to conventional wisdom, a common ground cg is an information state that contains all the propositions that are shared by the participants (for example, S and H). In more formal terms that means, an information state cg counts as common ground iff for each proposition ϕ holds (cf. Zeevat 1995):

(3) $cg \models \phi \Leftrightarrow cg \models B_S(\phi)\wedge B_H(\phi)$.

Here, B is a belief operator indexed to H and S, respectively. The belief logic we assume is Hintikka's (1962).

Let us write cg[α] for the common ground that results from cg by updating it with α. The crucial effects of the *maxims of quality* and their very special status within the overall theory can be formulated best in terms of conditions on updating the common ground. More concretely, let us claim that the first maxim of quality (*Do not say what you believe is false*) has the consequence that each possible state description m is consistent with the common ground cg[α]. Similarly, a consequence of the second maxim of quality (*Do not say what you lack evidence for*) is claimed to be that the informational content of the disjunction of the possible state descriptions m is contained in the common ground cg[α].

Let us use the abbreviation $\wp_Q(\alpha)$ for the set of possible state descriptions that are constrained both truth-conditionally (by means of $\mathbb{C}$) and by means of the Q-principle, i.e. $\wp_Q(\alpha) =_{def} \{m: [\alpha, m] \in \mathbb{C}$ and $[\alpha, m]$ satisfies the Q-principle$\}$. Analogously we have the definition $\wp_I(\alpha) =_{def} \{m: [\alpha, m] \in \mathbb{C}$ and $[\alpha, m]$ satisfies the I-principle$\}$. We simply write $\wp(\alpha)$ referring to the intersection of both conditions: $\wp(\alpha) =_{def} \wp_Q(\alpha) \cap \wp_I(\alpha)$. Using this notation we can give a partial formulation of the quality maxims:

(4) (a) Quality 1: for each $m \in \wp(\alpha)$: m is consistent with cg[α]
(b) Quality 2: $\wp(\alpha)$ is a non-empty set and $\bigvee\wp(\alpha)$ holds in cg[α]

Let us call an update *pragmatically licensed* iff it satisfies the conditions (4)(a,b). Now we call an utterance α *pragmatically anomalous* iff there is no pragmatically licensed update for it. Furthermore, a proposition ϕ is called a *conversational implicature* of α iff cg[α] $\models \phi$ for each pragmatically licensed update. If this relationship holds for each common ground cg we may speak of generalized implicature. Restricting the corresponding notions to specific classes of common grounds, we may define implicatures of the *particularized* variety.

3 Some Consequences of the Basic Mechanism

Let us now consider some simple examples to see the working of the mechanism. First consider Moore's paradox exemplified by the contrast between (5a) and (5b).

(5) (a) The cat is on the mat, but he doesn't believe it.
(b) ?The cat is on the mat, but I don't believe it.

The absurdity of (5b) falls out straight away as a case of pragmatic anomaly. The explanation immediately results from the formulation of the quality maxims in (4) and the conditions on common grounds (3). To see the crucial point, we have to show first that whenever α has a pragmatically licensed update cg[α], then the proposition $B_S(\text{sem}(\alpha))$ must be logically consistent. This assertion follows from the fact that the proposition $\bigvee\wp(\alpha)$ logically entails the proposition sem(α) and the fact that $\bigvee\wp(\alpha)$ is contained in the common ground cg[α]. The latter directly results from the condition (4b). As a consequence sem(α) must be contained in cg[α]. From our conditions on common grounds it follows that $B_S(\text{sem}(\alpha))$ must also be contained in cg[α] and therefore must be consistent.

Hintikka (1962) calls a proposition ϕ *doxastically indefensible* just in case the proposition $B_S\phi$ is inconsistent (with regard to his belief logic system). Using this notion, we can summarize our argumentation as follows: There can be no pragmatically licensed update for α in case the proposition sem(α) is doxastically indefensible. Utterances with doxastically indefensible sem(α) come out as *pragmatically anomalous* utterances according to the definition given above. It is a simple exercise to show that expressions of the form $\phi\wedge\neg B_S\phi$ are doxastically indefensible (i.e. $B_S(\phi\wedge\neg B_S\phi)$ is inconsistent). Consequently, the absurdity of (5b) comes out as a case of pragmatic anomaly.

Next, consider a simple example showing the generation of scalar and clausal implicatures. We consider the expression $\alpha = S_1$ *or* S_2 and the competing expression $\alpha' = S_1$ *and* S_2 and we assume that both expressions are of the same linguistic complexity: compl(α) = compl(α'). The derivation of the clausal and scalar implicatures of α is schematized in (6).

(6) (a) α: </S_1 or S_2/, $p\vee q$>
$\mathbb{C}(\alpha) = \{m_1, m_2, m_3\}$, where $m_1 = (p,q)$, $m_2 = (p,\neg q)$, $m_3 = (\neg p,q)\}$
(b) α': </S_1 and S_2/, $p\wedge q$>
$\mathbb{C}(\alpha') = \{m_1\}$
(c) $\wp_I(\alpha) = \{m_1, m_2, m_3\}$,
since $\underline{c}(\alpha,m_1) = \underline{c}(\alpha,m_2) = \underline{c}(\alpha,m_3)$
$\wp_Q(\alpha) = \{m_2, m_3\}$,
since $\underline{c}(\alpha',m_1) < \underline{c}(\alpha,m_1)$
$\wp(\alpha) = \wp_Q(\alpha)\cap\wp_I(\alpha) = \{m_2, m_3\}$
(d) if cg[α] is a pragmatically licensed update, then cg[α] $\models P_Sp, P_S\neg p, P_Sq, P_S\neg q$, ...(Quality 1: clausal implicatures)
cg[α] $\models B_S\neg(p\wedge q)$ (Quality 2: scalar implicature).
(P is the dual to the believe operator; $P_a\,\phi$ can be read as: *for all a believes, it is possible that* ϕ)

The derivation crucially rest on the assumption that the logically stronger expression S_1 *and* S_2 realizes the state description m_1 with higher probability than the logically weaker expression S_1 *or* S_2 and therefore can block this state description for the interpretation of S_1 *or* S_2.

It is worth noting that the present approach to "scalar implicatures" has some advantages over the traditional approach based on Horn-scales (cf. Gazdar 1979). In an exercise in his logic book McCawley (1993: 324) points out that the derivation of the exclusive interpretation by means of Horn-scales breaks down as soon as we consider disjunctions having more than two arguments. Consider the connectives AND and OR where both are construed as n-place operators, AND yielding truth when all n arguments are true and OR yielding truth when at least one argument is true. Clearly, as in the binary case we get for any number of arguments <AND, OR> as a Horn-scale which predicts that (7a) implicates (7b).

(7) (a) $OR(S_1, S_2, ..., S_n)$
(b) $NOT\ AND(S_1, S_2, ..., S_n)$

Unfortunately this prediction is too weak. The conjunction of (7a) and (7b) yields an formula which is true if any number of disjuncts smaller than n is true. This is correct for n = 2, but wrong for more arguments since a general account of the exclusive interpretation would have to predict the interpretation according to which it is true in case one (and only one) disjunct is. The utterance of (8) certainly does not invite you to take either one or two of the items mentioned.

(8) With the salmon you can have fries, rice or a baked potato.

It is easy to check that the current account yields the right result. As an example consider the case of three disjuncts $\alpha = OR(S_1, S_2, S_3)$. The derivation of the exclusive interpretation runs as above, but now based on the following alternatives $\alpha'_0 = AND(S_1, S_2, S_3)$, $\alpha'_1 = AND(S_1, S_2)$, $\alpha'_2 = AND(S_1, S_3)$, $\alpha'_3 = AND(S_2, S_3)$. Again the central point is that the stronger expressions realizes the relevant state descriptions with higher probability than the weaker expressions thereby blocking them for the interpretation of OR(S1, S2, S3).

It should be noted that we did not include the single disjuncts among the alternatives. This is motivated by the independent requirement (which any theory of Q-based implicatures has to make, but which is notoriously difficult to formalize) that the alternatives must contrast in view of an element which is qualitatively similar in a relevant sense. This is a general phenomenon. In spite of the entailment relation licensed by existential generalization a proper name as 'John' does not form a contrast class with a quantifier like 'some'. 'All' being a quantifier itself does.

The next class of examples deals with the phenomenon of (partial) lexical blocking. According to the Aronoff-Kiparsky tradition (e.g. Kiparsky 1982), 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*). Working independent of this tradition McCawley (1978) collects a number of further examples demonstrating the phenomenon of partial blocking outside the domain of derivational and inflectional processes. He observes for example that the distribution of productive causatives (in English, Japanese, German, and other languages) is restricted by the existence of a corresponding lexical causative. Whereas lexical causatives (e.g. (9a)) 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, (9b) could have been used appropriately when Black Bart caused the sheriff's gun to backfire by stuffing it with cotton).

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

Examples of this kind demonstrate a general pattern that Horn (1984) calls *the division of pragmatic labor*: "unmarked forms tend to be used for unmarked situations and marked forms for marked situations" (Horn 1984: 26).

Let us demonstrate now how the theory developed so far explains Horn's *division of pragmatic labor*. Consider two expressions α and α' that are semantically equivalent, i.e. $\mathbb{C}(\alpha) = \mathbb{C}(\alpha')$ and let us assume furthermore that $\mathbb{C}(\alpha)$ and $\mathbb{C}(\alpha')$, respectively, contain exactly two elements m_{dir} and m_{indir} of different complexity, say $c(\alpha,m_{dir}) = c(\alpha',m_{dir}) < c(\alpha,m_{indir}) = c(\alpha',m_{indir})$. In case that the expression α is linguistically less complex than the expression α', i.e. $compl(\alpha) < compl(\alpha')$, we can calculate the set $\wp_Q(\alpha)$ if it assumed that there is no expression α'' that expresses the content of α less costly than α itself. The application of (2a) simply yields $\wp_Q(\alpha) = \{m_{dir}, m_{indir}\}$. With this result at hand we can apply (2b) and get $\wp_I(\alpha) = \{m_{dir}\}$ (since $\underline{c}(\alpha,m_{dir}) < \underline{c}(\alpha,m_{indir})$). Consequently, we obtain $\wp(\alpha) = \{m_{dir}\}$, i.e., the unmarked form selects the unmarked situation.

Now consider the marked expression α'. In this case the application of (2a) yields $\wp_Q(\alpha') = \{m_{indir}\}$. This is due to the fact that a pair $[\alpha', m_i]$ can be blocked only by a less complex pair $[\alpha, m_i]$ if the latter satisfies the I-principle; thus, only $[\alpha', m_{dir}]$ is blocked but not $[\alpha', m_{indir}]$. Furthermore, it is easy to show that $\wp_I(\alpha') = \{m_{dir},$

m_{indir}}. Consequently, we obtain $\wp(\alpha') = \{m_{indir}\}$, i.e., the marked form selects the marked situation.

It is important to see that this explanation of the *division of pragmatic labor* doesn't rest on specific lexical stipulation or stipulation with regard to the costs, but is a general consequence of our formulation of the Q- and I-principle as presented in (2). According to earlier formulation (e.g. Atlas & Levinson 1981; Horn 1984), the Q- and I- (R-)based principles often directly collide and a general preference for the Q-principle has been stipulated. The present reformulation of the Q- and I-principle avoids this stipulation and predicts that in the "conflicting cases" the Q-principle yields a more restricting output than the I-principle.

4 Extending the Basic Mechanism

In the former sections, we have only considered a rather provisional explication of the constraint $\mathbb{C}$. In this section we consider a realization of $\mathbb{C}$ that seems refined enough to analyze some specific phenomena of lexical pragmatics. The main idea is to consider $\mathbb{C}(\alpha)$ as the set of abductive variants that can be generated from sem(α) by means of a specific common ground that includes crucial aspects of world and discourse knowledge.

In the following we consider sem(α) as a conjunction of positive literals and propose weighted abduction (Stickel 1989, Hobbs et al. 1993) as general method to specify sem(α) by exploiting Horn clause knowledge bases. The use of weighted abduction allows us to pair the abduced variants m_i with its proof costs. The earlier measure of the global costs $\underline{c}(\alpha, m_i)$ as given in (1) then should be replaced by an explicit account of those proof costs.

For the present purpose, we adopt Stickel's (1989) PROLOG-like inference system for generating abductive specifications and his mechanism for computing proof costs in a slightly simplified way. It is taken for granted that every literal in the initial formula is annotated with (non-negative) assumption costs c_i: $q_1^{c1}, \ldots, q_n^{cn}$. The knowledge base is assumed to provide formulas of the form $p_1^{\omega 1}, \ldots p_n^{\omega n}$ -> q, where the literals p_j in the antecedent are annotated with weights ω_j.

There are four inference rules that constitute abductive proofs and determine the assignment of concrete proof costs (for details, see Stickel 1989):

<u>Resolution with a fact</u>: If a current goal clause contains a literal that is unifiable with a fact in the knowledge base, then this literal is marked as proved. (The retention of a proved literal allows its use in future factoring).

<u>Resolution with a rule</u>: Let the current goal clause be ...$q'^{\,c}$ ... and let there be an axiom $p_1^{\omega 1}, \ldots p_n^{\omega n}$ -> q in the knowledge base. If q' and q are unifiable with most general unifier σ, then the goal clause ...$p_1^{c\omega 1}\sigma, \ldots, p_n^{c\omega n}\sigma$, $q'\sigma$... can be derived (where $q'\sigma$ is marked as proved). Obviously, we assume that the new assumption costs can be calculated by multiplying the corresponding weight factors with the assumption cost c of the literal q' in the old goal clause.

<u>Making an assumption</u>: Any unproved literal in a goal clause can be marked as assumed.

<u>Factoring</u>: If a literal q occurs repeatedly in a proof, each time with different costs, the occurrences of q are unified and the lowest cost is taken.

An abductive proof is complete when all literals are either proved or assumed. The cost measure for an abductive proof is the sum of all costs of the axioms involved in the proof plus the costs for the assumption of literals that are not proved. For the following, we will assume that all axiom costs are zero. Furthermore, we aim to bring our system as close as possible to a Bayesian network. As Charniak & Shimony (1990) have shown, this can be achieved when costs are interpreted as negative logarithms of certain conditional probabilities and when, besides other simplifying assumptions, no factoring occurs in the abductive proof. In the following, we will adopt the probabilistic interpretation of costs, but we will not refrain from using factoring. Factoring some literals obtained by backward chaining can be proven to be a very useful operation in natural language interpretation (cf. Stickel 1989).

It is now possible to incorporate the abductive component in the general pragmatic framework viewing natural language interpretation as inferences to pragmatically licensed updates. Let us for the sake of simplicity illustrate the incorporation of abduction by way of an elementary example. This gives us the opportunity to discuss some crucial differences between the present approach and the Hobbs-Stickel account where natural language interpretation is viewed as abductive inference to the *best explanation*. In order to simplify matters, we will exclude effects of blocking via the Q-principle. That means, we will assume that there are no expression alternatives α' that may block any interpretation of α.

Let us assume a knowledge base as presented in (10) and let us accept that all axiom costs are zero.

(10) **cg**:
$$\begin{array}{l} C^{2} \rightarrow S \\ D^{1} \wedge A^{0.5} \rightarrow C \\ D^{1} \wedge B^{0.5} \rightarrow A \\ \rightarrow D \end{array}$$

The diagram (11) shows the abductive inference graph in case (10) is taken as common ground and sem(α) = S is taken as the starting clause.

(11)

```
                                      S $10
                                      ▲
                                      |
                              C 2 -> S   [Assuming C:
                              ▲           2· $10=$20]
                              |
          D 1  ∧  A 0.5  -> C          [Assuming A:
          ▲       ▲                     0.5· $20=$10]
          |       |
          D       |
                  |
D 1 ∧ B 0.5 -> A                        [Assuming B:
▲                                        0.5· $10=$5]
|
D
```

The resulting set of abductive variants is presented in (12a) and the costs associated with these variants are given in (12b).

(12) (a) ℂ(α) = {A, B, C}
(b) c(S, A) = \$10, c(S, B) = \$5, c(S, C) = \$20
(c) ℘(S) = {B}
(d) cg[S] = cg ∪ {B}

Since we have assumed that there are no blocking alternatives, the condition (2a) becomes vacuous and the set ℘(S) is the set of cost-minimal variants, given in (12c). Since the expression B is consistent with cg, a pragmatically licensed update exists (satisfying the Quality conditions (4a,b)). It is given in (12d).

The Hobbs-Stickel account is looking for *minimal explanations*, that means it selects the cost-minimal variants from the set of the *consistent* abductive variants. This contrasts with the former view which first selects the cost-minimal variants from the set of all abductive variants and then checks them with regard to consistency. However, in the present case this makes no difference, since the minimal variant B is consistent with cg, and consequently it is at the same time the minimal explanation of S. The updating of cg by the minimal explanation gives the same result as already presented in (12d).

Now consider the common ground cg' given in (13), which is cg augmented by the clause `B ->` (i.e. by ¬B).

(13) **cg'**:

```
C 2 -> S
D 1 ∧ A 0.5 -> C
D 1 ∧ B 0.5 -> A
-> D
B ->
```

In this case we have the same abductive inference graph as shown before in (11), and we get the same abductive variants and the same costs associated with them. But now the cost-minimal variant B is inconsistent with cg'. From this fact it follows that there is no pragmatically licensed update for S with regard to cg'. In other words, S becomes pragmatically anomalous with regard to cg'.

(14)

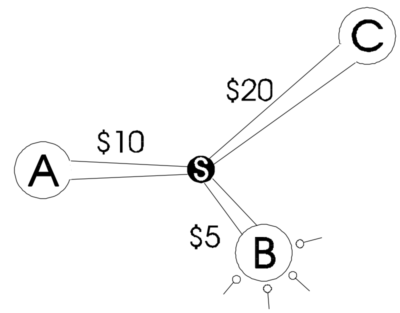

Now look at the Hobbs-Stickel account. It gives A as the minimal explanation (cf. the diagram (14)). This leads to the postulation of cg'∪{A} as update. Consequently, there is an important difference between the Hobbs-Stickel account and the present one. On the Hobbs-Stickel view there is an update in each case when the starting clause sem(α) is consistent with cg. The present account, on the other hand, yields a much more restricted notion of update. There is a pragmatically licensed update only in case one of the cost-minimal abduced variants is consistent with cg. If all cost-minimal variants are inconsistent with cg they can be seen as "blocking" any interpretation of the starting clause. As shown in the next section, this device is appropriate to capture cases of pragmatic anomalies in natural language interpretation.

From a computational point of view, the present approach looks well if it is assumed that the abductive machine generates the abductive variants in the order of its (estimated) costs. In this case, we have to assume simply that the abductive system stops if it has completed its first abductive proof. The result is then given to the consistency checker. In case the result is consistent, the system has found an interpretation. If not, the system may tell that it doesn't understand—the only interpretation it can find is a faulty one. Perhaps, there is a mechanism of accommodating the knowledge base that restores interpretability after all, but even then there is no possibility to access other variants than the cost-minimal ones.

The overall architecture of the Hobbs-Stickel account designates to access non-minimal variants when the minimal ones do not provide explanations. This feature makes processing less efficient, and it makes it difficult to discriminate between "good" and "bad" interpretations. In contrast, the present view of interpretation connects an efficient processing architecture with the possibility to give an explanation of pragmatic anomalies.

There is yet another important feature distinguishing the present account from the Hobbs-Stickel approach, the possibility of having non-cancelable implicatures. Let us call a conversational implicature ϕ of an utterance α in cg contextually cancelable iff there is a strengthening cg' of cg such that α is interpretable in cg' but ϕ is not longer a conversational implicature of α in cg'.

Obviously, the entailment $cg \cup \{sem(\alpha)\} \models \phi$ excludes the contextual cancelability of ϕ in cg. But what when we exclude simple entailments? Are conversational implicatures always cancelable? Is cancelability a necessary feature of conversational implicatures? According to the standard view the answer is yes.

With regard to our earlier example, B is a conversational implicature of S in cg and it is not entailed by $cg \cup \{S\}$. However, there is no (obvious) strengthening of cg that leaves the utterance α interpretable and defeats the proposition B. For example, if we strengthen cg by adding $\neg B$ (as in (13)) S will be pragmatically anomalous in the new context. This shows that on the current account cancelability is not a necessary feature of conversational implicature; some implicatures may be non-cancelable. The Hobbs-Stickel approach, on the other hand, is in agreement with the standard view (resting on the highly defeasible notion of minimal explanation).

The usefulness of cancellation as a test for conversational implicature has been attacked by Sadock (1978). Grice himself notes that cancelability doesn't hold for all kinds of conversational implicatures and mentions implicatures based on the Quality maxim as an exceptional case. Our discussion of Moore's paradox has demonstrated this in the context of the present theory. Langendoen's (1978) analysis of reciprocals shows other examples suggesting that cancelability is not necessary for conversational implicatures. Authors like Hirschberg (1991) on the other hand insist on taking cancelability a hallmark of conversational implicature. It seems, however, that this claim is not so much based on their treatment of conversational implicature but is rather a consequence the old dictum *semantics is strong and pragmatics is weak*. The present account to conversational implicature suggests that the borderline between semantics and pragmatics (conversational implicature) cannot be drawn by the condition of cancelability. The next section provides an analysis of the pragmatics of adjectives that may support this view.

5 The Pragmatics of Adjectives

One part of speech is especially suited for demonstrating the phenomenon of semantic underspecification: the adjective. It is well known that gradable adjectives as *large*, *short*, *quick*, and the like appear to take a fixed denotation only with respect to a certain class of objects. What is fast for a walker is not fast for a car and what is long for a truck is not long for a train. However, not only gradable adjectives but also adjectives that are commonly considered as *absolute* show a dependence upon the objects class. Quine (1960) for example 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 (1993) argues that an adjective 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, ... (Lahav 1993: 76).

Montague (1970) and others account for these facts by considering the attributive use of adjectives as fundamental and accordingly they do not treat adjectives as predicates, but rather as adnominal functors. Such functors turn the properties expressed by *train* into those expressed by *long train*. There are some notorious problems with this view. One concerns the analysis of the predicative use of adjectives. In this case an adjective must at least implicitly be supplemented by a noun. We moreover need various artificial assumptions (cf. Bierwisch 1989).

Another view of the semantics of adjectives has been called the *free variable view*. According to this view adjectives are represented by one-place predicate expressions. These expressions contain free variables which have the status of place holders for those aspects of interpretation which the grammar leaves unresolved. As an example, we can represent the adjective *long* (in its contrastive interpretation) 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 in a straightforward way:

The train is long translates (after λ-conversation) to 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 need not just be place holders for a comparison class as just indicated. The view can be generalized to include other types of free variables, for example a type of variable connected with the specification of the dimension of evaluation in cases of adjectives as *good* and *bad* or a type of variable connected with the determination of the object-dependent spatial dimensions in case of spatial adjectives as *wide* and *deep*.

The specification of free variables is necessary for a full interpretation of an utterance. We will demonstrate how the current theory yields an appropriate (contextual) specification by applying it to the kind of examples discussed by Quine (1960) and Lahav (1993).

(15) (a) The apple is red
(b) Its peel is red
(c) Its pulp is red
(d) APPLE(d)∧PART(d,x)∧COLOR(x,u)∧u=red
(e) APPLE(d)∧PART(d,x)∧PEEL(x)
∧COLOR(x,u)∧u=red

Our claim is that (15b) but not (15c) can be construed as a conversational implicature of (15a). Input of the analysis is the underspecified semantic representation given in (15d). One of the abductive specifications of this semantic input specifies x as the peel part of the apple (see (15e)). For the calculation of the corresponding costs we start with assumption costs as given in the first line of (16). Note that we take the assumption cost for the "slots" PART(d,x) and COLOR(x,u) as negligible with regard to the costs of the more "specific" elements of the representation. This contrasts with corresponding stipulations by Hobbs et al. (1993) but it agrees with the general picture that specificity is the primary determinator of the assumption costs. Furthermore, we refer to axioms of the form $q \leftarrow p_1^{\omega 1} \wedge p_2^{\omega 2}$, where the weights ω_1, ω_2 are monotonic functions of certain conditional probabilities: $\omega_i \propto prob(q|p_i)$ (cf., Hobbs et al. 1993). If the p_i are necessary conditions for q, then we have $\omega_1 + \omega_2 = 1$, and the weights ω_i can be interpreted to estimate the saliences of the features p_i with regard to p.

The diagram (16) shows that part of the abductive inference graph that is relevant for abducing the *red peel*-interpretation (15e) starting with (15d). The axiom in the second line of (16) can be seen as decomposing the concept of an apple into a peel part (salience α) and a residue, where the peel part is taken as a kind of slot-filler structure; γ may be interpreted as the salience of the part-relation for apples ($\gamma \ll 1$). In a similar vein, β may be interpreted as the salience of the color slot for the peels of apples.

Given the assumption that the color of the peel is more diagnostic for classifying apples than the color of other apple parts, for example, the color of the pulp, the *red peel*-specification is arguably the cost minimal specification. To make this point explicit, let us consider the calculation performed in (16). It crucially rest on the factoring operation which unifies the part- and color-slots of the predicate complex of the utterance with the corresponding slots that emerge while conceptually decomposing the subject term of the utterance. The *red peel*-specification comes out as the cost minimal specification if its total costs are smaller than the costs of any other specification. This corresponds to the condition $\alpha\beta > \alpha'\beta'$, where α' and β' are the parameters for any other apple part (e.g. for the pulp). Suppose that, as is rather plausible, this condition is satisfied, then the I-principle selects the *red peel*-interpretation and blocks the *red pulp*-interpretation. Consequently, we get (15b) as an conversational implicature, but not (15c).

(16)

$$\mathrm{APPLE(d)}^{const/2} \wedge \mathrm{PART(d,x)}^{\$0} \wedge \mathrm{COLOR(x,u)}^{\$0} \wedge \mathrm{u{=}red}^{const/2}$$

$$\mathrm{APPLE(d)} \leftarrow \mathrm{PART(d,x)}^{\alpha\gamma} \wedge \mathrm{PEEL(x)}^{\alpha(1-\gamma)} \wedge \mathrm{etc}^{1-\alpha}$$

$$\mathrm{PEEL(x)} \leftarrow \mathrm{COLOR(x,u)}^{\beta} \wedge \mathrm{etc}^{1-\beta}$$

R*ed Peel*-Variant

total costs: $const - \alpha\gamma - \alpha\beta(1 - \gamma) \approx const - \alpha\beta$

Note that the non-existence of the implicature (15c) doesn't *forbid* a discourse as (17) but rather *licenses* it.

(17) This apple is red. But not only its peel is red. Its pulp also is red.

In the case of (18a) analogous considerations give (18b) but not (18c) as a conversational implicature.

(18) (a) The apple is sweet
(b) Its pulp is sweet
(c) Its peel is sweet

It should be added that the present account evaluates utterances as (19) as pragmatically anomalous (assuming the former axioms and weights)

(19) ?This apple is red, but its peel is not (perhaps, its pulp is)

This qualifies implicatures like (15b) and (18b) as non-cancelable. It should be stressed that the fact that the present account doesn't postulate cancelability as necessary condition for conversational implicature is of remarkable importance for the sound treatment of such examples.

Finally consider the contrast between (20a) and (20b):

(20) (a) ?The tractor is pumped up
(b) The tires of the tractor are pumped up

The present account predicts (20a) as pragmatically anomalous. This prediction results from the fact that those parts of tractors that may be pumped on (the tires) are only marginally diagnostic for identifying tractors and therefore the corresponding interpretation can be blocked by specifications that refer to more salient parts. However, the latter specifications suffer from sort conflicts and therefore violate the condition (4). The details of this treatment can be found in Blutner (1996).

References

Atlas, J. and Levinson, S. 1981. *It*-Clefts, Informativeness and Logical Form. In P. Cole (ed.), *Radical Pragmatics*. New York: Academic Press.

Bierwisch, M. 1989. The Semantics of Gradation. In M. Bierwisch and E. Lang (eds.), *Dimensional Adjectives*. Berlin: Springer: 71-261.

Blutner, R. 1996. Lexical Pragmatics. Forthcoming.

Charniak, E., and Shimony, S.E. 1990. Probabilistic Semantics for Cost Based Abduction, Technical Report, CS-90-02, Dept. of Computer Science, Brown Univ.

Gazdar, G. 1979. *Pragmatics*. New York: Academic Press.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hintikka, J. 1962. *Knowledge and Belief. An Introduction to the Logic of the Two Notions*. Ithaca & New York: Cornell University Press.

Hirschberg, J. 1991. *A Theory of Scalar Implicature.* New York & London: Garland Publishing, INC.

Hobbs, J.R., Stickel, M.E., Appelt, D.E., and Martin, P. 1993. Interpretation as Abduction. *Artificial Intelligence* 63: 69-142.

Horn, L.R. 1984. Toward 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.

Kiparsky, P. 1982. Word-Formation and the Lexicon. In Proceedings of the 1982 Mid-America Linguistic Conference.

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.

Langendoen, D.T. 1978. The Logic of Reciprocity. *Linguistic Inquiry* 9: 177-197.

Levinson, S. 1987. Pragmatics and the Grammar of Anaphora. *Journal of Linguistics* 23: 379-434.

McCawley, J.D. 1978. Conversational Implicature and the Lexicon. In P. Cole (ed.), *Syntax and Semantics*, Vol. 9: *Pragmatics*. New York: Academic Press: 245-259.

McCawley, J.D. 1993. *Everything that Linguists have Always Wanted to Know about Logic but were Ashamed to Ask.* Chicago: The University of Chicago Press.

Montague, R. 1970. Universal Grammar. *Theoria* 36: 373-398.

Quine, W.V. 1960. *Word and Object*. Cambridge, Mass.: The MIT Press.

Sadock, J.M. 1978. On Testing for Conversational Implicature. In P. Cole (ed.), *Syntax and Semantics*, Vol. 9: *Pragmatics*. New York: Academic Press: 281-297.

Stickel, M.E. 1989. Rationale and Methods for Abductive Reasoning in Natural-Language Interpretation. In R. Studer (ed.), *Natural Language and Logic*. Berlin: Springer: 233-252.

Zeevat, H. 1995. The Common Ground. Forthcoming.