THE MECHANISMS OF DENIAL

BART GEURTS

University of Osnabrück and *Max Planck Institute for Psycholinguistics*

Normally, the negation of a sentence S serves to reverse the assertion that S would have made, and does not affect other types of information S would have conveyed, such as presuppositions, implicatures, and so on. Occasionally, however, it seems as if negation is directed precisely at presuppositions, implicatures, or even at purely formal aspects of a sentence (such as intonation, pronunciation, and so on). The following are cases in point: *Mary didn't visit the pizzeria in the Vatican, because there is no pizzeria in the Vatican./I'm not tipsy: I'm drunk./He didn't call the* POLICE*, he called the* POLICE.

I call such sentences DENIALS, and I argue against the unitarian approach advocated by Horn and van der Sandt, among others, according to which denial is a homogeneous phenomenon that calls for a unified analysis. According to the alternative theory proposed here, there are several mechanisms of denial, but each of these is needed for independent reasons, and therefore no ad hoc mechanisms are necessary.*

1. INTRODUCTION. By definition, presuppositions tend to be impervious to embedding under scope-bearing elements of whatever stripe, negation included. As is well known, however, under special circumstances negation may be directed at a presupposition.

(1) a. Kurt doesn't realize that his camels have been kidnapped.
b. Kurt DOESN'T realize that his camels have been kidnapped, because they HAVEN'T been kidnapped.

The factive verb *realize* triggers the presupposition that its complement is true, and in 1a this presupposition, which arises within the scope of a negation, is inherited by the matrix sentence: normally, 1a would be read as implying that Kurt's camels have been kidnapped. Normally, but not always: the same sentence occurring in the context of 1b does not imply in any way that Kurt's camels have been kidnapped.

It has been suggested that the difference between 1a and 1b may be accounted for by assuming that the former is semantically ambiguous: that it either suffers from an ambiguity of scope (although in this particular case it is perhaps not so clear what that might mean), or otherwise that negation is ambiguous between an operator that is transparent to presuppositions and one that blocks any presuppositions triggered in its scope. Such proposals have met with criticisms that nowadays are widely judged to be conclusive, and that I will not review here,[1] but even if it were true that 1a is semantically ambiguous between a presupposing and a nonpresupposing construal, we would still need to account for examples like 2 and 3.

(2) a. The room wasn't warm.
b. The room wasn't WARM: it was SWELTERING.
(3) a. The mice weren't killed.
b. The mice weren't KILLED: they were BUTCHERED.

By default, 2a will be construed as implying that the room was less than warm, but this inference must be overwritten in order to arrive at the intended interpretation of the same sentence as it occurs in 2b. The case of 3 is more puzzling still. If a speaker

* I should like to thank the editors and referees of *Language* for their very helpful comments on earlier versions of this paper.

[1] See, e.g., van der Sandt (1988) and Horn (1989) for references and discussion.

volunteers 3a, we would normally be entitled to the belief that the mice are still alive, but this doesn't hold for 3b. What is more, although it seems clear that the second half of 3b entails the negation of the first half, this sentence surely is not a contradiction in terms.

Prima facie, at least, the contrasts in the pairs in 2 and 3 are similar to the contrast in 1. In all these cases, the (a) sentences have a default reading which is overwritten in the corresponding (b) sentences, and in each case the (b) but not the (a) sentence would normally be used to correct an earlier statement or some contextually salient assumption. But even if postulating a semantic ambiguity would explain 1, it is rather doubtful that either some ambiguity of scope or lexical ambiguity of *not* will help to explain the data in 2 or 3.

Adopting a term used by Horn (1989), I will say that in the (a) examples in 1–3, negation is construed DESCRIPTIVELY. In the (b) examples, on the other hand, I will say that negation is used to make a DENIAL, adopting van der Sandt's (1991, 1998) terminology in this point (for Horn, these are 'metalinguistic' uses of negation).[2] Denials are marked vis-à-vis descriptive uses of negation: a denial is often escorted by a special (fall-rise) intonation contour, and generally requires further elucidation. In 2b, for example, it is only the second half of the utterance that explains why the speaker rejects 2a, and without this rectification the interpretation of this sentence would amount to guesswork (cf. Anscombre & Ducrot 1976, Horn 1989:402ff). Denials, thus understood, are the main theme of this paper.

Denials may be divided into the following four classes:

- proposition denials
- presupposition denials
- implicature denials
- form denials

For the time being this classification may be regarded as a purely expository device, because I do not want to claim that there is anything more than rough correspondence between it and the various empirical manifestations of denial. But I will claim later on that this taxonomy reflects the different mechanisms underlying denials.

I call a denial a PROPOSITION DENIAL if it serves to simply reject an earlier utterance, as in the following exchange:

(4) A: The cook is guilty.
B: The cook is not guilty.

Proposition denials will not play a prominent role in what follows, but right now the category is useful for establishing the following point. If it is agreed that 4 is in the same class as the (b) examples in 1–3, because in all these cases negation is used to make some sort of denial, it is evident that the division between descriptive negation and denial must be a fluid one. In 4, B uses a negative sentence for the purpose of correcting the previous statement, and in a different context he might have used the same sentence descriptively, for example, to announce a verdict he has arrived at. But between these paradigms of descriptive negation and denial there lie cases that are less clear-cut. For instance, if it is mutually known to the interlocutors that A believes that the cook is guilty, although A hasn't made this explicit, then B's *The cook is not guilty*

[2] Horn uses the term PREDICATE DENIAL in an entirely different way, viz. for a mode of predication (1989: 463ff). The label METALINGUISTIC NEGATION is originally due to Ducrot (1972). See Horn (1989:425ff) for a discussion of the differences between his own views and Ducrot's.

might still count as a denial. But if B isn't quite sure whether A holds this belief, then it is simply unclear whether or not B's utterance should count as a denial—unless one is prepared to cater to a notion of preemptive denial, which surely is not an appealing prospect. These observations demonstrate what on reflection should have been obvious from the beginning, namely that denial is a pragmatic category, and they illustrate, furthermore, that in practice the line between descriptive negation and denial may be hard to draw. This is because there are no deep differences between the two notions—or at any rate, that is one of the points that I shall be arguing for.

The category of presupposition denials is exemplified by 1b. Presupposition denials are comparatively easy to distinguish from others because there are quite reliable tests for diagnosing presuppositional expressions (see van der Sandt 1988 for extensive discussion). Proposition and presupposition denials are similar in that they often serve to downright reject a preceding utterance rather than amend it, as implicature and form denials typically do. The difference between proposition and presupposition denials is that the former are directed at the asserted content of the previous utterance, while the latter aim at its presuppositions.

I call 2b an IMPLICATURE DENIAL because I provisionally accept the Gricean (≠ Grice's) account of the standard interpretation of 2a, as developed by Horn (1972, 1989), Gazdar (1979), Levinson (1983), and others. In this article I use IMPLICATURE synonymously with CONVERSATIONAL IMPLICATURE, and I will mainly be concerned with a special class of conversational implicatures, that is, scalar ones. I take it, at least until §7, that the interpretation of the scalar predicate *warm* in 2a is the result of an interplay between the lexical meaning of *warm* and an implicature.[3] According to this view, the adjective's lexical meaning only imposes a lower bound on its interpretation, the upper bound being the result of a conversational implicature. Since *warm,* in effect, is taken to mean 'warm or more than warm', the default interpretation of the VP in 2a as 'less than warm' is immediately accounted for. However, in 2b the negation appears to function differently, namely to remove or block the upper-bounding implicature that would otherwise be associated with *The room wasn't warm.* Analogous observations apply to the following attested examples, reported by Horn (1989:382).

(5) a. Around here, we don't LIKE coffee, we LOVE it.
b. I'm not HAPPY he's gone—I'm ELATED.

Finally, I call a denial a FORM DENIAL if it is concerned with the form, style, or register of an utterance. I would place 3b in this category, but the following examples from Horn (1989:371) probably give a more vivid idea of what I have in mind.

(6) a. He didn't call the POlice, he called the POLICE.
b. I didn't manage to trap two monGEESE—I managed to trap two mon-GOOSES.
c. (Esker too ah coo-pay luh vee-and?)
Non, je n'ai pas 'coo-pay luh vee-and'—j'ai coupé la viande.
d. We didn't {have intercourse/make love}—we fucked.

Each of these examples is most likely to occur as a denial of an earlier utterance. Ex. 6a corrects the pronunciation of such an utterance, 6b the morphology, and 6c both. Finally, in 6d the truth conditions of the preceding utterance apparently are not at stake, either. What seems to be at stake, rather, is a certain choice of words that the speaker is evidently dissatisfied with. Prima facie such examples seem difficult to reconcile

[3] If the analysis that I propose in §7 is on the right track, the term IMPLICATURE DENIAL is actually misleading.

with the view that natural-language negation is unambiguous as well as truth functional—which is what I will be claiming, though.

To sum up the problem: it appears as if, under certain circumstances, negation can be directed against interpretative elements that normally would not be affected by it. This is the phenomenon that I want to account for in this paper. I begin by discussing two kindred proposals for dealing with denials: Horn's notion of metalinguistic negation; and a proposal put forward by van der Sandt, which by and large may be characterized as a more explicit version of Horn's analysis—with the proviso that van der Sandt doesn't require the assumption that negation is ambiguous between a descriptive and a metalinguistic (≈ denial) use. Apart from that, both theories are unitarian in the sense that they seek to account for all instances of denial in terms of a single key concept.[4] This unitarian approach leads both Horn and van der Sandt to regard denial as a metalinguistic device for objecting to full-blown utterances, real or imagined. According to this view, a denial applies to ALL the information conveyed by an utterance.

In contrast to this unitarian perspective, the theory that I advocate adopts a more eclectic stance. It repudiates the notion that denial is a phenomenon amenable to a unified treatment: there are several types of denial, whose underlying mechanisms are quite diverse. What these mechanisms have in common, however (and this is the second major difference with the unitarian approach), is that they are seldom aimed at complete utterances, although they may be directed at practically any part or aspect of a given utterance. In this sense, the analysis I propose is more local, or finer grained, than its unitarian competitor.

The range of facts that Horn and van der Sandt seek to cover is not as homogeneous as one might be tempted to think, as I will show. I propose that a distinction be made between implicature and form denials on the one hand and presupposition and proposition denials on the other. One argument in favor of this distinction is that in the implicature and form cases negation is typically used contrastively, while in the presupposition denials negation is generally noncontrastive, and in proposition denials it may be used either way. This difference is the surface reflex of a deeper distinction. Implicature and form denials affect the semantic interpretation of the expressions at which they are directed. In the context of 6a, for example, the noun *police* acquires a highly specific meaning, and the same holds for the occurrence of scalar *like* in 5a. The mechanisms at work in implicature and form denials are the same; they are introduced in §5. The differences between these two categories—the most important of which is that form denials are the only ones that are genuinely metalinguistic, in the sense that in the meaning of a form denial reference is made to a linguistic object—are discussed in §§6–7.

In §8, I turn to presupposition denials and argue that they are already accounted for by a presupposition theory proposed by van der Sandt (who doesn't employ it for this purpose, though). Formulated in classical terms, this theory assumes that regardless of the level of embedding at which they are triggered, presuppositions tend to break out of the scope of negation and other operators. However, under special circumstances a presupposition may be forced to remain in the position in which it was triggered, and I will argue that this is precisely what happens in presupposition denials. This line of explanation is quite different from the one I recommend for implicature and form

[4] In more recent work, Horn (1990, 1992) relaxes his strictly unitarian position in one point, but on the whole this is of relatively minor importance (cf. note 20). More whole-hearted defenses of a non-unitarian position are mounted by Kempson (1986) and Carston (1996).

denials, because in the case of presuppositions the semantic interpretation of the expression at which the denial is directed is not affected in any way: semantically speaking, there is no difference in interpretation between the two tokens of *realize* in 1a and 1b, for example.

2. METALINGUISTIC NEGATION. According to Horn 1985 and 1989, negation is not semantically but pragmatically ambiguous between a descriptive and what he calls a METALINGUISTIC use.[5] The latter corresponds to what I prefer to call DENIAL, save for the fact that Horn doesn't subsume proposition denials under his category of metalinguistic negation. When used descriptively, sentence negation represents a truth-functional operator. By contrast, in its metalinguistic use, negation is 'a device for objecting to a previous utterance on any grounds whatever, including the conventional or conversational implicata it potentially induces, its morphology, its style or register, or its phonetic realization' (1989:363). The following passage sums up Horn's position in rather explicit terms.

> Apparent sentence negation represents either a descriptive truth-functional operator, taking a proposition p into a proposition not-p (or a predicate P into a predicate not-P), or a metalinguistic operator which can be glossed 'I object to U', where U is crucially a linguistic utterance or utterance type rather than an abstract proposition. (Horn 1989:377)

Typically, the effect of metalinguistic negation is the rejection of an earlier utterance by another speaker, but it may also be wielded as a rhetorical device for correcting an utterance made by the same speaker or one that is somehow 'in the air'. Thus it is conceivable that example 7 (from Horn 1989:372) might be volunteered in the absence of an overt target in the preceding context.

(7) a. For a pessimist like you, the glass isn't half full—it's half empty.
b. I'm not a TrotskyITE, I'm a TrotskyIST.

In 7a, metalinguistic negation affects a scalar implicature, and in 7b it is directed against the use of the word *Trotskyite,* but the former might be felicitously uttered although nobody has claimed that the glass is half full, and an utterance of the latter doesn't require that someone has previously used the form in question.

Horn proposes three diagnostics for metalinguistic negation. First, metalinguistic negation cannot take the form of a prefix like *un-* or *non-,* as the following examples (Horn 1989:392) illustrate:

(8) a. The king of France is {*unhappy/not happy}—there isn't any king of France.
b. It {*is impossible/isn't possible} for you to leave now—it's necessary.

Since in 8a negation is incorporated in the adjective *unhappy,* it is unable to cancel the presupposition that there is a king of France. Similarly, the negation in *impossible* cannot suspend the scalar implicature that *possible* would normally give rise to. These observations are precisely what we would expect, given Horn's view on denial as a

[5] The sixth chapter of Horn 1989 is a revised and expanded version of Horn's 1985 article; accordingly, I will refer only to the former. A number of authors have wondered how Horn's notion of pragmatic ambiguity is to be understood (Foolen 1991, van der Sandt 1991, 1998, Carston 1996). Van der Sandt and Carston maintain, and I concur, that his own claims notwithstanding Horn's position entails that natural language negation is semantically ambiguous. In fact, as Carston points out, it involves a two-fold ambiguity: on Horn's view, negation is ambiguous between a descriptive and a metalinguistic interpretation, and furthermore, the material in the scope of a negative expression is construed differently depending on whether the negation is used descriptively or metalinguistically (Carston 1996:311).

metalinguistic device: a negative prefix is part of the word, and thus of the clause, in which it is incorporated, but metalinguistic negation 'operates, in effect, on another level from that of the rest of the clause in which it is superficially situated' (Horn 1989: 392), and it is plausible to assume that this is something an incorporated element cannot do.

The second diagnostic proposed by Horn hinges upon the distribution of polarity items. If metalinguistic negation doesn't operate within the clause but at a different level altogether, we shouldn't expect it to be able to either license negative polarity items or inhibit positive polarity items. This expectation is borne out by examples like the following (Horn 1989:370, 374, 396).

(9) a. Chris didn't manage to solve {*ANY/SOME} of the problems—he managed to solve ALL of them.
b. Chlamydia is not {*EVER/SOMETIMES} misdiagnosed, it is FREQUENTLY misdiagnosed.

Some and *sometimes* are positive polarity items which normally do not occur in the scope of negation, but as these examples illustrate, if negation is used metalinguistically, these items may occur in its scope, after all. Negative polarity items like *any* and *ever* normally occur in negative environments only,[6] but apparently metalinguistic negation does not create the kind of environment that admits such items.

Horn's third diagnostic for metalinguistic negation exploits the distribution of concessive and contrastive conjunctions with *but.* In some languages, the distinction between contrastive and concessive *but* is lexicalized (Pusch 1975, Anscombre & Ducrot 1977). Thus Spanish differentiates between *sino* and *pero,* and German between *sondern* and *aber,* where English only has *but.* Examples 10–12 (from Horn 1989:407f) illustrate this distinction.

(10) a. No es cierto, pero es probable.
b. Das ist nicht sicher, aber das ist wahrscheinlich.
c. It's not certain, but it's probable.
(11) a. No es probable, {*pero/sino} es cierto.
b. Das ist nicht wahrscheinlich, {*aber/sondern} das ist sicher.
c. It's not probable, but it's certain.
(12) a. Eso {*es inconsciente/no es consciente}, sino totalmente automatico.
b. Das ist {*unbewusst/nicht bewusst}, sondern ganz automatisch.
c. It's {*unconscious/not conscious} but (rather) totally automatic.

The contrast expressed by 10a–c is concessive: it is denied that 'it' is certain, but conceded that a weaker predicate, viz. *probable,* applies. In Spanish and German this concessive interpretation is signalled by a special form of conjunction, *pero* and *aber,* respectively. This analysis is corroborated by the examples in 11. While in 10 negation is used descriptively, in 11 its use is a metalinguistic one. If something is certain then, literally speaking, it is probable as well, and in these examples negation affects the upper-bounding scalar implicature which is associated with the weaker predicate. Accordingly, the contrast expressed in these examples is not a concessive one, and neither *pero* nor *aber* may be used. In their stead, the particles *sino* and *sondern* should have been used. It is these particles, therefore, that correlate with metalinguistic negation, as the contrasts in 12 confirm.

It would seem that Horn's dichotomy between descriptive and metalinguistic negation

[6] Actually, this is a considerable simplification. See Ladusaw 1996 for a recent overview.

correlates neatly with various empirical distinctions. That conclusion would be precipitate, however, because these diagnostics aren't actually reliable at all, either singly or collectively. The empirical distinctions they lay bare do not match with each other, and none of them coincides with the descriptive/metalinguistic dichotomy.

That the incorporation diagnostic does not test for metalinguistic negation per se can be seen from examples like 13.

(13) a. Not everybody was present.
b. Kurt is not very bright.
c. Gustav doesn't like Alma anymore.

For various reasons, none of these sentences have equivalents in which the negative element is morphologically incorporated, but we wouldn't want to conclude from this that they are instances of metalinguistic negation. On the other hand, there are cases that Horn would classify as metalinguistic in which the negation is incorporated, as in 14 for example.

(14) It is {impossible/not possible} that you have met the king of France, because there is no such person.

Horn's polarity test gives rise to similar problems. This can be shown, for example, by comparing Horn's examples in 9 with those in 15.

(15) a. Walter didn't give his ukulele to {?SOMEBODY/ANYBODY}: he never owned a ukulele.
b. Walter didn't {?SOMETIMES regret/regret at ANY time} that he betrayed his wife: he has always been faithful to her.

In these presupposition denials a negative polarity item may occur, whereas a positive polarity item produces an awkward effect. If Horn's polarity test were reliable, however, the converse should hold instead. It is clear enough that there are constraints on the distribution of polarity items and that there is some interaction between polarity and denial, but at present it is simply unclear what the nature of this interplay is. What is clear is that polarity is a poor diagnostic for denial.

Horn's third diagnostic has been discussed extensively by McCawley (1991), who argues convincingly that the notions of contrastive and metalinguistic negation are fundamentally distinct. Following Anscombre and Ducrot (1977), Horn suggests that 'the negation which (optionally) figures in the concessive . . . constructions is necessarily descriptive, while the negation required by the [contrastive] environments is typically understood as metalinguistic (1989:413). However, for one thing, contrastive environments are not always metalinguistic (as indeed Horn's guarded formulation indicates), as in 16.

(16) a. If Ramon hadn't been Spanish but French, he would still beat his donkey.
b. Wenn Ramon kein Spanier, *sondern* ein Franzose gewesen wäre, dann würde er noch immer seinen Esel schlagen.

In 16a negation is supposedly used descriptively, but nonetheless it occurs in a contrastive setting, as the German translation of this sentence (16b) demonstrates, in which *but* must be translated as contrastive *sondern* (rather than concessive *aber*). And as Pusch (1975) points out, there are many cases in which both *sondern* and *aber* are possible; these include scalar examples like 17.

(17) Die Arbeit ist nicht gut {sondern/aber} ausreichend.
The work is not good but satisfactory.

Many instances of metalinguistic negation, however, are not contrastive. This holds,

in particular, for the presupposition-denying uses, like 14. It appears therefore that, as the traditional terminology already implies, the distribution of concessive and contrastive *but*s distinguishes not between descriptive and metalinguistic negation, but rather between contrastive and noncontrastive environments and thus between contrastive and noncontrastive uses of negation. McCawley says that 'a correlation between contrastive and metalinguistic negation exists only because contrastive negation lends itself particularly easily to metalinguistic uses' (1991:189). This is surely correct, but it is important to note that contrastive negation tends to go together with certain varieties of metalinguistic negation only. Reverting for a moment to the classification I introduced earlier: contrastive negation very rarely occurs in presupposition denials, and is nearly always used in implicature and form denials.

There is an evident tension between Horn's claim that metalinguistic negation is directed at a complete utterance, on the one hand, and the observation that metalinguistic negation is typically focused on a selected portion of the utterance in question, on the other. Horn himself acknowledges that 'while serving to reject an entire utterance, metalinguistic negation focuses on one particular aspect of that utterance', and he offers the following examples (1989:434f):

(18) a. I'm advocating PROsecution, not PERsecution.
b. I called for a policeMAN, not a policeWOMAN.

There is not necessarily a problem here, but it is far from obvious that these two characteristics of metalinguistic negation are compatible. I will not pursue this point here because that would require the concept of metalinguistic negation to be more explicit than the one Horn is committed to. But I will return to this issue in the context of van der Sandt's similar but more explicit theory, where it does present a problem.

Horn characterizes metalinguistic negation as 'a device for objecting to a previous utterance on any grounds whatever' (1989:363), a description that seems particularly apt in view of the fact that even matters of register or style may be targeted with the help of metalinguistic negation. It is something of a surprise therefore, that certain aspects of the information content of an utterance cannot be objected to in this way. As Horn points out himself, relevance implicatures cannot be cancelled by means of metalinguistic negation. Consider examples 19 and 20 (discussed in Horn 1989:388f).

(19) a. He was able to solve the problem.
b. He wasn't able to solve the problem.
(20) a. She was clever enough to figure out the solution.
b. She wasn't clever enough to figure out the solution.

In many contexts, an utterance of 19a implies that, according to the speaker, the subject referent actually solved the problem, and similarly, an utterance of 20a may imply that the speaker believes that 'she' actually found the solution. These inferences, which Horn takes to be relevance implicatures, cannot be cancelled by way of metalinguistic negation, as the (b) sentences show: 19b cannot be construed as implying that 'he' was able to solve the problem but didn't solve it, and 20b cannot be read as implying that 'she' was sufficiently clever to figure out the solution but didn't bother to do so, or was prevented from doing so.

The same point can be made with the help of the following example.

(21) A: The door is open.
B: The door is not open—you can close it yourself if you wish.

A's utterance in 21 might be taken as an indirect request for B to close the door. But

even in a context in which it is apparent that this implication is intended, it would be strange for B to retort as in this dialogue.

Horn (1989:387ff) observes that there seems to be a difference in the way quantity implicatures and relevance implicatures interact with metalinguistic negation. As can be seen from ex. 22 (=5), quantity implicatures can be cancelled by metalinguistic negation.

(22) a. Around here, we don't LIKE coffee, we LOVE it.
b. I'm not HAPPY he's gone—I'm ELATED.

Examples 19–21 indicate that relevance implicatures, by contrast, cannot be cancelled by negation in any way. Horn suggests that the explanation for this contrast lies in the fact that relevance and quantity implicatures have different 'logics' (1989:388f). In both examples in 22, the speaker would be contradicting himself if the negation was interpreted descriptively. For instance, supposing that, semantically speaking, *happy* merely imposes a lower bound on its argument's degree of happiness, the descriptive interpretation of *not happy* must be 'less than happy'. But on this interpretation 22b would be a contradiction in terms, and therefore the negation 'self-destructs' and 'must in effect be sent back through, whence the marked, metalinguistic quality of the negation' in 22b and 22a (Horn 1989:388). In 19b and 20b, however, the negation of the relevance implicature is perfectly consistent with a descriptive interpretation of the negation. If the person in question didn't solve the problem, then 19b might be one of the possible reasons, and the same holds, mutatis mutandis, for 20b. Thus, the negations in 19b and 20b 'never spontaneously self-destruct, to get sent back through to be interpreted metalinguistically as cancelers of implicata' (Horn 1989:389).

This solution is unsatisfactory for two reasons at least. First, if Horn's reasoning is correct, it should be possible to cancel a relevance implicature by means of negation, provided care is taken that a descriptive reading is ruled out, as in 23.

(23) a. He wasn't able to solve the problem—he could have solved it but he didn't.
b. She wasn't clever enough to figure out the solution—she could have done it but she didn't.

In both 23a and 23b there is a contradiction if the negation is read descriptively, and if Horn's argument were sound, one should expect the negation 'to get sent back through to be interpreted metalinguistically' in both cases. But this doesn't seem to happen: both discourses are simply awkward.

Secondly, one of the premises of Horn's argument is false. Horn assumes that the problem is to explain why quantity implicatures can but relevance implicatures cannot be cancelled by means of negation. However, the following example shows that not even all quantity implicatures can be cancelled in this way.

(24) A: Jones may have gone to Panama.
[Implicature: A isn't sure where Jones went.]
B: Jones may not have gone to Panama: I'm convinced you actually believe that he went to Costa Rica.

In 24 speaker A implicates that he doesn't know where Jones went—more accurately, that the information he takes to be true doesn't decide this issue. Speaker B denies A's statement, making it clear that he is objecting to A's quantity implicature. But obviously B should have employed other means for expressing this.

Metalinguistic negation is so called because it applies not to propositions but to utterances or utterance types. Hence it would seem that metalinguistic negation has the

same status as the logicians' turnstyle: '$\vdash (\varphi \vee \neg \varphi)$' is not an expression *of* but *about,* say, classical logic; it is an expression of the metalanguage. Therefore, Horn's characterization of metalinguistic negation leads one to expect that it cannot occur in embedded positions, an expectation that doesn't materialize, as 25 shows.

(25) a. In America, people don't eat to[mɑ:]toes but to[meɪ]toes.
b. Barney claims that Mary didn't eat two but three bananas.
c. Mary is convinced that the king of France isn't bald because there is no king of France.

In 25a metalinguistic negation is used within a locative frame, and in 25b and 25c it occurs within the scope of an attitude verb.[7] Yet intuitively speaking there are no relevant differences between these examples and the ones previously seen. The same point can be made with the help of the pattern in 26, discussed in Anscombre & Ducrot 1977.

(26) If Walter hadn't had three children but four . . .

Apparently, metalinguistic negation is just as much part of the language as is ordinary descriptive negation. But this conclusion is inconsistent with Horn's notion of metalinguistic negation as an 'operator which can be glossed 'I object to U', where U is crucially a linguistic utterance or utterance type rather than an abstract proposition' (Horn 1989:377).

3. NEGATING ECHOES. Van der Sandt's (1991, 1998) notion of denial is closely related to Horn's metalinguistic negation, although there are some nontrivial differences, the most important difference being that van der Sandt does not accept that negation is ambiguous in any way. While he agrees with Horn that there is a functional difference between (what Horn would call) descriptive and metalinguistic negation, van der Sandt's position is that negation always has the same meaning, semantically as well as pragmatically. Regardless of whether *not* φ is used descriptively or to correct a preceding utterance, *not* has the same content in either case. It is its operand, i.e. φ, which according to van der Sandt may be determined in two different ways. Normally speaking, in order to work out the meaning of 27a in a context c, we first determine the meaning that 27b would have expressed in c, and apply the negation operator to the resulting proposition:

(27) a. The king of France is not bald.
b. The king of France is bald.

Since the presupposition that there is a king of France is not part of the propositional content of 27b, it is not affected by the negation operator, and thus 27a normally expresses that there is a king of France who is not bald. However, if 27a occurs as a denial, as in the following discourse,

(28) A: The king of France is bald.
B: The king of France is not bald: there is no king of France.

the operand of the negation (27b) does not have its ordinary interpretation. Instead, it serves to take up, or echo, the previous utterance, and in so doing it picks up ALL the information conveyed by this utterance, including the presupposition that France has a king. Thus the first negation in B's contribution is directed, inter alia, against the presupposition of A's utterance because its argument picks up all the information con-

[7] The relevant reading of 25c is the one on which the attitude verb takes scope over the whole conjunction.

veyed by that utterance, presupposition and all. In precisely the same way, implicatures and formal properties of preceding utterances may enter into the scope of a denial.

Van der Sandt's conception of denial can be made more precise as follows.[8] We model a discourse as a sequence of contexts $c_0, \ldots, c_n$. In each context c_i the common ground between speaker and hearer is given by GIVEN(c_i), which may be construed as a set of possible worlds, for example. Suppose that in context c_i a sentence φ is uttered. Then PROP(φ, c_i), PRES(φ, c_i), and IMP(φ, c_i) denote the propositional content, the presuppositions, and implicatures of φ in c_i, respectively, while FORM(φ, c_i) contains a miscellany of information conveyed by φ in c_i, for example, that (the speaker assumes that) the words in φ are pronounced in certain ways, that they belong to a register that is appropriate to the occasion, and so on. The information that φ contributes to c_i may be defined as in 29a: CONT(φ, c_i) is the sum of the four sets of information associated with φ in c_i MINUS the information that was already given in c_i before φ was uttered.

(29) a. CONT(φ, c_i) = (PROP(φ, c_i) + PRES(φ, c_i) + IMP(φ, c_i) + FORM(φ, c_i)) − GIVEN(c_i)

b. PROP(*φ, c_i) = CONT(φ, $_{i-1}$), provided φ was uttered in c_{i-1} by another speaker; otherwise PROP(*φ, c_i) is undefined

Ex. 29b gives the semantics of the 'echo operator': the propositional content of *φ at c_i is the sum of all information conveyed by φ at c_{i-1}. Van der Sandt assumes that *φ is undefined if φ didn't occur in the previous context or was uttered by the same speaker. In this respect his theory is not only more explicit but also more restrictive than Horn's, and some of the trouble that van der Sandt's analysis runs into is caused by this restriction.[9]

The difference between ordinary (descriptive) negation and denial can now be accounted for if we assume that in the latter case but not in the former the echo operator is involved: while ordinary negated sentences are of the form $\neg\ \varphi$, denials are of the form $\neg\ *\varphi$. Thus van der Sandt's notion of denial almost perfectly fits Horn's description of metalinguistic negation as 'a device for objecting to a previous utterance on any grounds whatever' (Horn 1989:363). The fit is not quite perfect, though, because van der Sandt doesn't have to assume that negation is ambiguous in any sense, although denials are metalinguistic in that they echo the preceding utterance. But this is arguably a virtue of his analysis, because it allows van der Sandt to account for the examples in 25 and 26, which are problematic for Horn, as we have seen. Since van der Sandt's analysis requires only ordinary descriptive negation, these examples present no problems to his theory.

In van der Sandt's theory only one type of negation operator is called for, which is truth functional. Horn, presumably, would object to this on the grounds that a truth-functional operator could never apply, for example, to phonetic information.[10] But there really is no deep mystery about this. Consider 30.

[8] This quasi-formalization is a simplified version of van der Sandt's own proposal.

[9] The similarity between van der Sandt's notion of ECHOING and Carston's (1996) ECHOIC USE is mainly verbal: van der Sandt's echo operator picks up everything that is conveyed by the preceding utterance; Carston's notion of echoic use is much more restrictive.

[10] For instance, referring to Carston and van der Sandt, Horn objects that 'any such attempt . . . to propositionalize not only upper-bounding implicata but the stylistic, connotative, and mechanical aspects of utterances that fall within the scope of marked negation . . . would seem to be self-defeating, representing a kind of category mistake: an 'echoic use' is not the sort of beast to which a truth-functional operator applies.' (Horn 1992:172, cf. Horn 1989:434) This objection is echoed by Foolen (1991:228) and criticized by Carston 1996: 327ff).

(30) A: Your wife reminds me of Frankenst[i^{y}]n.
B: She doesn't remind you of Frankenst[î^{y}]n: she reminds you of Frankenst[á^{y}]n.

The propositional content of A's utterance is that B's wife reminds him of Frankenstein, and his pronunciation of the monster's name doesn't enter into this. But by uttering his line in this particular way, A commits himself to the truth of at least one further proposition, to wit that the name is pronounced 'Frankenst[î^{y}]n'. Therefore, this proposition is part of the information conveyed by A's utterance, and because B echoes this utterance it lands within the scope of an ordinary, truth-functional, negation operator. Therefore, what is actually negated is the proposition that the name is pronounced 'Frankenst[î^{y}]n', and of course this IMPLIES that B criticizes A's pronunciation. If this is what is meant if one says that negation may 'apply to' the pronunciation of an expression (and I can see no other way of making sense of such claims), then examples like 30 can very well be analyzed (e.g. along the lines suggested by van der Sandt) without sacrificing the assumption that natural-language negation is univocally truth functional.

In the previous section I indicated that there is a tension between Horn's description of metalinguistic negation, on the one hand, and the fact that this type of negation typically focuses on part of the utterance under denial, on the other. In van der Sandt's theory this tension becomes more acute. According to van der Sandt, a denial rejects ALL the information that was conveyed by the preceding utterance.[11] Intuitively, however, although this may sometimes be the case, in general denials leave at least part of the target utterance intact. This observation is illustrated by many of the examples cited above, two of which are repeated here for convenience.

(31) a. I'm not a TrotskyITE, I'm a TrotskyIST. (= 7b)
b. I'm advocating PROsecution, not PERsecution. (= 18a)

In these examples, the speaker objects to an aspect of the preceding utterance without rejecting it wholesale. This observation is in line with the fact that, within a denial, negation may be restricted in its scope, as in 32.

(32) Some of my best friends aren't TrotskyITES but TrotskyISTS.

Here, negation is most plausibly construed as being within the scope of the subject term, which is not what one would expect if it were directed against a complete utterance. Furthermore, denials often carry presuppositions of their own. Thus in 31b the speaker presupposes that he is advocating something; his actual statement concentrates on what he is advocating. Similarly, if someone utters 33, he presupposes that there is a Canadian empress.

(33) The king of France didn't have dinner with the Canadian empress: France doesn't have a king.

Van der Sandt predicts, however, that the negation in 33 takes scope over both the presupposition that there is a king of France and the presupposition that Canada has an empress.

In the same vein, consider the following examples:

[11] This view is tied up with the unitarian outlook and the position on implicature denials that Horn and van der Sandt share. The argument (which neither author makes explicit) goes as follows: conversational implicatures are licensed by complete utterances only, not by any of their parts; so implicatures can be denied only by quoting or echoing complete utterances; but since denial is a uniform phenomenon, it must always work this way. See §7 for further discussion of this point.

(34) A: All professors were intoxicated.
B: Some of them weren't intoxicated but drunk.
(35) A: So you caught two little fish?
B: I didn't catch two little fish: they were sturdy ten pounders!

In 34, the denial contains a pronoun that picks up its reference from the utterance that is being objected to, while in 35 a pronoun refers back to a discourse entity that is introduced within the scope of a denial. But if van der Sandt is right in claiming that a denial always serves to reject a complete utterance, one should not expect such phenomena to be possible.

These observations are irreconcilable not only with the letter but with the very spirit of van der Sandt's analysis. It is crucial to his concept of denial that it apply to (the information conveyed by) a complete utterance, because one of van der Sandt's (and Horn's) premises is that a denial may affect conversational implicatures and presuppositions. And since it is utterances, not parts of utterances, that carry implicatures and presuppositions, van der Sandt is forced into the assumption that the target of a denial is always a full-blown utterance. It is precisely this assumption which causes trouble in the examples above. Consider 33, for example. Intuitively, we would like to say that the definite NP *the Canadian empress* performs its ordinary presuppositional function in this case: the speaker merely wants to question the existence of a French monarch. However, according to van der Sandt, this NP is part of an echo; the first part of 33 is of the form $\neg$ *(the king of France had dinner with the Canadian empress), and it is hard to see how an echo, properly so called, could have presuppositions of its own. At this point van der Sandt might want to say that the *-operator doesn't literally echo anything, and that presuppositions may be triggered within its scope. Out of the frying pan, into the fire: for then he would have to explain why 33 doesn't presuppose that there is a king of France, although it contains a definite NP that triggers this presupposition. And if the contexts created by the *-operator are transparent to presuppositions, shouldn't we expect them to give rise to conversational implicatures, as well? Etcetera. If such questions could be answered, a viable theory might emerge, but it would be rather different from the one proposed by van der Sandt.

Finally, van der Sandt's assumption that denials always echo explicit utterances is surely too strong (cf. Carston 1996:312). Many of the examples discussed above might very well occur in the absence of an overt target. For example, in a talk on the pronunciation of English, the following might naturally occur noncorrectively:

(36) Until the end of the 18th century, Englishmen didn't [dɑ:ns] but [dæns].

Following a suggestion offered by Anscombre and Ducrot (1977:27), van der Sandt might want to say that in such cases, a denial aims at a virtual utterance ('affirmation virtuelle'). But this solution gives rise to a panoply of puzzling questions: Do virtual utterances have virtual presuppositions and virtual implicatures? May the speaker of a virtual utterance be himself a virtual person, or must he be one of the interlocutors? Must the set of virtual utterances be consistent? And so on. In brief, the concept of virtual utterance is a murky one, at best. Furthermore, it is intuitively not very plausible that 36 involves an echo of a virtual utterance of, say,

(37) Until the end of the 18th century, Englishmen [dɑ:nst].

But, something like this does appear to be called for if we want to rescue the notion that all denials are echoic.

4. VARIETIES OF DENIAL. The theories proposed by Horn and van der Sandt are unitarian in the sense that they seek to explain denial in terms of a single key concept.

In Horn's account, this concept is metalinguistic negation, in van der Sandt's it is echoing. Thus, both theories are committed to the claim that we are dealing with a collection of phenomena subsumed by a single unifying principle. I believe that this claim is wrong. In my view, there is no natural class of denials or instances of metalinguistic negation, and no single principle to account for the phenomena that Horn and van der Sandt are concerned with. I would argue instead that their data are held together only by the somewhat peripheral circumstance that they involve negative expressions that are often used correctively—in particular, to correct an earlier utterance. But this tendency is not indicative of a single underlying principle, at whatever level of analysis.

I will try to show that, although there is no such unifying principle, denials can always be accounted for in a principled way. That is to say, all instances of denial may be explained without special-purpose theoretical machinery. In particular, it is not necessary to assume that negation is ambiguous in any sense. Semantically as well as pragmatically, negation always functions descriptively; in this respect, my position is closer to van der Sandt's than to Horn's, but on the whole is far removed from both.

In §1 I proposed a four-fold classification of denials:

(38) PROPOSITION DENIALS
A: It's raining.
B: No, it is not raining.

(39) PRESUPPOSITION DENIALS
A: Fred has stopped smoking.
B: Fred hasn't stopped smoking: he never did smoke.

(40) IMPLICATURE DENIALS
A: Julius had six beers.
B: He didn't have six beers: he had at least seven.

(41) FORM DENIALS
A: Kurt swallowed a whole to[mɑ:]to.
B: He didn't swallow a to[mɑ:]to but a to[meɪ]to.

This classification primarily reflects my conviction that different mechanisms underlie these categories of denial. To be sure, there are significant correlations between this division, on the one hand, and surface phenomena, on the other, but as far as I can see these correlations are imperfect. I have argued for example that implicature and form denials typically involve a contrastive use of negation, whereas presupposition denials are typically noncontrastive and proposition denials may be used either way, though these rules are by no means hard and fast. For example, presupposition denials may be used contrastively, as in

(42) Barney didn't take his WIFE to Acapulco—he isn't even married—but his GIRLfriend.

This type of example is somewhat special, presumably because in most cases an utterance which is thought to suffer from presupposition failure will be rejected downright, and not be amended, as in 42. Nonetheless, I maintain that the interpretative processes underlying this sentence are the same as in other instances of presupposition denial.

There is a further reason why the classification I propose is reflected but imperfectly in the observable data. It is that denials can sometimes be understood in more than one way, without these alternatives corresponding with clearly distinct readings.

(43) A: So your fiancée finally ditched you?
B: No, she did not ditch me. As a matter of fact, we'll be married by the end of the month.

Here, B's response may be construed as a straightforward proposition denial but also as a form denial, in which case B would be taken as echoing part of A's utterance. There may be differences between these two construals, but they are extremely subtle; they certainly make no difference to the truth conditions of the sentence, and in practice it may be impossible to tell which reading is intended. But as I understand it such observations do not show that the classification I propose is incorrect.

In the remainder of this article I will present detailed analyses of form, implicature, and presupposition denials (in that order). I will not discuss proposition denials, because I consider this class to be unproblematic: if speaker A says φ, speaker B may correct him by tokening the negation of φ. As far as I can see, there is nothing mysterious about this. Compared with the other varieties of denial, it seems to me, proposition denials do not present any problems of interest.

5. SEMANTIC TRANSFER AND POLYSEMY. In §§6 and 7, I shall present analyses of implicature and form denials that crucially depend upon the assumption that the contextual interpretation of an expression need not coincide with any of its lexical meanings (if it is a single word) or with any of its compositional meanings (if it is a nonlexical expression). I refer to this phenomenon as SEMANTIC TRANSFER. I don't think that it is controversial to claim that semantic transfer is a fact of life, but I think its importance is widely underrated and that it is too often dismissed as a fringe phenomenon. To demonstrate its importance, I will first outline the essential facts about semantic transfer and the related notion of polysemy.

A word is standardly considered to be polysemous if it has several senses that are conceptually related to one another, and this is how I will use the term POLYSEMY, too. I speak of semantic transfer if, on a given occasion, a word acquires a contextual meaning that is not encoded in its lexical entry, although typically it will be related to one of its lexical meanings. Thus understood, semantic transfer and polysemy are allied concepts, which cannot always be kept apart. If a given instance of semantic transfer recurs often enough, it is likely to result in polysemy. This is how words grow new meanings: first a word comes to be used in a certain novel way by semantic transfer, then this nonce meaning is lexicalized alongside the original meaning of the word, and eventually the original meaning may even be ousted by the derivative. Thus, semantic transfer can be the beginning of polysemy. It may sometimes be difficult, if not impossible, to determine if the contextual meaning of a word was selected from its lexical entry or modulated by semantic transfer.

If one consults a dictionary for the meaning of an ordinary word, it is quite likely that the search will yield more than one sense. More often than not, these senses will be related, which is to say that, according to most dictionaries, many words (and, presumably, most of the more frequent ones) are polysemous in that they have several related senses. Although lexicographers have been criticized and even derided for their allegedly haphazard collections of word meanings, I don't think one would want to claim that they have exaggerated the EXTENT of the phenomenon. On the contrary, in the literature on lexical semantics it has become something of a truism to observe that polysemy is the rule rather than the exception, especially among words in everyday use.[12]

Although it is not a lawlike phenomenon, polysemy is systematic to a considerable

[12] See, for example Nunberg 1978, 1979, Bierwisch 1983, Herskovits 1986, Lakoff 1987, Sweetser 1990, Wierzbicka 1994, Copestake & Briscoe 1995.

degree. There are numerous patterns of polysemy that recur both within and across languages. Some instances are shown in 44; see Nunberg 1978:55ff for more examples.

(44) INDIVIDUAL: He smashed the light bulb.
GENERIC: The light bulb was invented by Edison.
COUNT: George caught a fish.
MASS: I suspect that there's fish in this salad.
CONTAINER: She engraved the glass with her name.
CONTENTS: She drank two glasses.

These sense pairs have at least two features in common. First, in each of these cases there is a fairly definite intuition as to which sense is the basic or primary one: it is the first member of every pair. There doesn't seem to be one single explanation for these judgments. I can think of reasons why the count sense of *fish* should be prior to its mass interpretation, but none of these explains why the individual sense of *light bulb* is more basic than its generic sense. Nonetheless, it is worth mentioning that we have such intuitions at all, because they probably determine our judgments on what is the ordinary or literal sense of a word. Secondly, in each pair there is what I will call a salient functional link from the primary to the secondary sense—given the world knowledge that is generally available among speakers of the language, the primary sense offers a useful pointer towards the nonprimary sense (see Nunberg 1978, 1979). For instance, there is a correspondence between the zoological and the consumer concepts of fish. This correspondence is a functional one in the mathematical sense: there is a function that maps the former concept to the latter. This doesn't take us far because functions are cheap, and for any given purpose there are generally more functions than one needs. But there are strict constraints on the functions that can map between the senses of a word. In particular, such functions must be salient in the sense that they are part of our everyday knowledge, and may therefore be taken to be part of the common ground between speaker and hearer. Although there are many functions that map the concept of fish onto the concept of giraffe, none of these will be a salient function in the sense indicated, and therefore it is not to be expected that the noun *fish* (or its counterpart in any other language) will be polysemous between 'fish' and 'giraffe'.[13] The concept of salient function could be explored in greater detail but for my purposes here it is sufficient if something like this notion is involved in the kind of polysemy exemplified by the pairs in 44. (See, especially, Nunberg 1978, 1979, 1995 for further discussion.)[14]

In a commonsense theory of linguistics the relation between the lexical entry of a

[13] There are forms of polysemy that seem to work quite differently. For example, the link between the two senses of an expression may be analogical. Thus, in many languages perception verbs have derived senses for designating various mental activities, as in 'I see your point' or 'Do you hear me?' Such analogical or metaphorical transfers have been studied by Lakoff (1987) and Sweetser (1990); they will not concern us here.

[14] Although polysemy is often considered to be a lexical phenomenon, it is not restricted to single words, as the following examples show.

(i) I suspect that there is some freshly caught fish in this salad.
(ii) She drank two glasses that her husband had filled for her.

In (i) it is presumably not just the noun *fish* that acquires a mass interpretation, because the adjectival phrase *freshly caught* requires that its head have a count meaning. It seems likely, therefore, that it is the phrase *freshly caught fish,* rather than any of its parts, that is construed in the culinary sense. Parallel remarks apply to the object NP of (ii). So polysemy is not just a lexical phenomenon, although I will concentrate on lexical instances of polysemy.

word and its use is deceptively straightforward: someone who needs to understand a word as it is being used in a given context will have to look it up, either in his own mental lexicon or in a dictionary, and should the word turn out to have more than one meaning, then the meaning appropriate to the occasion must be selected. If this view were right, the relation between the lexical entry of a word and its uses would be one of selection: each time a word is used it is sufficient to select one of its lexical meanings. This view surely isn't right in all cases, nor even acceptable as an idealized picture (see Nunberg 1978, Bosch 1984, 1985). In support of this assessment, I could cite many of the data collected by Nunberg (1978), the most famous of which probably is

(45) The ham sandwich is waiting for his check.

Nunberg observes that this may make sense, for example, when uttered by one waiter addressing another, and in order to arrive at this construal the interpretation of at least some element of 45 must be contextually modulated. Nunberg himself suggests that the referent of *the ham sandwich* is used to pick up another referent, the man who ordered the ham sandwich. However, as Sag (1981:285) observes, this will not work for examples like 46.

(46) a. There are five ham sandwiches sitting at table 9.
b. Ham sandwiches generally prefer to sit by the window.
c. We haven't seen any ham sandwiches for more than a week now.

These data indicate that what is required is a transfer in the meaning of *ham sandwich,* rather than just a shift in reference. I will assume that Sag is right about this and that 45 is to be analyzed along the same lines. Thus, in certain contexts at least, 45 is interpreted as follows:

(47) The ⟨person who ordered {the/a} ham sandwich⟩ is waiting for his check.

Here and in the following I use angle brackets to hint at a possible contextual interpretation of a given expression, in this case *ham sandwich.* (This notation is not to suggest that the expression in question is an abbreviation of sorts.)[15] Readers who feel that 45 is an aberrant example, which can safely be dismissed as a marginal case, might want to take a look at the collection compiled by Nunberg (1979), which contains plenty of similar cases; here are a few of Nunberg's examples.

(48) a. He can hit the ball two *football fields.*
b. He made a pile in *radio.*
c. He hit a home run two *games* ago.
d. *IBM* went up three points last week.
e. I love some kinds of liver; *chicken* is tasty.
f. *The Times* asked if he could have an exclusive interview.

In each of these examples, the italicized expression must undergo semantic transfer in order to mesh with the rest of the sentence.

6. 'QUOTATION' AND FORM DENIALS. According to the unitarian theories proposed by Horn and van der Sandt, all denials are metalinguistic in that they crucially refer to other utterances. In my view this is only partly right because, first, denials are not always metalinguistic in nature, and second, to the extent that an element of quotation is involved, neither Horn's nor van der Sandt's account is satisfactory, because they

[15] For proposals as to how such semantic transfers might be incorporated into a classical model-theoretic semantics, see Sag 1981 and Bartsch 1987.

both misrepresent the workings of quotation. In order to explain this and clear the ground for a better treatment, I turn to the subject of quotation.

The distinction between using and quoting an expression might appear to be a straightforward one: 'The name of a name or other expression is commonly formed by putting the named expression in single quotation marks; the whole, called a *quotation,* denotes its interior' (Quine 1951:23f). As Quine defines it, quotation is a perfectly simple matter, but Quine is not describing what speakers actually do; rather, he is prescribing what they should do, and he holds a decidedly negative view on the everyday practice of quotation. 'Confusion of sign and object is original sin, coeval with the word' (Quine 1969:15).

I want to take a closer look at some natural language phenomena that are best regarded as instances of quotation, by which I mean that in some way reference to linguistic objects is involved. It will become apparent that these phenomena are generally not in conformity with Quine's recommendations, because in many cases the referent of a quotation cannot be equated with the expression that is being quoted. As a matter of fact, it is not at all unusual that the referent of a quotation cannot plausibly be construed as a linguistic object at all, although reference to a linguistic object is involved, and I will argue that quotation phenomena cannot be accounted for unless we allow for the possibility that the meaning of a quotation is determined by, among other things, the kind of contextual influences discussed in the previous section.

One obvious respect in which linguistic usage deviates from Quine's prescriptions is that there are no generally accepted formal devices for distinguishing between use and mention. There are typographical conventions for marking quotations, but they are not applied systematically. In speech, of course, there mostly is no difference between use and mention, although occasionally intonation is called upon to assist in clarifying the speaker's intentions. All things considered, language users don't seriously try to make formal distinctions between linguistic expressions and their names. Instead they simply use an expression to refer to itself, whenever the need arises. It follows from this that ALL expressions of any language are equivocal in a way: besides their ordinary meaning(s), they can also be used to designate themselves. I am not sure whether this is an instance of polysemy or not, but I believe we find polysemy even if we consider the quotation sense on its own.

It is easily shown that the meaning of a quotation depends on the context in which it occurs. Consider:

(49) a. 'Romance' consists of two syllables.
b. 'Romance' is both a substantive and a verb.

If it is understood that we are speaking of English, 49a and 49b are both true. But if it is the author's intention to quote the Dutch word *romance,* neither is correct, since Dutch *romance* is unequivocally a trisyllabic substantive. And fixing the language doesn't always suffice to pin down the meaning of *romance* either, as 50 demonstrates.

(50) a. 'Romance' contains three vowels and four consonants.
b. 'Romance' contains two vowels and four consonants.

Both 50a and 50b may be true: it depends on whether one is interested in letters or sounds. This difference can be accounted for if we assume that the quotation of *romance* may express different meanings on different occasions. Then 50a might be equivalent to something like 51a, and 50b to 51b, for example.

(51) a. ⟨The string 'romance'⟩ contains three vowels and four consonants.

b. ⟨The Received Pronunciation of 'romance'⟩ contains two vowels and four consonants.

Another type of quotation, which is extremely common in texts about foreign languages and cultures, is exemplified by 52.

(52) a. The organization of society was by families and clans with a sharp distinction between *eorls,* a kind of hereditary aristocracy, and the *ceorls* or simple freemen. (Baugh & Cable 1978:49)

Apparently, this is a specimen of the confusion between use and mention that Quine complains about. On the one hand, *eorl* and *ceorl* are not part of the vocabulary of Modern English, and Baugh and Cable clearly intend to convey to their readership the meanings of these Old English words—i.e. they want to say something ABOUT these words. But on the other hand, these words are at the same time USED to say something about their referents (in Old English). Thus described, it would seem that 52a is a hopeless muddle. However, we can explain how this sentence nonetheless manages to make sense if we assume that, in the context of 52a, the words *eorls* and *ceorls* acquire the following nonce meanings:

(52) b. The organization of society was by families and clans with a sharp distinction between ⟨a class of people called in OE 'eorls'⟩, a kind of hereditary aristocracy, and the ⟨class of people called in OE 'ceorls'⟩ or simple freemen.

Here the contextual meanings of *eorls* and *ceorls* refer to linguistic expressions as well as to social classes. But the latter are their referents, properly so called: the sentence says something about two classes of people, referring to OE expressions by the way. Example 53a may be construed along the same lines as 52a.

(53) a. The rooms of [Timofey Pnin's] Waindell period looked especially trim in comparison with one he had had in uptown New York, midway between Tsentral Park and Reeverside. (Nabokov, *Pnin*)

b. The rooms of [Timofey Pnin's] Waindell period looked especially trim in comparison with one he had had in uptown New York, midway between ⟨what Pnin referred to as 'Tsentral Park'⟩ and ⟨what Pnin referred to as 'Reeverside'⟩.

Apparently, it is the author's intention to slip in samples of Pnin's pronunciation while conveying information on the location of Pnin's abode in New York. Thus 53a is strictly parallel to 52a: even if we count names as part of the language, *Tsentral Park* and *Reeverside* are not English expressions, but nonetheless the author can use these expressions to refer to Central Park and Riverside, respectively. Also, because he makes it understood that these expressions are to echo Pnin's speech, the author is not committed to the assumption that *Tsentral Park* and *Reeverside* are formally correct ways of identifying their referents: if a speaker uttered 53a, it would be inappropriate to correct the speaker's pronunciation of these words (except on the grounds that it fails to capture Pnin's accent).

Each of these cases is metalinguistic in the sense that reference is made to a linguistic object. But the last couple of examples aren't PURELY metalinguistic; rather, they are about linguistic and extralinguistic entities at the same time. An element of quotation is involved in each case, but it is only part of our intuitive understanding of these examples. There may be more than one way of capturing this blend of use and mention, but an account in terms of semantic transfer suggests itself rather strongly.

If we compare examples like 52a and 53a with Horn's paradigm instances of metalin-

guistic negation, the similarity is quite striking. In particular, there doesn't seem to be any relevant difference at all between form denials and the examples I have discussed in this section.

(54) a. The formulator of relativity theory wasn't Einst[ɪ̂ʸ]n but Einst[á͡ʸ]n. (Horn 1989:379f).

b. The formulator of relativity theory wasn't ⟨a man called 'Einst[ɪ̂ʸ]n'⟩ but ⟨a man called 'Einst[á͡ʸ]n'⟩.

(55) a. He didn't call the POLICE, he called the POLICE. (= 6a)

b. He didn't call the ⟨official body whose name is pronounced 'POLICE'⟩, he called the ⟨official body whose name is pronounced 'POLICE'⟩.

Sentence 54a would typically serve to correct someone's pronunciation of the name *Einstein,* and this may be accounted for if we assume that, in the context of 54a, 'Einst[ɪ̂ʸ]n' is construed as 'a man called ''Einst[ɪ̂ʸ]n'' ' and 'Einst[á͡ʸ]n' as 'a man called ''Einst[á͡ʸ]n'' '. There may be alternative ways of construing these forms, but the only thing that matters for now is that it is possible to make perfectly good sense of 54a, even adopting the standard construal of negation, if we accept that the contextual meanings of 'Einst[ɪ̂ʸ]n' and 'Einst[á͡ʸ]n' are the outcome of semantic transfer.[16]

My analysis of 54a carries over, mutatis mutandis, to 55a and to all cases in which negation is deployed for the purpose of objecting to the realization, style, or register of a given expression—i.e. the prototypical instances of metalinguistic negation. (Other varieties of denial call for different analyses.) In my view, there are no pertinent differences between these instances of denial and specimens of quotation like 52 or 53. Note in particular that an example like 55a refers to linguistic and nonlinguistic objects at the same time: a speaker who volunteers this sentence explicitly claims that the referent of *he* called the police, and in doing so comments upon the pronunciation of the word *police.* In this sense, 55a isn't purely metalinguistic, although its interpretation does involve a metalinguistic component—which is precisely what we found in 52a and 53a, too.[17]

The treatment I advocate for 54 and 55 differs in various ways from the approaches taken by Horn and van der Sandt. First, unlike Horn, I don't have to assume that negation is ambiguous; the negative morphemes in 54 and 55 carry their ordinary truth-functional meaning. In this respect, my proposal is in line with van der Sandt's. Second, I don't require the assumption that the unit of denial is the sentence, as both Horn and van der Sandt must assume, because they start out from the premise that in denials only complete utterances can be quoted. I have shown how to account for the intuitively obvious fact that smaller units can be quoted as well, and therefore this proposal doesn't suffer from the drawbacks that are inherent to the coarser-grained unitarian analyses. For one thing, my proposal is consistent with the observation that a denial typically focuses on a selected portion of the offending utterance: if a denial is concerned with

[16] Note that in order to determine the contextual meaning of 'Einst[ɪ̂ʸ]n', the hearer may require the cue provided by 'Einst[á͡ʸ]n'. The same holds for 'POLICE' and 'POLICE' in 55a. In many cases, the intended meaning of a denial simply cannot be determined without the speaker's subsequent rectification or clarification.

[17] If I am not mistaken, here lies an important difference between my analysis and Carston's. According to Carston 1996, a denial is always echoic in that it 'reports what someone else has said or thought and expresses an attitude to it' (Carston 1996:320; cf. Sperber & Wilson 1986:238f, Wilson & Sperber 1988: 145ff). Apart from the fact that I do not agree with Carston that echoic use is an essential ingredient of denial, my analysis of form denials diverges from hers in that on my view the expression the negation operator is aimed at in form denials is intended to refer to linguistic objects as well as to objects in the world, as illustrated by 54 and 55.

a single mispronounced word, for example, the speaker will only have to quote this word, and it is only natural that he then should focus on it too. By contrast, if we adopt the unitarian stance, this simple fact may be quite difficult to account for, because this position entails that metalinguistic negations or denials must be directed at complete utterances (cf. §§2–3). Also, and for the same reason, my proposal is consistent with the fact that denials may occur in embedded positions, as illustrated by 56 (= 36), a fact which causes trouble especially for Horn's approach, as I argued in §2.

(56) Until the end of the 18th century, Englishmen didn't [dɑ:ns] but [dæns].

This example also illustrates the last difference between my theory and van der Sandt's that I want to mention here. There is nothing in my analysis that leads us to expect that form denials must be directed at a preceding utterance. As I observed in §3, 56 might quite naturally occur, for example, in a textbook on comparative or historical linguistics—an observation that is irreconcilable with van der Sandt's approach (and presumably with Horn's as well). But according to my analysis, the fact that the contextual meaning of an expression is in some sense metalinguistic does not entail that it must pick up an expression used in the previous discourse. This is illustrated already by the examples 52 and 53, and confirmed by 56. It may be that form denials are TYPICALLY directed against explicit utterances, but that will not be an essential feature of this variety of denial. Rather, it must be explained by the circumstance that uttering a sentence like 54a, for instance, happens to be useful only (or mainly) in contexts in which someone has just referred to Einstein with 'Einst[ɪ̈ˀ]n' (see also Carston 1996).

7. IMPLICATURE DENIALS. I turn now to what I have provisionally called IMPLICATURE DENIALS, which are exemplified by the following:

(57) a. Fred didn't buy five sheep but six.
b. She isn't intelligent: she is brilliant.

This type of denial plays a crucial role in Horn's and van der Sandt's unitarian theories, because they take as their starting point the standard Gricean account of scalar expressions, according to which the interpretation of a scalar expression is lower bounded by its lexical meaning and, in most contexts at least, upper bounded by a conversational implicature (see Horn 1972, 1989, Gazdar 1979, Levinson 1983). Since, on the one hand, conversational implicatures are licensed by utterances, not by sentences or their parts, while, on the other hand, in the examples above negation appears to cancel an implicated upper bound, one of the main problems with denials, if we accept the Gricean analysis, is to explain how a conversational implicature lands within the scope of a negation operator. Horn and van der Sandt seek to solve this problem by assuming that the upper bounding inferences cancelled in 57a and b do not originate as conversational implicatures in utterances of THESE sentences, but in preceding utterances (or, perhaps, utterances which are just 'in the air'), which these utterances quote, implicatures and all.

I am not convinced that the Gricean theory of scalar expressions is true, nor am I convinced that it is false. For present purposes, however, it is irrelevant whether it will prove to be true or false, and if it is true it is consistent with an analysis of denials that does not require the assumptions that denials are always metalinguistic and that they only apply to complete utterances.

I will defend the position that, at least in such contexts in which the sentences in 57 would normally be uttered, the scalar expressions *five* and *intelligent* are construed as having an 'exactly' sense; their contextual meanings are 'exactly five' and 'intelligent but not brilliant', respectively. A scalar expression might conceivably acquire its 'ex-

actly' reading in either of two ways. This reading, if it emerges, might be the outcome of semantic transfer, or it might be that it is one of the lexical meanings of a scalar expression. For reasons discussed in §5, I will not take a firm stance in this matter, but I want to establish that, in the context of a denial, scalar expressions may have an 'exactly' meaning. I will tentatively suggest, however, that for some scalars at least, this is also the primary (and therefore, presumably, lexicalized) meaning.

According to the Gricean analysis of scalar expressions, the literal or conventional meaning of 58a is that the cat ate at least five chocolates.

(58) a. The cat ate five chocolates.
b. The cat ate exactly five chocolates.

In most contexts however 58a would be taken to implicate that the cat didn't eat more than five chocolates. Here I am not so much interested in the train of reasoning that secures this result as in the form of its conclusion, which is that 58b is entailed by the literal meaning of 58a in conjunction with a quantity implicature derived from an utterance of this sentence. Thus the argument stops short of what I take to be its natural conclusion, that the hearer takes one further step, and infers that the speaker intended to use the adjective *five* meaning 'exactly five'. That is to say, we may accept the Gricean argument as far as it goes, but assume in addition that it results in a semantic transfer of the adjective *five.*

There are a number of considerations that militate in favor of this suggestion. First, it is intuitively plausible. If it turns out that the speaker actually intended to convey no more than the information that the cat ate at least five chocolates, we do say things like: 'When you said five, I thought you meant exactly five'. This holds not only for scalar implicatures, but for practically all generalized conversational implicatures. Grice introduces the notion of generalized conversational implicature as follows:

> Sometimes one can say that the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature. (Grice 1989:37)

Here Grice suggests that conversational implicatures may be associated with PARTS of utterances. Strictly speaking, this cannot be right, because if conversational implicatures are 'carried' at all, then they will have to be carried by UTTERANCES rather than by their parts. However, this passage testifies to a natural tendency to attribute such implicatures to certain key expressions. Apparently, whenever an expression α is escorted by a conversational implicature φ, and it is felt that α plays a key role in the derivation of φ, then it is natural to pretend as if φ had actually become part of α's contextual meaning. So natural, in fact, that in the more informal parts of their communications, virtually all implicature theorists have spoken in these terms, at one time or another.

It may be argued, however, that the attribution of implicatures to subsentential units cannot be merely a manner of speaking, even if we accept the Gricean framework without reservations, because such attributions are a prerequisite for the conventionalization of implicatures. It is widely accepted that conversational implicatures may become conventionalized.[18] A textbook example of this is the English conjunction *since,* which started out having only a temporal meaning, but developed an additional causal sense (e.g., Hopper & Traugott 1993:74). This development may be described as the

[18] The possibility of conventionalization was already alluded to by Grice in the first of his William James lectures (Grice 1989:39). See also Morgan 1978, Levinson 1983:165f, Horn 1989:343ff, Hopper & Traugott 1993:75ff.

conventionalization of what started out as a conversational implicature, and it must have proceeded via an intermediate stage during which the habit developed to use *since* in a causal sense, although this option was not yet incorporated into its lexical entry.

I have not argued so far that scalar expressions have or may have multiple lexical meanings. I have merely tried to establish that even if we adopt the Gricean stance, it will have to be granted that scalar expressions may be construed, in the contexts in which they occur, as having an 'exactly' reading, which would be the outcome of a conversational implicature. This is not enough to account for implicature denials, however. I want to claim that the scalar expressions in 57 are construed as having 'exactly' readings, and these readings could not be the result of conversational implicatures. Therefore, I will have to assume that a scalar expression may have an 'exactly' reading even in a context in which there is no conversational implicature to secure the upper-bounding inference. But this assumption is not at all implausible. Suppose that the lexical entry of a scalar expression only lists an 'at least' sense. Then one of the related concepts of this lexical meaning surely will be the corresponding 'exactly' sense. As noted in §5, there is a salient function from the 'at least' to the 'exactly' sense, and therefore it is only to be expected that semantic transfer from the former to the latter may occur.[19]

Thus far I have assumed, in accordance with the Gricean view, that all scalar expressions have the same type of lexical meaning, an 'at least' sense. However, once it is accepted that the contextual meaning of scalar expressions (as of any other type of expression) may be mediated by semantic transfer, it is no longer necessary to stick to this assumption. We could also consider the possibility that the 'exactly' sense is lexicalized, and that the 'at least' sense, whenever it occurs, is the result of semantic transfer. After all, the 'at least' and the 'exactly' sense of each scalar expression uniquely determine EACH OTHER, so if we have either one, it is straightforward to derive the other. I want to suggest, however, that in this regard the category of scalar expressions is not uniform, and that in some scalars the 'exactly' sense is primary, while in others it is the 'at least' sense which is basic. Into the first category I would place cardinal expressions and perhaps determiners like *a few,* while adjectives like *warm, bad,* and *intelligent* would go into the second category.

Cardinal expressions, in particular, behave in a way that sets them apart from other scalar expressions.[20] As Sadock (1984) points out, the distinction is most vivid in arithmetical statements like 'Two times three is five', which should come out true on the strength of its truth-conditional meaning alone, if the lexical meaning of all cardinal expressions *n* were 'at least n'. But even if we ignore such cases, it may be argued that the combinatorics of cardinals is captured more adequately on the assumption that the primary meaning of all cardinals *n* is 'exactly n' (see Koenig 1991). Consider, for example, the expression *more than five sheep.* What does *five* mean in this context? If it

[19] In certain respects, the analysis of scalars and implicature denials presented is in line with the views taken by relevance theorists such as Kempson (1986), Carston (1988, 1996), and Foolen (1991). The central assumption that I share with these authors is that the meaning of scalar expressions may be affected by pragmatic factors, so that they may acquire a genuine 'exactly' meaning. What distinguishes my theory from Kempson's and Carston's is that I propose to view the process that determines the meaning of a scalar in terms of polysemy and/or semantic transfer. This is what allows me to claim that in some scalars the 'exactly' sense is primary while others default to the 'at least' sense.

[20] Cf. Sadock 1984, Horn 1989:250f, and especially Horn 1992, which concedes that cardinals differ from other scalar expressions, and suggests that cardinals should be exempted from the standard neo-Gricean treatment.

really meant 'five or more', then *more than* would have to have a rather counterintuitive semantics (something like a function that takes a scale σ into a scale σ' that differs from σ only in that the lower bound of σ has been cut off). Similar remarks apply to expressions like *less than, at least, approximately,* and so on. Moreover, none of these premodifiers combines comfortably with gradable adjectives like *warm: more than warm* and *at least warm* are rather odd, save perhaps in special contexts, and *approximately warm* doesn't make sense at all. Such observations are better accounted for if we may assume that the primary meaning of a cardinal expression (as opposed to that of other scalars, such as *warm*) is the 'exactly' one.

The position that I have outlined so far is that the lexical meaning of a scalar is either an 'at least' or an 'exactly' sense, while the other sense is the result of semantic transfer. Therefore, in a given context, a scalar expression may have either interpretation, which brings us to the question of how the hearer manages to select the intended interpretation. Several factors presumably play a role here, the most important of which is that the hearer will try to extract a useful message from the speaker's contribution. As a consequence, the negated predicates in the implicature denials in 57 will get an 'exactly' interpretation; on an 'at least' interpretation, each of these would be inconsistent. A second factor particularly relevant to denials is that there are various linguistic means for steering the hearer towards one interpretation or the other. One class of modifiers, for example, requires its scalar arguments to have an 'exactly' interpretation. More relevant here is the fact that there are more or less conventional means of enforcing contrastive interpretations, which in the case of implicature denials require 'exactly' interpretations. The lexicons of some languages (see §2) distinguish between concessive and contrastive conjunctions, a distinction which in English is indicated by intonation and/or syntactic form (see Horn 1989:402ff). I have argued that the distinction between concessive and contrastive conjunction is not, as Horn has suggested, an adequate test for denial. But there can be no doubt that there are linguistic means for signalling contrast, and in certain contexts these may serve to direct the hearer towards a denial interpretation. For reasons discussed in §4, this will generally hold for implicature as well as for form denials, but not for presupposition denials.

This analysis of implicature denials is quite similar to the analysis of form denials that I proposed in the previous section. Both accounts hinge on the notion of semantic transfer (which may or may not have resulted in polysemy); the main difference is that form denials always involve a metalinguistic element that may be lacking in the scalar cases: implicature denials do not (or not necessarily) require that reference be made to linguistic objects in any way. In view of this similarity, my claims about the advantages of my theory over its unitarian competitors (see §6) apply here as well: my analysis is more local, more in line with observations about focusing, has no problems with embedded denials, and so on. I offer two further pieces of evidence. Horn, in his discussion of scalar implicatures (1989:240f), points out that, in some cases at least, scalar predication is crucially dependent upon contextual factors. To illustrate this, he cites the following attested examples (1989:241).

(59) a. Overt *antifeminism,* if not *homosexuality,* may be the result of such experience in the male.
b. In the Netherlands the crowds [for the Pope] were small, the welcome *lukewarm* if not *cold.*
c. Most photographers were *inarticulate* if not *subhuman.*

Surely, *homosexuality* does not imply *antifeminism* by virtue of its lexical meaning,

but in the context of 59a, these two words have to be interpreted in such a way that the latter is semantically weaker than the former. An analogous interpretation also applies to the italicized expressions in 59b and 59c. These observations do not fit comfortably with the Gricean line on scalar predicates that Horn advocates, because this is a form of context dependence for which there is no natural niche in Grice's doctrine. In contrast, the theory of scalars proposed here has no trouble at all with this type of phenomenon, because it takes the context dependence of meaning as its starting point.

I argued above that one of the most serious problems of the unitarian position is that it leads us to expect that denials should be possible in many cases in which they are clearly inappropriate. Examples 19 and 21 are repeated here as 60 and 61.

(60) a. He was able to solve the problem.
b. He wasn't able to solve the problem.

(61) A: The door is open.
B: The door is not open—you can close it yourself if you wish.

Example 60b shows that even in a context in which its affirmative counterpart 60a would imply that the subject solved the problem, this sentence cannot be used to convey the information that 'he' was able to solve the problem but didn't do so. Similarly, even in a context in which it is perfectly clear that A intends to make an indirect request, B cannot reject this request as he attempts to in 61. I have argued that there is nothing in the unitarian theory to explain why this should be so, but on the account that I defend the explanation is obvious. My analysis of implicature denials involving scalar expressions turns on the assumption that, in the context of a denial, a scalar is construed as having an 'exactly' interpretation. Compare this with exx. 60 and 61. It is intuitively evident that if 60a is used to convey that 'he' actually solved the problem, this has no consequences for the contextual interpretations of, say, *able* or *was able to solve the problem.* The latter is not then construed as 'was able to and in fact did solve the problem'. Put otherwise, the inference that the subject of 60a solved the problem is not incorporated into the contextual interpretation of any of its parts, and therefore cannot get under the scope of the negation, as 60b confirms. Similarly for 61, where A's utterance might convey that A wants B to close the door, but this inference does not affect the semantic interpretation of any part of the sentence uttered, and therefore it cannot be negated. This is a further argument in favor of an analysis that applies more locally than the unitarian one.

The standard argument in favor of the Gricean analysis of scalar expressions (and of other natural language phenomena) is that it is more parsimonious than the alternative account, which assumes that scalar expressions are ambiguous between an 'at least' and an 'exactly' reading. This is a dubious argument. Apart from the fact that it suggests that ambiguist analyses present the only possible alternative to the Gricean doctrine, it is doubtful that considerations of lexical economy should be given so much weight. Implicature theorists have always been wont to emphasize that senses should not be multiplied beyond necessity; this is what Grice has called Modified Occam's Razor. Although I have no objections to this slogan, I don't think that Occam's name should be associated with it. Lexical meanings are not on a par with scientific laws or natural kinds. If they exist in any way at all, they are individuals, and a scientific theory does not become more complex if we double or triple the number of individuals in its ontology. Senses should not be multiplied beyond necessity, but postulating two senses instead of one should not be compared, if only tacitly, with, say, proliferating levels

of syntactic description (see Lewis 1973:87 on the difference between qualitative and quantitative parsimony).

These remarks are beside the point, however, because I am not committed to the assumption that scalar expressions are ambiguous. I do hold that different tokens of a scalar expression may express different meanings, but that is not the same as saying that it has two lexical meanings. So the choice between the unitarian theory and my theory cannot be decided by comparing the sizes of their respective lexicons. Nor should either one be accused of being dependent on ad hoc premises. Straightforward methodological considerations, it seems, will not force a decision. The matter will have to be decided by comparing the explanatory potential of the two theories; as I have argued, such a comparison speaks against the unitarian approach.

8. PRESUPPOSITIONS. There are at least two ways in which denials and presuppositions interact. First, although presuppositions generally escape from the scope of any operator, including negation, in denials negation may be employed to cancel a presupposition. Second, denials, especially when used contrastively, give rise to presuppositions of their own.

At first glance, presupposition and implicature denials appear to be similar:

(62) a. Barney didn't paint his Buick pink: he doesn't own a Buick.
b. Barney owns a Buick.

Normally speaking, someone who uttered the first half of 62a would presuppose that 62b is true, but in this particular context the presupposition that Barney owns a Buick is cancelled. This pattern is reminiscent of the interaction between negation and scalar predicates, because scalar predicates normally have an 'at least' reading in the scope of negation, but not if negation is used as denial, for then they acquire the 'exactly' interpretation. However, I shall argue, against the unitarian view, that despite this similarity presupposition and implicature denials are quite different things. The only feature they have in common is that they require special interpretations and are accordingly marked, but otherwise they work in entirely different ways. This is not to say that we need a special-purpose explanation for presupposition denials. As in the case of implicature and form denials, presupposition denial can be explained entirely on independent grounds. It is just that these are different from what we need in the implicature and form cases.

In my opinion, all instances of presupposition denial are readily accounted for by a theory of presupposition which, for reasons that will soon become clear, I will refer to as the BINDING THEORY OF PRESUPPOSITION. The binding theory was originally proposed, curiously enough, by van der Sandt,[21] although he doesn't employ it in his theory of denial, presumably because it would force him to give up his unitarian stance on denials. I will try to show that, once we have given up our unitarian prejudices, the binding theory provides a perfectly natural account of presupposition denial.

If a speaker presupposes φ he assumes that, or pretends as if, φ is given. What does it mean to say that some piece of information φ is taken to be given? Part of the standard answer to this question is that φ can only be given if it is entailed by the common ground of the discourse (see, for example, Stalnaker 1973, 1974, Karttunen 1974, Heim 1983). Presupposed information is peculiar, however, in that the speaker may PRETEND as if it is given: should the common ground not entail that φ is the case, then his

[21] See van der Sandt 1989, 1992 and van der Sandt & Geurts 1991. For further developments and modifications of the theory, see Zeevat 1992, Sæbø 1993, Geurts 1995, 1998, and Geurts & van der Sandt 1998.

audience will in general be prepared to extend it so as to include φ—or in other words, the presupposition that φ is 'accommodated' to the common ground (Stalnaker 1973, 1974, Lewis 1979).

This, however, is only part of the answer. Suppose we have been told that Barney has a Buick, and the speaker goes on to claim that 'Barney's Buick is pink'. His utterance presupposes that Barney has a Buick, and if presupposition did come down to entailment, we would only have to check whether this follows from the common ground—which by supposition is the case. But surely this cannot be the whole story. We would like to say that the speaker intends the definite noun phrase *Barney's Buick* to refer to the same individual he introduced earlier on in the discourse, and as far as the notion of entailment goes there is no need to say anything of the sort. Hence, givenness cannot be reduced to entailment.

It is part of the traditional wisdom about definite noun phrases that they may be used anaphorically, to refer back to an object that is already part of the conversational record. The binding theory generalizes this idea to all presupposition-inducing expressions, so that, in effect, the notion of presupposition comes to subsume that of anaphora. Technically speaking, the theory is an extension of DISCOURSE REPRESENTATION THEORY (Kamp 1981, Kamp & Reyle 1993), and rests upon three principal assumptions. The first is that anaphora is a species of presupposition, and that the standard presupposition-inducing expressions (such as definite noun phrases, factives, aspectual verbs, and so on) differ from pronominal anaphors mainly in that they possess a richer semantic content. This difference explains why presupposition inducers, unlike anaphoric pronouns, may be interpreted by way of accommodation, which is the second key notion of the theory. Finally, it is assumed that the process of accommodation is subject to certain constraints.

Formulated in procedural terms, the binding theory predicts that if an utterance contains a presupposition-inducing element the hearer will initially attempt to bind the presupposition to a suitable antecedent, just as he would try to bind an ordinary anaphor. If the presupposition cannot be so bound, it will be accommodated, i.e., it will be inserted somewhere in the DISCOURSE REPRESENTATION STRUCTURE (DRS). In general the number of sites at which a presupposition may be accommodated is larger than one, and it is assumed, in Heim's terminology, that GLOBAL ACCOMMODATION is preferred to LOCAL ACCOMMODATION (Heim 1983). That is to say, if there are several DRSs in which a given presupposition might be accommodated, the least embedded DRS is preferred.

To illustrate these principles, let us consider exx. 63a and 63b.

(63) a. If anybody spoke the truth, then it was the Prime Minister who spoke the truth.
b. If it was the Prime Minister who spoke the truth, then the Queen has been lying.
c. Somebody spoke the truth.

Both 63a and 63b contain an *it*-cleft, which triggers the presupposition that 63c is true, but only a speaker who utters 63b presupposes that somebody spoke the truth; if he utters 63a he does not commit himself to this assumption. These observations are explained as follows. Let us assume, to begin with 63a, that at some stage in the interpretation of a token of this sentence, the DRS shown in 64a has been constructed.

(64) a. [: [x: x spoke the truth] ⇒ [z, u: PM z, u spoke the truth, u = z]]

In 64a and the following examples, underlining serves to indicate the presuppositions

triggered in a given sentence. In this case there are two such presuppositions: that someone spoke the truth (single underscore) and that there is a Prime Minister (double underscore). We now try, first, to link the information in [z: PM z] to an antecedent, but since no suitable antecedent is available, this presupposition will have to be accommodated. There are three DRSs in which it might be accommodated: the DRS in which the presupposition originates (the consequent of the conditional), the antecedent, and the principal DRS. But since, in general, there is a preference for global accommodation, the presupposition is accommodated at top level:

(64) b. [z: PM z, [x: x spoke the truth] ⇒ [u: u spoke the truth, u = z]]

Secondly, we try to link the information in [u: u spoke the truth] to a suitable antecedent, which may be done by equating x and z or, equivalently, by substituting one for the other (for perspicuity, I suppress the now redundant condition which is underlined in 64b).

(64) c. [z: PM z, [x: x spoke the truth] ⇒ [: x = z]]

This DRS adequately represents the preferred interpretation of 63a. In particular it does not follow from 64c that somebody spoke the truth, and thus it is correctly predicted that someone who utters 63a does not thereby commit himself to the assumption that 63c is true. The presupposition that originates in the consequent of the conditional is bound in the antecedent, just as a personal pronoun might have been bound, and in this sense is blocked.

The antecedent of 63b contains two presupposition inducers (the interpretation of the consequent may be left out of the account here), the initial representation of this sentence being something like 65a.

(65) a. [: [z, u: PM z, u spoke the truth, u = z] ⇒ [the Queen has been lying]]

Neither of the presuppositions triggered in the antecedent of the conditional can be bound, and both will therefore have to be accommodated. In this case, there are two DRSs in which they might be accommodated, and as before it is predicted that they will be accommodated in the principal DRS, which yields 65b.

(65) b. [z, u: PM z, u spoke the truth, [: u = z] ⇒ [the Queen has been lying]]

In contradistinction to 64c, this DRS does entail that someone spoke the truth, and thus the theory accounts for the observation that a speaker who volunteers 63b commits himself to the assumption that 63c is true.[22]

To sum up, the binding theory of presupposition adheres to the following principles:

[22] The binding theory neither is nor incorporates a theory of scope. To begin with, the notion that, in principle, presuppositions want to be bound would not be part of a theory that sought to deal with presuppositions in terms of scope. But even if we leave the crucial notion of binding out of account, there is an important difference. Consider the following:

(i) Many boys flirted with some teachers.
(ii) [many x: boy x] ([some y: teacher y](x flirted with y))
(iii) [some y: teacher y]([many x: boy x](x flirted with y))

In (i) it is natural to give wide scope to *many boys,* as in (ii), and possible to give *some teachers* wide scope, as in (iii). But in either case, *many boys* (or *some teachers*) may be read as presupposing that there are or were boys (teachers). So the question of whether or not an expression gives rise to a presupposition is distinct from the question of whether it has wide scope. The fundamental difference between presupposition projection and scope is that the former is a pragmatic phenomenon, which can only be explained with reference to a level of discourse representation, whereas the latter is, to some degree at least, a grammatical phenomenon, in the following sense. If we speak about an expression α in terms of scope, we are referring to some α-sized unit, i.e. α itself or some semantic entity corresponding with α. But if we speak about a presupposition triggered by α neither α itself nor its semantic representation need be at issue. It is this difference which rules out a treatment of presupposition in terms of scope. For further discussion of this point, see Geurts 1998 and Geurts & van der Sandt 1998.

(i) a presupposition must either be bound to a suitable antecedent or be accommodated; (ii) if a presupposition can be bound, this option is preferred to accommodation; (iii) if a presupposition must be accommodated, it is preferably accommodated in the least embedded DRS. For our current purposes, it is the third assumption which is the crucial one. It says that there is a preference for accommodating a presupposition (if it must be accommodated) globally—i.e. in the least embedded DRS. This preference holds only as long as all things are equal. If they are not, it may be overwritten. There are various circumstances that may force a presupposition to be accommodated locally (at some embedded level in the DRS); see Geurts 1995:28ff for more extensive discussion of this point. To begin with, a presupposition may be prevented from surfacing at a certain level because the resulting DRS would be inconsistent, as in ex. 66.

(66) a. Harry is a bachelor, so it wasn't Harry's wife who shot the burglar.
b. [x: Harry x, bachelor x, ¬ [u: u is x's wife, u shot the burglar]]
c. [x, u: Harry x, bachelor x, u is x's wife, ¬ [: u shot the burglar]]

If the presupposition in 66b should be accommodated at top level, the resulting DRS, 66c, could never be true, and in fact 66b seems to capture the intended interpretation of 66a rather well as it stands, so it is plausible to assume that a presupposition may be prevented from floating up to some superordinate DRS in the face of imminent inconsistency.

It is a general constraint on presupposition projection that the consistency of the discourse representation be preserved in the process. In 66a, the presupposition that Harry has a wife is blocked because it is inconsistent with the speaker's explicit claim that Harry is a bachelor. However, there are other, more subtle, sources of potential inconsistency that may force a presupposition to accommodate locally. One such source is the information the speaker conveys about his own knowledge; 67 is a case in point.

(67) a. I don't know if Harry is married . . .
b. but the woman we saw him with last night certainly wasn't his wife.

By uttering 67 the speaker does not presuppose that Harry has a wife, although 67b contains a definite NP that triggers this presupposition. Hence the presupposition that Harry has a wife is prevented from accommodating in the principal DRS and is forced to accommodate within the scope of the negation instead. Why? The answer is fairly obvious: if the presupposition had been accommodated globally, the resulting DRS would contradict the speaker's explicit claim that he doesn't know if Harry is married.

In 67 the speaker asserts that he doesn't know if Harry is married, but he might have preferred to convey this information in a more implicit way, as in 68.

(68) a. It is possible that Harry is married . . .
b. but the woman we saw him with last night certainly wasn't his wife.

In 68a the speaker doesn't say in so many words that he doesn't know if Harry has a wife, but nonetheless we infer from his utterance that this is so, and accordingly the presupposition that Harry has a wife will be accommodated locally. My claim is that presupposition denials are not substantially different from the instances of local accommodation that I have discussed so far.[23]

(69) a. It wasn't Harry's wife who shot the burglar: Harry is a bachelor.

[23] Despite the fact that it is superficially similar to the binding theory, what I have elsewhere called the SATISFACTION THEORY (e.g., Karttunen 1974, Heim 1983) cannot account for examples like 69a–c. See Geurts 1995, 1996, for further discussion.

b. Harry is a bachelor, so it wasn't Harry's wife who shot the burglar. (= 66a)

c. If Harry is a bachelor, then it wasn't Harry's wife who shot the burglar.

In my view, there are no fundamental differences between 69a, on the one hand, and 69b and 69c, on the other: in all cases the presupposition that Harry has a wife is accommodated locally, because otherwise some form of inconsistency would arise. Sentence 69a is different from 69b and 69c in just one respect: where in 69b and 69c, the information that Harry is a bachelor precedes the negated sentence, in 69a this information is presented afterwards. This has some consequences for the process of interpretation, but not for its outcome.

In 69b and 69c the information that Harry doesn't have a wife is part of the local context in which the definite NP *Harry's wife* is evaluated (at the point at which the NP is evaluated it is already given that Harry is a bachelor), and therefore the issue of where the corresponding presupposition must be accommodated is settled right away. In 69a, by contrast, the hearer doesn't know for sure that the speaker intends this presupposition to be accommodated locally until he has received the second sentence; the final interpretation of the first half of 69a is partly dependent upon the interpretation of the second half. Of course, this is not a peculiarity of presupposition denials; it also occurs in implicature denials, as in 70, for example.

(70) The soup isn't hot: it's scalding.

Here the adjective *hot* cannot get its final interpretation until the second part of the utterance is in. It is to signal this, presumably, that the intonation contour that generally accompanies denials like 69a and 70 is a marked one, with a final rise within the negative clause (Horn 1989:374; also Ladd 1980:145ff). In cases like 69b and 69c, by contrast, marked intonational devices are not required.

Although there are certain similarities between implicature and presupposition denials, as we have just seen, the differences are more important. In my analysis of implicature and form denials, semantic transfer plays the central role; in the analysis of presupposition denials it doesn't figure at all. An expression doesn't change or shift its meaning just because a presupposition it induces lands within the scope of a negation operator. Rather, what happens is the following. The binding theory predicts that, if a presupposition must be accommodated, it may in general be accommodated at several sites. In this sense, the theory predicts that the interpretation of presuppositional expressions is ambiguous, if you will. (But note that if we call this ambiguity, practically all occurrences of anaphoric pronouns will be ambiguous, too, because in the majority of cases there is more than one antecedent that a pronoun may link up to.) The theory predicts, among other things, that this 'ambiguity' will be resolved in favor of global accommodation, which is therefore the default option, although under special circumstances local accommodation may be preferred. The meaning of a presuppositional expression remains the same regardless of whether the presupposition it triggers is accommodated globally or locally. For example, although the presupposition triggered by the *it*-cleft is bound in 63a and accommodated in 63b (and, consequently, the latter but not the former implies that someone spoke the truth), the meaning of the *it*-cleft is the same in both cases.

This account of presupposition denial has a number of advantages, of which I want to mention two. First, it is to be had for free, since it is just a corollary of a theory of presupposition that is motivated on independent grounds. Second, it explains examples like 71 (= 33), which, as I argued in §3, unitarian theories have trouble with.

(71) The king of France didn't have dinner with the Canadian empress: France doesn't have a king.

Normally speaking, an utterance of 71 would presuppose that Canada has an empress; the speaker's objection is restricted, apparently, to the notion that France should have a king. A unitarian theory is likely to have problems with this because the type of negation it requires is directed at complete utterances, rather than parts or aspects of utterances. My theory, however, has no problems at all with this type of example.

(72) a. [: ¬ [x, z: king of France x, empress of Canada z, x dined with z]]
b. [z: empress of Canada z, ¬ [x: king of France x, x dined with z]]

Supposing that something like 72a is the semantic representation of the first part of 71, the binding theory predicts that both presuppositions will be accommodated globally, but since global accommodation of the first presupposition (that France has a king) would be inconsistent with the speaker's assertion that there is no king of France, this presupposition will be accommodated locally. Hence, the interpretation that is predicted for 71 is 72b, and this interpretation entails that Canada has an empress, which is what we wanted.

I conclude with a brief comment on a presupposition that typically arises in contrastive (or reparative) denials, such as the following, modeled after one of Kempson's (1986) examples.

(73) Julius didn't lose a finger but an entire hand.

A speaker who utters this would focus on the NP *a finger,* making it understood that he accepts that Julius did lose something, a body part say, although it was not just a finger. This inference (that Julius lost something) I consider to be of a presuppositional nature, because it exhibits the projection behavior that is typical of presuppositions.

(74) a. More likely than not, Julius didn't lose a finger but an entire hand.
b. Walter believes that Julius didn't lose a finger but an entire hand.

In these examples, 73 is embedded in a nonentailed position, but intuitively we would nonetheless infer from 74a or 74b that, according to the speaker, Julius lost something. Such observations are easily explained if we assume that in a focus-background structure, (an existential instantiation of) the background is presupposed.[24] Since in 73 a focus-background division is indicated, within the scope of the negation operator, in which the semantic correlate of *a finger* is focus and 'Julius lost x' is backgrounded, we predict that the initial semantic representation of the denial in 73 is something like 75a.

(75) a. [u: Julius u, ¬ [x, y: body part x, u lost x, finger y, u lost y]]
b. [u, x: Julius u, body part x, u lost x, ¬ [y: finger y, u lost y]]

The binding theory predicts that, in the absence of evidence to the contrary, the presupposition in 75a must be accommodated globally, so that the interpretation that we predict is 75b. As I have just argued, this is the correct result. So the binding theory not only accounts for presupposition denials: it also helps to explain the presuppositional effects of focusing in general, and of focusing in denials, in particular.

REFERENCES

ANSCOMBRE, J. C., and O. DUCROT. 1976. L'argumentation dans la langue. Langages 42.5–27.

[24] This analysis is developed in greater detail in Geurts & van der Sandt 1998.

——, and ——. 1977. Deux *mais* en français? Lingua 43.23–40.
BARTSCH, RENATE. 1987. Context-dependent interpretations of lexical items. Foundations of pragmatics and lexical semantics, ed. by Jeroen Groenendijk, Dick de Jongh, and Martin Stokhof, 1–26. Dordrecht: Foris.
BAUGH, ALBERT C., and THOMAS CABLE. 1978. A history of the English language. 3rd edn. London: Routledge and Kegan Paul.
BIERWISCH, MANFRED. 1983. Semantische und konzeptuelle Repräsentation lexikalischer Einheiten. Untersuchungen zur Semantik, ed. by R. Ruzicka and W. Motsch, 61–99. Berlin: Akademie-Verlag.
BOSCH, PETER. 1984. Lexical learning, context dependence, and metaphor. Bloomington: Indiana University Linguistics Club.
——. 1985. Context dependence and metaphor. The ubiquity of metaphor, ed. by Wolf Paprotté and René Dirven, 141–76. Amsterdam: John Benjamins.
CARSTON, ROBYN. 1988. Implicature, explicature, and truth-theoretic semantics. Mental representation: The interface between language and reality, ed. by Ruth M. Kempson, 155–81. Cambridge University Press.
——. 1996. Metalinguistic negation and echoic use. Journal of Pragmatics 25.309–30.
COPESTAKE, ANNE, and TED BRISCOE. 1995. Semi-productive polysemy and sense extension. Journal of Semantics 12.15–67.
DUCROT, O. 1972. Dire et ne pas dire. Paris: Hermann.
FOOLEN, AD. 1991. Metalinguistic negation and pragmatic ambiguity: Some comments on a proposal by Laurence Horn. Pragmatics 1.217–37.
GAZDAR, GERALD. 1979. Pragmatics: Implicature, presupposition, and logical form. New York: Academic Press.
GEURTS, BART. 1995. Presupposing. Stuttgart: University of Stuttgart dissertation.
——. 1996. Local satisfaction guaranteed. Linguistics and Philosophy 19.259–94.
——. 1998. Presuppositions and anaphors in attitude contexts. Linguistics and Philosophy, to appear.
——, and ROB A. VAN DER SANDT. 1998. Domain restriction. Focus: linguistic, cognitive, and computational perspectives, ed. by Peter Bosch and Rob A. van der Sandt. Cambridge University Press, to appear.
GRICE, H. PAUL. 1989. Studies in the way of words. Cambridge, MA: Harvard University Press.
HEIM, IRENE. 1983. On the projection problem for presuppositions. West Coast Conference on Formal Linguistics 2.114–26.
HERSKOVITS, A. 1986. Language and spatial cognition. Cambridge: Cambridge University Press.
HOPPER, PAUL J., and ELIZABETH C. TRAUGOTT. 1993. Grammaticalization. Cambridge: Cambridge University Press.
HORN, LAURENCE R. 1972. On the semantic properties of logical operators in English. Los Angeles, CA: UCLA dissertation.
——. 1985. Metalinguistic negation and pragmatic ambiguity. Language 61.121–74.
——. 1989. A natural history of negation. Chicago: University of Chicago Press.
——. 1990. Showdown at truth-value gap: Burton-Roberts on presupposition. Journal of Linguistics 26.483–503.
——. 1992. The said and the unsaid. Proceedings of SALT II, ed. by Chris Barker and David Dowty, 163–192. Columbus: Ohio State University.
KAMP, HANS. 1981. A theory of truth and semantic representation. Formal methods in the study of language, ed. by Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 277–322. Amsterdam: Mathematical Centre Tracts 135.
——, and UWE REYLE. 1993. From discourse to logic. Dordrecht: Kluwer.
KARTTUNEN, LAURI. 1974. Presupposition and linguistic context. Theoretical Linguistics 1.181–94.
KEMPSON, RUTH M. 1986. Ambiguity and the semantics-pragmatics distinction. Meaning and interpretation, ed. by Charles Travis, 77–103. Oxford: Blackwell.
KOENIG, J.-P. 1991. Scalar predicates and negation: Punctual semantics and interval interpretations. Chicago Linguistic Society 27(2).331–44.

LADD, D. ROBERT. 1980. The structure of intonational meaning. Bloomington: Indiana University Press.
LADUSAW, WILLIAM A. 1996. Negation and polarity items. The handbook of contemporary semantic theory, ed. by Shalom Lappin, 321–41. Oxford: Blackwell.
LAKOFF, GEORGE. 1987. Women, fire, and dangerous things. Chicago: University of Chicago Press.
LEVINSON, STEPHEN C. 1983. Pragmatics. Cambridge: Cambridge University Press.
LEWIS, DAVID K. 1973. Counterfactuals. Cambridge, MA: Harvard University Press.
——. 1979. Scorekeeping in a language game. Journal of Philosophical Logic 8.339–59.
MCCAWLEY, JAMES D. 1991. Contrastive negation and metalinguistic negation. Chicago Linguistic Society 27(2).189–206.
MORGAN, J. L. 1978. Two types of convention in indirect speech acts. Syntax and semantics, vol. 9: Pragmatics, ed. by Peter Cole, 261–80. New York: Academic Press.
NUNBERG, GEOFFREY. 1978. The pragmatics of reference. Bloomington: Indiana University Linguistics Club.
——. 1979. The non-uniqueness of semantic solutions: Polysemy. Linguistics and Philosophy 3.143–84.
——. 1995. Transfers of meaning. Journal of Semantics 12.109–32.
PUSCH, LUISE F. 1975. Über den Unterschied zwischen *aber* und *sondern* oder die Kunst des Widersprechens. Syntaktische und semantische Studien zur Koordination, ed. by István Bátori, Luise F. Pusch, Jurij L. Levin, Werner Abraham, Wolfram Bublitz, and Manfred von Roncador, 45–62. Tübingen: Gunter Narr.
QUINE, W. V. O. 1951. Mathematical logic. Rev. edn. Cambridge, MA: Harvard University Press.
——. 1969. Ontological relativity and other essays. New York: Columbia University Press.
SADOCK, JERROLD. 1984. Whither radical pragmatics? Meaning, form, and use in context: linguistic applications, ed. by Deborah Schiffrin, 139–49. Washington, DC: Georgetown University Press.
SÆBØ, KJELL JOHANN. 1993. Anaphoric presuppositions, accommodation, and zero anaphora. Presupposition, ed. by Hans Kamp, 135–92. University of Amsterdam: DYANA-2 Report R2.2.A, part 2.
SAG, IVAN. 1981. Formal semantics and extralinguistic content. Radical pragmatics, ed. by Peter Cole, 273–94. New York: Academic Press.
SPERBER, DAN, and DEIRDRE WILSON. 1986. Relevance: Communication and cognition. Oxford: Blackwell.
STALNAKER, ROBERT C. 1973. Presuppositions. Journal of Philosophical Logic 2.447–57.
——. 1974. Pragmatic presuppositions. Semantics and philosophy, ed. by Milton K. Munitz and Peter K. Unger, 197–213. New York: New York University Press.
SWEETSER, EVE. 1990. From etymology to pragmatics. Cambridge: Cambridge University Press.
VAN DER SANDT, ROB A. 1988. Context and presupposition. London: Routledge & Kegan Paul.
——. 1989. Presupposition and discourse structure. Semantics and contextual expression, ed. by Renate Bartsch, Johan van Benthem, and Peter van Emde Boas, 267–94. Dordrecht: Foris.
——. 1991. Denial. Chicago Linguistic Society 27(2):331–44.
——. 1992. Presupposition projection as anaphora resolution. Journal of Semantics 9.333–77.
——. 1998. Denial and negation. Linguistics and Philosophy, to appear.
——, and BART GEURTS. 1991. Presupposition, anaphora, and lexical content. Text understanding in LILOG, ed. by Otthein Herzog and Claus-Rainer Rollinger, 259–96. Berlin: Springer-Verlag.
WIERZBICKA, ANNA. 1994. Semantic primitives across languages: A critical review. Semantic and lexical universals, ed. by C. Goddard and A. Wierzbicka, 445–500. Amsterdam: John Benjamins.
WILSON, DEIRDRE, and DAN SPERBER. 1988. Representation and relevance. Mental representation: The interface between language and reality, ed. by Ruth M. Kempson, 133–53. Cambridge: Cambridge University Press.

ZEEVAT, HENK. 1992. Presupposition and accommodation in update semantics. Journal of Semantics 9.379–412.

University of Osnabrück
49069 Osnabrück, Germany
[geurts@cl-ki.uni-osnabrueck.de]

[Received 2 August 1996;
revision received 19 June 1977;
accepted 8 September 1997]