

Referring Expressions and Conversational Implicature*

Andrew Kehler
University of California San Diego

Gregory Ward
Northwestern University

1. Introduction

English, like all natural languages, provides its speakers with an impressive array of referring expressions from which to choose. The question is, of the many forms that are available (e.g., indefinite description, definite description, proper name, bare NP, demonstrative, pronoun, null), which factors guide a speaker in selecting from among these alternatives to facilitate successful communication? As is well known, the choice is far from random but instead heavily constrained by such properties as information status (including salience and givenness), uniqueness, individuability, and hearer's (assumed) knowledge of the world, among a variety of others.

Starting with Grice (1975), it has been noted by several authors that a speaker's choice of referring expression can lead to inferences that display the hallmarks of conversational implicature (Prince 1981b; Hawkins 1991; Gundel, Hedberg, and Zacharski 1993). In considering a sentence of the form *X is meeting a woman this evening*, for example, Grice suggests that anyone using such an expression:

...would normally implicate that the person to be met was someone other than X's wife, mother, sister, or perhaps even close platonic friend. (p. 56)

He adds:

When someone, by using the form of expression *an X*, implicates that the *X* does not belong to or is not otherwise closely connected with some identifiable person, the implicature is present because the speaker has failed to be specific in a way in which he might have been expected to be specific, with the consequence that it is likely to be assumed that he is not in a position to be specific. (p. 57)

There are, of course, many ways in which a speaker can 'fail to be specific' in choosing what to say. A classic instance in Gricean pragmatics involves SCALAR relations based on Grice's (1975) Maxim of Quantity:

(1) MAXIM OF QUANTITY

Q1: Make your contribution as informative as is required (for the current purposes of the exchange).

Q2: Do not make your contribution more informative than is required.

*We gratefully thank Barbara Abbott, Betty Birner, and Larry Horn for valuable comments on earlier drafts.

As shown by Horn (1984), Hirschberg (1991), Levinson (2000), inter alia, the implicatures associated with linguistic scales are explainable in terms of the Maxim of Quantity, whereby a speaker's use of a 'weaker' expression positioned on a Quantity-based scale (e.g., *some*) conversationally implicates that the speaker was not in a position to use a 'stronger' expression (e.g., *all*). Previous analyses have in fact posited that certain properties encoded by referring expressions are similarly governed by Quantity-based scales, and hence give rise to conversational implicatures. Evidence for such implicatures is provided by the hallmark diagnostics of CANCELLABILITY and REINFORCEABILITY (Grice 1975, Horn 1984, Sadock 1978, Hirschberg 1991). For instance, in example (2), from Hawkins 1991, an implicature associated with an indefinite noun phrase is felicitously cancelled without contradiction. Likewise in (3), a naturally-occurring example from Gundel et al. (1993, henceforth GHZ), an analogous implicature is felicitously reinforced without redundancy:

- (2) I met *a student* before class. *A student* came to see me after class as well – in fact it was the same student I had seen before.
[=Hawkins' ex. 12 with follow-on introduced in the text]
- (3) But forged provenance papers still did not mean that the kouros was fake....The Getty decided that the fake documents were not reason enough to ask Mr. Becchina, the Basel dealer who had sold the kouros, to take back the sculpture. (Attempts by The Times to reach Mr. Becchina were unsuccessful.) Then last April, an independent scholar in London, Jeffrey Spier, was shown a photograph of a fake torso of a kouros, belonging to A BASEL DEALER (NOT MR. BECCHINA) that looked similar to the Getty's sculpture.
[= GHZ's ex. 51, emphasis theirs]

Hawkins (1991) and GHZ (1993) offer analyses of referring expressions that are governed by Quantity-based scales defined in terms of UNIQUENESS and COGNITIVE STATUS, respectively. In this paper we review these analyses, and argue that the implicatures evident in examples such as (2) and (3) are based instead on the notion of FAMILIARITY. Such NONFAMILIARITY IMPLICATURES result from a speaker's decision not to use a referring expression that would ordinarily suggest that the hearer is already familiar with the intended referent.

2. Three Proposals

We begin with a brief overview of the Hawkins and GHZ proposals, and then present our own analysis based on the notion of familiarity. In Section 3 we examine various data designed to test the predictions of each of the three analyses under consideration.

2.1. The Hawkins Analysis

Hawkins (1991) presents an implicature-based analysis of the contrast between the definite and indefinite articles in English (*the* and *a/some*, respectively). For Hawkins, the difference between these two forms hinges on whether the intended referent is unique within a contextually-determined 'association set' called a P-SET. The notion of a P-set is crucial for Hawkins because a definite NP such as *the professor* may be used felicitously – e.g., if a single professor has been recently mentioned, or if a university-level class has been mentioned with which a unique professor can be inferentially associated – even though there are obviously multiple professors in

the universe. With the notion of a P-set in hand, Hawkins' system is characterized by the following four implicatures (1991:429):¹

- I. *THE: conventional implicature: P-membership*
The conversationally implicates that there is some subset of entities, {P}, in the universe of discourse which is mutually manifest to S & H on-line and within which definite referents exist and are unique.
- II. *THE: I-implicature: P-membership*
If there is a given P-set of entities in the universe of discourse which is mutually manifest to S & H on-line and in which the definite referent or referents can exist and be unique, then the use of *the* (conversationally) I-implicates that these referents are indeed members of this P-set.
- III. *A/SOME: Q-implicature: non-uniqueness*
If the speaker could use *the* and instead selects *a* or *some*, then he (conversationally) Q-implicates non-uniqueness, i.e. that there is at least one entity satisfying the description of the indefinite NP and non-identical to the individual or set of individuals whose existence is entailed by this indefinite NP.
- IV. *A/SOME: I-implicature: P-membership*
If there is a given P-set of entities in the universe of discourse which is mutually manifest to S & H on-line and in which the indefinite referent or referents may exist, then the use of *a* or *some* (conversationally) I-implicates that these referents are indeed members of this P-set.

Implicature III is the central one for our purposes. This formulation of the implicature situates <*the, a/some*> on a HORN SCALE (Atlas and Levinson 1981), that is, an entailment-based implicational hierarchy. The semantic basis for this Horn scale is the (binary) concept of uniqueness. As such, the use of *a/some* instead of *the* conversationally implicates, by the Maxim of Quantity, that the referent cannot be uniquely identified by the hearer within some mutually manifest P-set.

For Hawkins, example (2), repeated below as (4), supports the existence of this scale:

- (4) I met *a student* before class. *A student* came to see me after class as well – in fact it was the same student I had seen before. [= (2)]

In his analysis, the use of the second occurrence of *a student* conversationally implicates that the student is not unique within any P-set that is mutually manifest to the speaker and hearer. This implicature is cancelled by the *in fact* clause, in light of its uniqueness in the P-set that includes those entities that have been previously mentioned in the discourse.

2.2. The GHZ Analysis

Whereas Hawkins posits a Horn scale for the articles *the* and *a/some* based on the concept of uniqueness, GHZ go considerably further by placing a range of different categories of referring

¹The second and fourth of these implicatures make reference to Levinson's (1987) notion of I-implicature, which is a refinement of the second half of Grice's Maxim of Quantity. As such, it represents a subcategory of Horn's (1984) R-implicatures. These implicatures are not our primary concern here, however, and will not be considered further.

expressions on such a scale based on a notion of COGNITIVE STATUS. Specifically, they propose a GIVENNESS HIERARCHY consisting of six possible cognitive statuses that a referent may have. Below each cognitive status are the (English) referential expressions that encode it.

in focus	>	activated	>	familiar	>	uniquely identifiable	>	referential	>	type identifiable
<i>it</i>		<i>that</i> <i>this</i> <i>this N</i>		<i>that N</i>		<i>the N</i>		indef <i>this N</i>		<i>a N</i>

(GHZ, p. 275)

The cognitive statuses are defined as follows, along with an example of each:

TYPE IDENTIFIABLE: The addressee is able to access a representation of the type of object described by the expression.

(5) *I couldn't sleep last night. A dog (next door) kept me awake.*

REFERENTIAL: The speaker intends to refer to a particular object or objects.

(6) *I couldn't sleep last night. This dog (next door) kept me awake.*

UNIQUELY IDENTIFIABLE: The addressee can identify the speaker's intended referent on the basis of the nominal alone.

(7) *I couldn't sleep last night. The dog (next door) kept me awake.*

FAMILIAR: The addressee is able to uniquely identify the intended referent because he already has a representation of it in memory (in long-term memory if it has not been recently mentioned or perceived, or in short-term memory if it has).

(8) *I couldn't sleep last night. That dog (next door) kept me awake.*

ACTIVATED: The referent is represented in current short-term memory.

(9) *My neighbor has a dog. This dog kept me awake last night.*

IN FOCUS: The referent is not only in short-term memory, but is also at the current center of attention.

(10) *A dog was in the front yard last night. It kept me awake.*

The weakest status is TYPE IDENTIFIABLE, which requires only the hearer's presumed familiarity with the kind denoted by the noun (or common noun phrase) in question. Stronger values to the left must satisfy that condition but also impose additional requirements; as such, each cognitive status logically implies the weaker values to its right in the hierarchy. For example, an IN FOCUS referent (at the current center of attention) is necessarily also ACTIVATED (in short-term memory),

FAMILIAR (in memory), and so forth. Thus, a form normally associated with a particular cognitive status could in principle be used to refer to an entity with a stronger one. However, in a survey of data across several languages, GHZ found that with some exceptions, each form was found almost exclusively to be associated with the status with which it is correlated. One of these exceptions in English is the case of definite descriptions, which were found not only with UNIQUELY IDENTIFIABLE referents, but with all stronger statuses.

GHZ address this correlation and the exception for definite descriptions by appealing to both parts of Grice's Maxim of Quantity, repeated below:

(11) MAXIM OF QUANTITY

Q1: Make your contribution as informative as is required (for the current purposes of the exchange).

Q2: Do not make your contribution more informative than is required. [= (1)]

GHZ argue that the fact that certain types of referring expressions (e.g., demonstratives) are not typically used to refer to referents having a stronger cognitive status is the result of Q1-based implicatures; i.e., their use conversationally implicates that the stronger status does not hold (i.e., an instance of a Q-BASED IMPLICATURE in Horn's (1984) terms). In particular:

Interaction of this maxim [=Q1] with the Givenness Hierarchy correctly predicts that (a) an indefinite article will normally not be used for referents that are uniquely identifiable, since this form explicitly signals only type identifiability, and (b) for referents that are in focus, an unstressed personal pronoun or zero, which explicitly delimits the set of possible referents to those that are in focus, will normally be chosen over a demonstrative pronoun, which gives less information about cognitive status because it only requires that the referent be activated. (GHZ, p. 299)

The fact that definite descriptions are exceptionally found not only for UNIQUELY IDENTIFIABLE referents but also for referents with all stronger statuses lead GHZ to appeal instead to Q2 for those cases. Unlike bare demonstratives and pronouns, definite descriptions typically contain the descriptive content necessary to uniquely identify the referent. According to GHZ, an explicit signal of a stronger status is therefore unnecessary, per the second half of the maxim (i.e., by an R-BASED IMPLICATURE in Horn's (1984) terms). This, they claim, follows from Atlas and Levinson's (1981) observation that Q2 induces stereotypical interpretations. Since most referents that are UNIQUELY IDENTIFIABLE are also FAMILIAR, the use of a definite description that signals the former will generally conversationally implicate the latter.²

GHZ cite Hawkins' example (4) as providing support for the Givenness Hierarchy. In their analysis, the use of the second occurrence of *a student* conversationally implicates that its referent is merely TYPE IDENTIFIABLE. This implicature is subsequently cancelled in light of the fact that the referent also holds the stronger status of ACTIVATED. GHZ also cite example (3), repeated below as (12), as a source of further support:

²GHZ concede, however, that this story cannot explain why full NPs are found instead of bare demonstratives and pronominals for cases that are (at least) ACTIVATED. They briefly address this issue in their footnote 32 (1993:300), citing ambiguity resolution and global focus shift as possible reasons why a fuller NP might be used instead of a pronoun. But the felicity of simple examples shows that there must be more to the story:

(i) I was talking to a guy yesterday, and the guy said...

Note that the NP *a guy* must give rise to an IN FOCUS referent, since a pronominal reference in the second clause would also be felicitous.

- (12) But forged provenance papers still did not mean that the kouros was fake.... The Getty decided that the fake documents were not reason enough to ask Mr. Becchina, the Basel dealer who had sold the kouros, to take back the sculpture. (Attempts by The Times to reach Mr. Becchina were unsuccessful.) Then last April, an independent scholar in London, Jeffrey Spier, was shown a photograph of a fake torso of a kouros, belonging to A BASEL DEALER (NOT MR. BECCHINA) that looked similar to the Getty's sculpture. [= (3)]

According to their account, the use of *a Basel dealer* in (12) implicates that the referent is merely TYPE IDENTIFIABLE, which is then reinforced by a denial that an existing ACTIVATED entity is the intended referent.

2.3. An Alternative Analysis

We have so far described two analyses that characterize implicatures that result from particular choices of referring expressions. For Hawkins, implicatures in examples such as (4) derive from the uniqueness restriction entailed by the use of *the*, such that the use of *a/some* implicates non-uniqueness within the relevant P-set. For GHZ, the implicatures in examples such as (4) and (12) result from the speaker's decision to use an expression associated with a particular cognitive status, which licenses the inference that the referent does not satisfy the constraints associated with stronger statuses.

There is, however, another possible explanation for the implicatures associated with examples such as (4) and (12), namely one which characterizes the relevant inferences as NONFAMILIARITY IMPLICATURES. To make the notion of familiarity concrete we appeal to GHZ's criterion, repeated here:

FAMILIAR: The addressee is able to uniquely identify the intended referent because he already has a representation of it in memory (in long-term memory if it has not been recently mentioned or perceived, or in short-term memory if it has).

Importantly, familiarity is distinct from uniqueness. The definite description *the book that John is reading this week* may constrain reference to a single entity, but one that is nonetheless not familiar to the hearer.

Our familiarity-based analysis (henceforth FA) simply states that in unexceptional contexts, the speaker's failure to use a referring expression that indicates hearer-familiarity conversationally implicates that the referent is nonfamiliar. Under this account, the implicatures associated with (4) and (12) are based on nonfamiliarity rather than with uniqueness or cognitive status. Specifically, in (4) a nonfamiliarity implicature associated with the second occurrence of the phrase *a student* is felicitously cancelled by the speaker's subsequent use of the NP *the same student I had seen before* to identify it with the (now FAMILIAR) referent of the first occurrence. In example (12), a nonfamiliarity implicature associated with the NP *a Basel dealer* is felicitously (i.e. non-redundantly) reinforced by the phrase *not Mr. Becchina*, indicating that the referent is not one that was previously mentioned (and hence FAMILIAR).

To sum up so far, both Hawkins and GHZ utilize entailment-based implicational hierarchies in their analyses of referential form, in which linguistic expressions are associated with different points on a scale. Hawkins limits his analysis to the distinction between *a/some* and *the* in English, whereas GHZ address a broader range of forms in several languages. Our

familiarity-based analysis, on the other hand, assimilates the facts at hand to a subcase of Grice's notion of specificity: When a speaker fails to select a referring expression that would be expected to indicate familiarity, the referent is conversationally implicated to be nonfamiliar. In the next section we examine a variety of predictions each analysis makes with respect to the generation of conversational implicatures.

3. Referring Expressions and Quantity-Based Conversational Implicature

In the previous section, we saw that the three analyses under discussion offer different explanations for the implicatures that result from the choice of referential form, such as those evident in examples (4) and (12). To distinguish among the proposals, we will need to consider additional data.

We begin by comparing the FA to Hawkins' analysis. First, we agree with Hawkins that uniqueness can serve as a coherent semantic basis for implicature cancellation. Larry Horn (p.c.) in fact points out that numerous examples attest to the existence of a pragmatically robust *<the,a>* scale; examples (13a-c) are three of a great many that can be found on the internet:

- (13) a. Over the nineteenth century, Britain became a, if not the, world power.
[eserver.org/cultronix/sigel/]
- b. Decision making is a, but not the, fundamental construct in design.
[dbd.eng.buffalo.edu/papers/DR.position.htm]
- c. I find it humbling that the Torah, uniquely in the religious literature of mankind, begins by setting out these themes – insisting that parenting is not only a, but the supreme religious act.
[http://www.chiefrabbi.org/speeches/parent.htm]

The felicitous cancellations in these examples provide evidence for the existence of a Horn scale defined in terms of uniqueness.

Nonetheless, we disagree with Hawkins that example (4) – in which, unlike (13a-c), the relevant noun phrases differ with respect to more than just their respective articles – belongs to this class of examples. Instead, we argue that it naturally patterns with other examples that are captured by the FA but which either require extensions to Hawkins' analysis or receive a considerably less natural treatment within it. First, we find that a greater range of opposing linguistic forms than merely *the* and *a/some* participate in the types of implicatures found in (4) and (12). Consider (14) and (15):

- (14) A: Did any students bother you during office hours today?
B: Yes, a student came by. In fact, it was that weird guy who sits in the back of the class.
- (15) A: Did any students bother you during office hours today?
B: Yes, this student came by. In fact, it was HIM. [pointing to a salient guy standing nearby]

In both of these exchanges, the status indicated by the first referring expression suffices to answer A's question; an implicature associated with that form is then cancelled with the stronger

form in a subsequent utterance. These cancellations do not seem qualitatively different from the one exhibited in (4), yet they are not covered by an analysis that opposes only *the* and *a/some*. On the other hand, as with (4), the felicitous cancellations in examples (14) and (15) are captured by the FA.³

Second, the FA provides a more natural explanation for the facts in cases in which the uniqueness and familiarity properties of the relevant noun phrases diverge. The first includes cases that involve two referring expressions which differ with respect to familiarity but not uniqueness. Example (16) is such a case:

- (16) A: Has John read any good books lately?
B: Yes, the book that John is currently reading – in fact the one I showed you yesterday when we were at the bookstore – is reportedly a real page-turner.

Example (16) is felicitous as long as (the speaker believes that) the hearer does not know which book John is currently reading (i.e., the definite merely signals uniqueness), but does know which book the speaker showed him yesterday. This case patterns with examples (4), (14), and (15) in exhibiting a cancellation of nonfamiliarity.

Despite the fact that the two relevant noun phrases in B's utterance in (16) do not differ with respect to uniqueness, Hawkins' account may provide an analysis of this case, although we believe it is considerably less intuitive than the FA analysis and requires appeal to two extensions beyond the four implicatures listed previously. First, although the noun phrase *the book that John is currently reading* in (16) is NOT mutually manifest to the speaker and hearer with respect to any P-set in the scenario considered here, Hawkins allows such manifestness to be freely accommodated by the hearer:

As I see it, [such] definite descriptions are extending the members of otherwise mutually known P-sets to include a novel entity, whose existence within the set is none the less accepted by the hearer... The speaker is simply telling the hearer to extend the relevant P-sets on these occasions, and expects the hearer to accept these extensions, presuming that they are at least compatible with what he knows and that he has no reasons not to accept them. (p. 411)

Such novel entities are precisely those that can be felicitously mentioned with a *the*-NP but that are excluded by our notion of familiarity. Second, Hawkins remarks that the selection of a given choice of P-set on a particular occasion of use is itself based on implicature, and can be cancelled in favor of another P-set. He offers the example *I bought a book and talked to the author about it – not the author of the book, I mean the author we were just talking about*, in which the P-set containing entities that have an associative relationship with the previously mentioned book is cancelled in favor of, say, the universe of all entities mutually manifest to the speaker and hearer, within which there is presumably only one author that has just been talked about.

³Although Hawkins' analysis is limited to the opposition between *the* and *a/some*, it is possible that it could be extended to cover these cases as well. Revising the position he took in earlier work (Hawkins 1978), Hawkins suggests that it would be desirable to view demonstratives and pronominals as sharing the same uniqueness requirements as definite descriptions. This would require that the differences between these forms be characterized solely in terms of how "the parameters of uniqueness are set by different (though partially overlapping) conventions of use" (1991:416, fn. 7). Although he does not offer such an analysis, examples (14) and (15) could be captured by uniqueness implicatures if such an analysis were to prove tenable. The same reasoning presumably applies to categorizing indefinite *this* with *a* as not encoding a uniqueness requirement.

The cancellation in example (16) could thus be characterized as one in which the P-set that is extended to accommodate the noun phrase *the book that John is currently reading* – a P-set presumably consisting of all entities that are mutually manifest to the speaker and hearer – is subsequently cancelled in favor of the same P-set *without* the just-accommodated entity. While an argument for such an analysis could conceivably be constructed, example (16) receives a more natural and intuitive explanation by the FA analysis, which provides a clear-cut basis for the cancellation in terms of familiarity.

The second scenario that distinguishes the two proposals includes cases in which two referring expressions differ with respect to uniqueness but not familiarity. Consider (17), a variant of (4):

- (17) I met a student before class. #A student came to see me after class as well – in fact it was the student I met with three days ago.

The attempted cancellation of a nonuniqueness implicature in (17) is bizarre in a context in which the hearer has no knowledge of the speaker having met a student three days ago (let alone which student). (On the other hand, it is perfectly felicitous if the hearer does have such knowledge.) To account for this fact in Hawkins' analysis, the principle that allows for the accommodation of unfamiliar entities into P-sets would have to be suppressed in cancellations of the sort attempted in (17). Yet any argument along these lines would at best only indirectly capture the fact that the felicity of (17) hinges precisely on whether the definite noun phrase *the student I met with three days ago* allows the hearer to identify the referent as one that is familiar to him. These facts therefore suggest that familiarity, rather than uniqueness, is the basis for the cancellations in these examples.

We now compare the predictions of the FA with those made by the GHZ analysis. The Givenness Hierarchy fares well with respect to examples (14) and (15), since each involves a strengthening from a form associated with a weaker cognitive status to one associated with a stronger status.⁴ (We will return to examples (16)–(17) momentarily.) To differentiate between the two analyses, we therefore need to consider alternations between forms that differ with respect to GHZ's cognitive statuses but not with respect to familiarity.

Let us first consider the alternation between TYPE IDENTIFIABLE *a NP* phrases (henceforth '*a*-NPs') and REFERENTIAL indefinite *this NP* phrases (henceforth '*this*-NPs'). Despite being opposed on GHZ's implicational hierarchy, their alternation does not appear to exhibit the hallmark of quantity implicature that we have seen in previous examples. Consider the attempted cancellation in (18):

- (18) A: I'm sick of driving you around everywhere. You should buy a car.
B: #Yeah, I'm going to buy a car today, in fact, this car, and it's in great condition.
[on the REFERENTIAL interpretation of *this car*]

The cancellation of TYPE IDENTIFIABLE (which does not require there to be a specific car that is going to be bought) in favor of REFERENTIAL (which does) seems highly awkward, even though it would be perfectly reasonable for B to confirm A's suggestion as stated and then strengthen it by indicating that a particular car has already been identified.⁵ GHZ address the lack of such implicatures in a footnote:

⁴That is, from TYPE IDENTIFIABLE to FAMILIAR in (14), and from REFERENTIAL to ACTIVATED in (15). Note that GHZ associate accented pronominals such as HIM in example (15) with the status ACTIVATED.

⁵Barbara Abbott (p.c.) points out that such examples become more acceptable if new information is added to the cancelling NP:

Note, however, that the indefinite article does not implicate nonreferentiality. This is so because, with the exception of indefinite *this* in colloquial English, the languages we examined have no separate form that signals referentiality. The indefinite article would thus be the strongest possible form for coding something which is referential but not uniquely identifiable. (1993:296, fn. 26)

We understand this quotation to suggest that such implicatures SHOULD arise between *a*-NPs and *this*-NPs in colloquial English, since both forms are available.⁶ Yet corpora of naturally-occurring English colloquial speech contain many instances of *a*-NPs that are used to refer to entities that satisfy the criteria of the REFERENTIAL status. (GHZ's data support this also; they found that only one of 42 references to entities with this status used a *this*-NP, although their corpus included other sources besides colloquial speech.) Consider the following passage from the Switchboard corpus:

- (19) Yeah, I have a [(throat-clearing) I have a] friend. [She's studying, + she's going into] law, another friend going into law, and, [she's [[[I, + I,] + I'm really,] + I'm really] concerned about, you know, she's just going to be, [you know,] what's going to happen when she gets there.

The referring expression *a friend* in the first sentence meets the criterion for the REFERENTIAL status, and in fact its referent is mentioned again several times using a pronoun. Although a *this*-NP would have been felicitous here, there is no hint of any uncooperativeness or Maxim-flouting on the part of the speaker in failing to use one. Nor would any form of cancellation or reinforcement be felicitous.

A possible response would be to appeal to Q2, as Maclaran (1980, endnote 9) does. Specifically, if the identity of the referent is unimportant, then an explicit signal of referentiality is unnecessary; indeed the predicate that accompanies the referring expression (*have*) in (19) entails that the reference is to a particular person.⁷ But then we need an explanation for why the passage remains felicitous if we were to replace *a friend* with *this friend*, i.e., why this would not violate Q2. As far as we can tell, there are no cases in which a *this*-NP MUST be used instead of an *a*-NP for a reference to be pragmatically felicitous, that is to say, without coming across as

-
- (i) B: Yeah, I'm going to buy a car today, in fact, this blue sportscar that I saw at John's car dealership, and it's in great condition.

There are at least two reasons to suspect that this is not a cancellation of a cognitive status implicature, however. First, the felicity of the putative cancellation is unaffected by the use of an *a*-NP instead of a (REFERENTIAL-signalling) *this*-NP, as in (5):

- (ii) B: Yeah, I'm going to buy a car today, in fact, a blue sportscar that I saw at John's car dealership, and it's in great condition.

Likewise, the example remains felicitous if *going to buy* is replaced with *bought*, in which case the REFERENTIAL status of the car has already been entailed at the time the apparent cancellation occurs:

- (iii) B: Yeah, I actually bought a car today, in fact, a/*this* blue sportscar that I saw at John's car dealership, and it's in great condition.

We therefore conclude that the felicity of the putative cancellation in (i) derives from the additional descriptive content (*not just any car, but a blue sportscar...*), and not the distinction between the two cognitive statuses TYPE IDENTIFIABLE and REFERENTIAL.

⁶Although it is worth noting that the very fact that English is the only language of those GHZ studied that has a form that corresponds specifically to REFERENTIAL, coupled with the fact that it is found only in colloquial speech, might be taken to cast doubt on whether REFERENTIAL actually constitutes a separate cognitive status.

⁷Although note that replacing the predicate in (19) with one that does not entail referentiality – e.g., *I need to call a friend* – does not in turn require that the object be marked with a *this*-NP, even though it remains REFERENTIAL.

uncooperative or generating conversational implicatures. This casts doubt on the notion that these referring expressions are part of a Gricean Q-based scale that supports generalized quantity implicature.⁸

This pattern of ‘missing’ implicatures extends to the relationship between other pairs on the right end of the Givenness Hierarchy, that is, between the TYPE IDENTIFIABLE and UNIQUELY IDENTIFIABLE categories (ex. 17, repeated below as ex. 20) and the REFERENTIAL and UNIQUELY IDENTIFIABLE categories (ex. 21):⁹

- (20) I met a student before class. #A student came to see me after class as well – in fact it was the student I met with three days ago. [= (17)]
- (21) I met a student before class. #This student came to see me after class as well – in fact it was the student I met with three days ago. [on the REFERENTIAL interpretation of *this student*]

Thus, the evidence suggests that pairs of categories on the Givenness Hierarchy that are associated with nonfamiliarity do not participate in the type of conversational implicatures that GHZ’s system would predict.

The same appears to be true of the three statuses on the left end of the hierarchy that are associated with familiarity. The following cases are markedly odd:

- (22) a. The neighbors to our right just got a Doberman like the one the neighbors behind us have. #That dog, in fact it, kept me awake last night. (ACTIVATED _ IN FOCUS)
- b. I don’t like our neighbor’s dog. #That dog that bit you once, in fact it, was barking last night. (FAMILIAR _ IN FOCUS)
- c. Contractors are putting up a fence around the yard of my neighbors with the Doberman Pinscher. #That dog that bit you last year, in fact, that/this dog/that dog, won’t be able to get out of the yard anymore. (FAMILIAR _ ACTIVATED)

In (22a), for instance, *that dog* is semantically compatible with either of the two Dobermans mentioned; in GHZ’s account, the referent should be the dog denoted by *the one the neighbors behind us have* by implicature, since the one denoted by *a Doberman* is presumably IN FOCUS. While this is in fact the preferred referent, an implicature analysis predicts that this reference assignment should be cancellable, but (22a) shows that this is not the case.¹⁰ Similar arguments apply for (22b-c).

⁸While we offer no analysis of *this*-NPs here, a common claim is that they provide a means to establish a hearer-new referent as a discourse topic (Maclaran 1980, Prince 1981a, Wald 1981, Gernsbacher and Shroyer 1989). As such, *a* and *this* are not distinguished by the same kind of properties that differentiate other forms on the Givenness Hierarchy. We return to this point at the end of this section.

⁹Crucially, we use UNIQUELY IDENTIFIABLE here to mean “UNIQUELY IDENTIFIABLE but not FAMILIAR”; in examples (20) and (21), we assume that the referent of the definite description is unknown to the hearer. As we previously noted for (20), these examples become impeccable if the referent is actually familiar.

¹⁰Note that while the entity denoted by *the one the neighbors behind us have* is (at least) ACTIVATED, none of the referring expressions that are associated with that status in the Givenness Hierarchy – i.e., *that*, *this*, or *this dog* – would be preferentially interpreted as referring to that entity. For this reason we use *that dog*, as the form *that*-N is commonly used to refer to ACTIVATED entities despite its being aligned with the weaker FAMILIAR status in the GHZ system. Generally speaking, there appear to be many ACTIVATED entities in natural language corpora that cannot be felicitously referred to with a demonstrative pronoun, e.g. the referent of *the table* in (i):

Finally, we return to example (16), repeated below as (23):

- (23) The book that John is currently reading – in fact the one I showed you yesterday when we were at the bookstore – is reportedly a real page-turner.

Recall that example (23) is felicitous as long as (the speaker believes that) the hearer does not know which book John is currently reading (i.e., the definite merely signals uniqueness), but does know which book the speaker showed him yesterday at the bookstore. Interestingly this cancellation occurs despite the fact that a form that signals only UNIQUELY IDENTIFIABLE is used in BOTH referring expressions. That is, under GHZ's analysis one might expect, since it is precisely the implicature of nonfamiliarity that is being cancelled, that a form correlated with the stronger cognitive status would be chosen for the second referring expression (presumably a *that*-NP, thereby signaling FAMILIAR). However, as we see from (23), an NP signalling FAMILIARITY is not necessary. GHZ's Q2-based explanation for why *the*-NPs are commonly used to refer to FAMILIAR referents – i.e., that the stronger form is informationally otiose – does not work well here, since the very purpose of the cancellation is to indicate the hearer's familiarity with the referent. Thus, the facts concerning cancellation with respect to nonfamiliarity implicatures appear to be orthogonal to the choice of referring expression, although clearly there are examples that crosscut GHZ's hierarchy.

It therefore appears that of the cases we have discussed (excluding 13a-c), only those captured by the FA display evidence for conversational implicatures. All of those cases for which the classic cancellation test for conversational implicature succeeds involve a strengthening from a presumption of nonfamiliarity to one of familiarity.¹¹ We have seen that Hawkins' analysis predicts such implicatures only for a subset of this otherwise uniform set of cases, whereas GHZ's analysis predicts that we would find implicatures in many more situations than we actually do.

We believe that the empirical limitations of these analyses derive from the nature of the semantic scales on which their respective hierarchies are based. To elaborate, the generation of Q-based scalar implicatures requires a single unidimensional scale along which various values on that scale are ordered with respect to each other, where these values, according to Levinson (2000), “must be from the same semantic field, ‘about’ the same semantic relations, and thus in conceptually salient opposition” (p. 80). The classic Horn scales commonly discussed in the literature – *<all, most, many, some>*, *<and, or>*, *<..., 3, 2, 1>*, *<necessarily, possibly>*, and so forth – all share this property. As we have shown, we find that Hawkins' concept of uniqueness, while being a coherent semantic notion on which to base a scale, does not straightforwardly account for the full range of conversational implicatures that we have seen. GHZ's analysis, on the other hand, is more complicated precisely because there appears to be no unifying semantic dimension underlying the Givenness Hierarchy. While GHZ briefly characterize the notion of cognitive status as “information about location in memory and attention state” (1993:274), on

(i) The Doberman_i is sitting next to the table_j. It_i / # It_j / # That_j / The table_j is tall and brown. GHZ's association of demonstrative pronouns with the ACTIVATED category therefore seems to underspecify the constraints on their use.

¹¹To be clear, we do not deny the existence of other properties of referring expressions that might give rise to conversational implicatures (uniqueness likely being among these, as we have already indicated). Indeed, the example of Grice's discussed in the introduction – in which *X is meeting a woman this evening* implicates that the woman is not X's wife, mother, sister, or close platonic friend – does not appear to involve a nonfamiliarity implicature, at least under GHZ's definition of FAMILIAR, which we are assuming here. GHZ (1993:296, fn. 26) claim that this example involves an implicature that the referent is not uniquely identifiable, but (20) and (21) cast doubt on this being the right analysis. Further discussion would require a detailed analysis of the properties of possessive NPs (which are not included in the Givenness Hierarchy), which would take us too far afield.

closer inspection it is difficult to identify a single, consistent semantic field on which the statuses can be aligned. On the left end of the hierarchy, the operative notion relates to the level of ACTIVATION (Chafe 1976, Lambrecht 1994) of the referent in the cognitive state of the hearer, that is, the distinction between whether the referent is merely FAMILIAR to the hearer (i.e., in either long- or short-term memory), ACTIVATED (in short-term memory) or IN FOCUS (i.e., at the current center of attention). All of these cognitive statuses are equivalent with respect to hearer-knowledge, i.e., all are HEARER-OLD in the sense of Prince (1992). Moving to the right end of the hierarchy, however, degree of hearer activation cannot be what distinguishes entities that are TYPE IDENTIFIABLE from those that are REFERENTIAL: Since both correspond to entities that the speaker assumes are new to the hearer (i.e., HEARER-NEW in Prince's (1992) terms), they cannot differ with respect to activation in the hearer's cognitive state at the time that the referring expression is uttered. Likewise for those referents that are UNIQUELY IDENTIFIABLE (but not FAMILIAR): As hearer-new entities, these are distinguished from those that are only REFERENTIAL in that the addressee is expected to construct a (unique) representation of the referent on the basis of the referring expression alone. It is hard to imagine a single uniform dimension that can relate these different statuses such that they would be expected to generate Quantity-based implicatures in a systematic way. It is therefore not surprising that we do not find evidence in the data for the type of broad-based system of implicatures that the GHZ spectrum would predict.

4. Conclusion

To summarize, previous work by Hawkins and GHZ has appealed to examples such as (2) and (3) to support their respective implicature-based theories of referential form. A further examination of the data using the classic diagnostics for conversational Quantity-based scalar implicature, however, has revealed that these implicatures do not derive from a scale of uniqueness (à la Hawkins) nor cognitive status (à la GHZ). Instead, we have argued that the pattern reflects the existence of nonfamiliarity implicatures, in which a speaker's failure to use a referring expression that indicates hearer-familiarity conversationally implicates that the referent is in fact nonfamiliar to the hearer. Nonfamiliarity implicatures can thus be seen as a subclass of those implicatures that result when a "speaker has failed to be specific in a way in which he might have been expected to be specific" (Grice 1975:57).

We therefore find the evidence for entailment scales in the English referential system to be limited in scope. Of course, this is not to say that familiarity is in any sense the dominant factor that determines choice of referential form. Indeed, we suspect that choice of referential form is governed by a system of heterogeneous factors across many dimensions. From this brief discussion of referring expressions and Quantity-based implicature, it should be clear that no single dimension – whether it be uniqueness, givenness, or familiarity – is adequate in and of itself, and that a complete account of the relationship that holds among the various referring expressions of a language will have to take additional factors into account.

References

- Atlas, J.D. and S.C. Levinson. 1981. *It*-clefts, informativeness, and logical form: Radical pragmatics (revised standard version). In P. Cole, ed., *Radical Pragmatics*. New York: Academic Press. 1–61.
- Chafe, W.L. 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view.

- In C.N. Li, ed., *Subject and Topic*. New York: Academic Press. 25–55.
- Gernsbacher, M.A. and S. Shroyer. 1989. The cataphoric use of the indefinite *this* in spoken narratives. *Memory and Cognition* 17:536–540.
- Grice, H.P. 1975. Logic and conversation. In P. Cole and J. Morgan, eds., *Syntax and Semantics, Volume 3: Speech Acts*. New York: Academic Press. 41–58.
- Gundel, J.K., N. Hedberg, and R. Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language* 69:274–307.
- Hawkins, J.A. 1978. *Definiteness and Indefiniteness*. Atlantic Highlands, NJ: Humanities Press.
- Hawkins, J.A. 1991. On (in)definite articles: Implicatures and (un)grammaticality prediction. *Journal of Linguistics* 27:405–442.
- Hirschberg, J. 1991. *A Theory of Scalar Implicature*. New York: Garland.
- Horn, L.R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In D. Schiffrin, ed., *Meaning, Form, and Use in Context*. Washington, DC: Georgetown University Press. 11–42.
- Lambrecht, K. 1994. *Information Structure and Sentence Form*. Cambridge: Cambridge University Press.
- Levinson, S.C. 1987. Pragmatics and the grammar of anaphora: A partial pragmatic reduction of binding and control phenomena. *Journal of Linguistics* 23:379–434.
- Levinson, S.C. 2000. *Presumptive Meanings*. Cambridge, MA: MIT Press.
- Maclaran, R. 1980. On two asymmetrical uses of the demonstrative determiners in English. *Linguistics* 18:803–820.
- Prince, E.F. 1981a. On the inferencing of indefinite-*this* NPs. In A.K. Joshi, B.L. Webber, and I.A. Sag, eds., *Elements of Discourse Understanding*. Cambridge: Cambridge University Press. 231–250.
- Prince, E.F. 1981b. Toward a taxonomy of given-new information. In P. Cole, ed., *Radical Pragmatics*. New York: Academic Press. 223–255.
- Prince, E.F. 1992. The ZPG letter: Subjects, definiteness, and information-status. In S. Thompson and W. Mann, eds., *Discourse Description: Diverse Analyses of a Fundraising Text*. Philadelphia/Amsterdam: John Benjamins. 295–325.
- Sadock, J.M. 1978. On testing for conversational implicature. In P. Cole, ed., *Syntax and Semantics, Volume 9: Pragmatics*. New York: Academic Press. 281–297.
- Wald, B. 1981. Referents and topic within and across discourse units: Observations from current vernacular english. In F. Klein-Andreu, ed., *Discourse Perspectives on Syntax*. New York: Academic Press. 91–116.

