# Discourse model representation of referential and attributive descriptions

Kristine H. Onishi and Gregory L. Murphy *University of Illinois at Urbana-Champaign, IL, USA* 

Definite descriptions (as in *The murderer of Smith is insane*) can have at least two interpretations: a *referential* one, in which insanity is predicated of a particular individual who killed Smith, and an *attributive* one, in which insanity is predicated of whoever it is that killed Smith. Experiment 1 manipulated shared knowledge and focus on specific entities, the verb in the sentence, and whether the description was definite or indefinite. Each factor influenced interpretation of the description. Experiment 2 confirmed that changing the verbs alone affected reference choice. Experiments 3 (ratings) and 4 (reading times) indicated that both referentially and attributively introduced entities are conceptually singular (better as antecedents of singular than plural pronouns) while generically introduced entities are conceptually plural. Thus, the difference between the discourse representation underlying referential and attributive interpretations does not hinge on a difference in the number of tokens being instantiated.

### INTRODUCTION

When communicating, people often need to talk about the world around them. Specifically, they often mention things and events, and to do so they use expressions that refer to these entities. One particular type of referring expression is a definite description (a description containing *the X*). These

Requests for reprints should be addressed to Kristine H. Onishi, Department of Psychology, University of Illinois, 603 East Daniel St., Champaign, IL 61820, USA. Email: konishi@s.psych.uiuc.edu

This research was supported by NIMH grants MH41704 and 1 T32MH19990-01. Thanks to Amanda Seidl and Ellen Prince for helping to get this project started and to Kathryn Bock, Susan Garnsey, and MurphyLab for comments, suggestions, and advice. We would also like to thank two anonymous reviewers for helpful comments on previous versions of this paper.

expressions can have at least two different interpretations (Donnellan, 1966, 1978; Searle, 1979). For example, a sentence containing a definite description (as in 1a; adapted from Donnellan, 1966, see Roberts, 1993):

1a. The murderer of Smith is insane.

can be interpreted *referentially* (Donnellan, 1966) when it is understood to mean something like (1b).

1b. I think the particular person just convicted of Smith's murder is crazy.

The same sentence has an *attributive* interpretation when its understood meaning is more as in (1c).

1c. I think anyone who would have killed Smith must be crazy.

The distinction between referential and attributive descriptions was originally drawn to criticise Russell's (1905) account of definite descriptions. Roughly speaking, Russell believed that the description *the man drinking a martini* in a sentence (2; adapted from Donnellan, 1966):

2. The man drinking a martini is going to be ill.

meant that: (1) there is exactly one man drinking a martini and (2) anything that is the man drinking a martini is going to become ill. (Of course, there may be more than one man drinking a martini in the world, so *the* must be interpreted relative to a discourse context.) Russell's analysis describes the attributive interpretation, because the person who will be ill is identified by virtue of the description (martini drinking), but it fails to account for referential uses.<sup>1</sup>

There has been some confusion as to whether Donnellan was arguing that the difference between referential and attributive uses was semantic or pragmatic. On a semantic view, referential and attributive uses express different propositions. On a pragmatic view, both uses express the same proposition, but differ in their conversational implicatures (Recanati, 1993, ch. 15). Whether the difference is semantically based or pragmatically

<sup>&</sup>lt;sup>1</sup> The difference between referential and attributive interpretations can be seen most clearly when the description turns out to be wrong (Donnellan, 1966). For example, if a sentence with a definite description (2 in text) is interpreted attributively and it turns out that no one fits the description, no new information is added to the discussion. If understood attributively, the addressee takes the description as essentially a definition, picking out whoever happens to fit it, and saying nothing when nothing fits. In contrast, if the sentence is interpreted referentially, it does not matter if no one fits the description—it still applies to the particular person under discussion. If understood referentially, the addressee selects a particular referent using the description as a tool to find it.

based, addressees may interpret attributive and referential uses differently. As seen in examples (1b) and (1c), speakers could intend very different things with the same description (e.g., Clark, 1979; Gibbs, 1984), so the present study will examine how addressees interpret speaker meaning, rather than sentence meaning. Specifically, it will examine the questions of (1) what cues addressees might use to determine which interpretation speakers intend by the use of a description, and (2) how each of these interpretations might be instantiated in a discourse representation of the text.

# Cues to referential and attributive interpretations

The first major goal of this paper is to investigate whether interpreting a description differently has consequences for judgements about referents and to examine what sorts of cues addressees can use to determine the speaker's intent. Since a single sentence (as 1a) can be interpreted either referentially or attributively, what information does the addressee use to choose the appropriate interpretation? A variety of cues may be available.<sup>2</sup>

Donnellan (1966) hints at one potential cue by pointing out that one of the differences between referential and attributive descriptions is that they have different presuppositions. According to him, a referential use (e.g., of 1a) presupposes that (1) something fits the description (Smith has a murderer), and (2) there is a particular something that fits the description (Jones is that murderer). An attributive use only has the first presupposition. This difference in the presuppositional requirements for referential and attributive interpretations suggests that the easier it is to come up with something that satisfies the second presupposition (that a particular something fits the description) the easier a referential interpretation would be. That is, the more emphasis there is on individuals within a context, the easier it would be to interpret the entire sentence referentially.

<sup>&</sup>lt;sup>2</sup> Some suggested cues have relied on specific lexical items. For example, it has been claimed that if *whoever* can be inserted felicitously (3a), the description should be interpreted attributively (Kripke, 1979; Roberts, 1993).

<sup>3</sup>a. The man who broke into the car (whoever he is) is mad. However, it is possible for a sentence to contain *whoever* (3b) and still be interpreted referentially (Searle, 1979), so the *whoever* test cannot be definitive.

<sup>3</sup>b. That man I see there, breaking into the car, whoever he is, is mad.

Another suggested cue was the presence of demonstrative *that*—which would indicate that the description should be interpreted referentially (3b; Roberts, 1993). A general problem with such cues is that they are often not present (e.g., 1a), leaving open the question of what indicates the intended interpretation in their absence. The effect of these overt lexical cues will not be examined in this paper.

The few psychological studies that have examined the referential/ attributive distinction offer hints about additional cues.<sup>3</sup> Another potential cue may be familiarity or fame (Boland & Dell, 1991; Johnson-Laird & Garnham, 1980; Mueller, 1988; Mueller-Lust & Gibbs, 1991), though the direction of its effect is unclear. From the addressee's point of view, attributive interpretations should be easier when the description is about a role that is already established in semantic memory (e.g., the pope, the president), but more difficult when the role in the description is not as familiar (e.g. Urbana High School Most Valuable Player). This is because expressions of reference can point to roles (attributive interpretations) or to specific tokens (referential interpretations; Boland & Dell, 1991). From the speaker's point of view, it has been suggested that more knowledge makes a referential interpretation easier. Johnson-Laird and Garnham (1980) have suggested that a sentence can only be used referentially when the speaker has enough knowledge to talk about the referred-to individual in at least two independent ways (e.g., with a proper name and also with a description or with two independent descriptions). So, the more the speaker knows about the referent, the easier referential interpretations should be, whereas the familiarity argument made above predicts that the more the addressee knows the easier an attributive interpretation would be.

A third potential cue to the referential/attributive distinction may be the type of predicate used (Ortony & Anderson, 1977; see also Boland & Dell, 1991). For example, sentences with predicates as in (4a) were assumed (with no explanation) to be referential while sentences as in (4b) were assumed to be attributive (Ortony & Anderson, 1977).

- 4a. Nobel/The inventor of dynamite had a fine beard.
- 4b. Nobel/The inventor of dynamite profoundly influenced warfare.

Although predicate type may be important, it cannot be the only cue to referential or attributive interpretation since a single sentence (e.g., 1a)

<sup>&</sup>lt;sup>3</sup> Ackerman (1979) claimed to have investigated two cues that might lead people to make attributive rather than referential interpretations, however his distinction does not seem to correspond to the one suggested by Donnellan (1966). Ackerman's distinction depended on the acceptability of erroneous descriptions. For example, if a character was looking for a book with "a picture of a cat" on the cover and the subject was NOT willing to accept a book with a picture of a dog, an attributive interpretation was coded. On the other hand, accepting the dog book indicated a referential interpretation, since the subject acted as if the description could be an imprecise fit, but still refer to the object (Ackerman, 1979, p. 5), which is permissible only for referential uses. However, Ackerman's distinction does not seem to hinge on a difference in interpretation of the noun phrase, but on how willing subjects were to accept a mistaken description. In all cases the characters in the scenarios were searching for a particular entity (e.g., the missing math book, rather than any book that happened to fit the description), so all his cases seem to be referential.

can be interpreted either way. Thus several mechanisms have been hinted at as potential cues to determining whether a sentence was used referentially or attributively. To date, none has been directly tested.

Two aspects of the referential/attributive distinction are to be addressed in this paper. Experiments 1 and 2 will examine some of the potential cues to the referential/attributive distinction that have been suggested. First, Experiment 1 used a forced choice paradigm to distinguish between referential and attributive interpretations. It manipulated the emphasis on particular individuals (as suggested by Donnellan, 1966) as well as the amount of knowledge the speaker is mutually known to have (as suggested by Johnson-Laird & Garnham, 1980). A subsidiary question addressed in Experiment 1 is the similarity of definite and indefinite descriptions. Donnellan considered only definite descriptions (Cole, 1978) and some have argued that the distinction does not hold for indefinite (a/n X)descriptions (Kripke, 1979; Russell, 1905). However, others have argued for a functionally equivalent distinction (Roberts, 1993). Experiment 1 will examine whether the interpretation of both definite and indefinite descriptions are affected by the same sort of cues. Finally, verbs will be manipulated in both Experiment 1 and in Experiment 2 (which used a different task) to examine the effect of predicate on referential/attributive interpretations (as suggested by Ortony & Anderson, 1977).

Experiments 3 and 4 investigated the discourse representations of referential and attributive interpretations, specifically attempting to determine whether attributively introduced noun phrases (NPs) are interpreted as conceptually singular or plural.

# **EXPERIMENT 1**

Mueller-Lust demonstrated that readers are sensitive to contexts that are referentially or attributively biased (Mueller, 1988), but it is not clear what factors give listeners cues as to speaker intent, since she did not vary her contexts such that the influence of individual factors could be distinguished. Also, there is no direct psychological evidence that different interpretations have implications for what the referent of the description is taken to be.

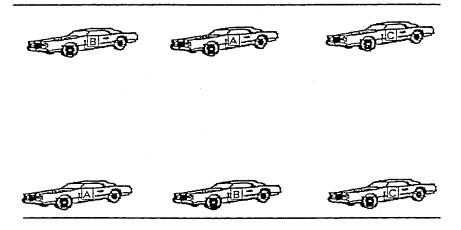
To address the question of whether the referential/attributive distinction has consequences for reference, subjects in Experiment 1 listened to scenarios and were asked to choose the referent they thought the characters in the scenario were talking about. To investigate what factors affect interpretation, the scenarios were systematically varied on several dimensions, including context and type of verb. Definiteness of description was also varied to examine whether definite and indefinite descriptions are

influenced by the same factors, which might suggest that the referential/attributive distinction is found in both.

### Method

Subjects. Twenty-seven subjects were recruited from the University of Illinois Introductory Psychology subject pool or from the university community and received class credit or cash remuneration, respectively, to participate in this and other unrelated studies. All subjects in this and other experiments reported speaking English as their first language. Data from two subjects were not included because they made the same choice for every scenario (attributive) and from one subject who did not choose either the referential or attributive option in 9 of the 28 scenarios. Upon debriefing, she reported basing her responses on an irrelevant visual quality, not on the presented scenarios.

Materials. Twenty-four target sentences were embedded in short scenarios about races in which there were three racers and two viewers watching the race. The target sentence was uttered by viewer 1 to viewer 2. Each race scenario was depicted by two schematic pictures of races involving either three foxes or three cars (Figure 1). One picture showed the race early on (Early picture), the other showed the same race at a later time (Later picture). Depicting the race at two different times allowed for changes in the interpretations of the target description (as described



**Figure 1.** Examples of car pictures used in Experiment 1. The first row represents the race when the description was heard (Early picture) and the second row represents the race at the time of the question (Later picture). Subjects were told that racers were moving to the left, so in the Early picture, car B is winning.

below). Each Early picture was accompanied by a description, and each Later picture by a question to the subject (see Table 1 for example scenarios). The descriptions with the Early picture consisted of context sentences followed by the target sentence.

The three racers were labelled A, B, and C. Subjects were told that letters served to identify individual racers. In the Early picture, one racer fit the description (winning), while the other two racers did not. In the Later picture, the racers were always shown in a different configuration, so that a different racer was winning. The order of the two losing racers was counterbalanced.

For the descriptions that accompanied the Early picture, there were three possible contexts: Individual, Property, and Neutral. These contexts preceding the target description were varied in two ways. First, they differed in their emphasis on the importance of the individual entity. Increasing the emphasis on specific entities should increase referential interpretations (Donnellan, 1966). Second, they differed in the amount of shared knowledge between characters in scenarios. Johnson-Laird and Garnham's (1980) discussion suggests that shared knowledge may increase referential interpretations, because addressees would be able to determine that speakers had enough knowledge to be referential in their intent. These two factors were combined in the hopes of finding any effect of context.

The Individual context was expected to bias the subject towards a referential reading of the target sentence by stressing the importance of a particular racer for the viewers and indicating that the viewers had a high degree of shared knowledge as well as the ability to distinguish individual racers. The Property context emphasised the attributes of some racer rather than a particular racer, while suggesting that the viewers had an intermediate degree of shared knowledge and that they might not be able to differentiate particular racers. The Neutral context indicated that there was little shared knowledge between the two viewers, while giving no indication as to whether particular racers or attributes of racers were more important to the speaker (see Table 1).

Another variable that might affect interpretation is the type of predicate (Ortony & Anderson, 1977). Though not explicitly suggested by previous studies, one intuitive possibility is that the degree of perceived stability of a property may affect referential/attributive interpretations. Target sentences included one of two verbs, own or like, that differed in how likely it was that their object would change over the course of the presented scenario. Own (a permanent verb) was chosen because, for the duration of the race, what was owned was unlikely to change, and so the use of own should increase referential interpretations of the target sentence. The object of like (a temporary verb) could more plausibly change during the

TABLE 1 Example scenarios from Experiment 1

	Own	Like			
Individual Context	[show the Early picture] Jane and Rich went to the races together all the time and could recognise the cars individually. They were excited because Jane had just bought a car and was having it raced for the first time. It was the middle of a close race and Jane shouted,	[show the Early picture] Devorah and Helen went to the races together all the time and cou recognise the foxes individually. They were excited because Devora was planning to buy a fox and had decide which one she liked most. I was the middle of a close race and Devorah shouted.			
Target question	"I own the/a car that is winning." [show the Later picture] This picture is later in the race. Which car do you think Jane owns?	"I like the/a fox that is winning." [show the Later picture] This picture is later in the race. Which fox do you think Devorah likes?			
Property Context	[show the Early picture] Tim and Ann were good friends, but neither had ever been to the races before. Tim had just won a fox in a lottery and was trying to guess which one was his. By the afternoon, they had figured it out. Tim said,	[show the Early picture] Nicole and Andy were good friends, but neither had ever been to the races before. They watched several races carefully to decide what made a good racing car. By the afternoon, they had figured it out. Nicole said,			
Target question	"I own the/a fox that is winning." [show the Later picture] This picture is later in the race. Which fox do you think Tim owns?	"I like the/a car that is winning." [show the Later picture] This picture is later in the race. Which car do you think Nicole likes?			
Neutral Context Target question	[show the Early picture] Fred enjoyed going to the car races by himself every week. He had just come back from the concession stand in the middle of a race. The man sitting next to him turned and said out of the blue, "I own the/a car that is winning." [show the Later picture]	[show the Early picture] Nadia enjoyed going to the fox races by herself every week. She had just come back from the concession stand in the middle of a race. The man sitting next to her turned and said out of the blue, "I like the/a fox that is winning." [show the Later picture]			
	This picture is later in the race. Which car do you think the man owns?	This picture is later in the race. Which fox do you think the man likes?			
Property biased filler	[show the Early picture] Knut and Nancy had only been to the fox races together once before. Even though they didn't bet any money, they were having a good time just cheering for different foxes. Knut said,				
Target question	"Let's cheer for whichever fox is winning." [show the Later picture] This picture is later in the race. Which fox do you think Knut is cheering for				

<sup>\*</sup> All scenarios were presented once with foxes and once with cars to each subject.

race. For example, if I like whichever fox is winning, early in the race I may like fox A, but later in the race, fox B. The use of *like* should either bias towards an attributive reading or leave the interpretation more open to influence by other factors.

Finally to examine whether similar factors affected both definite and indefinite descriptions, target sentences also used one of two articles, the or a. The (definite) should have a tendency to lead to a referential interpretation of the target sentence since the "definite article indicates that in some sense a particular reference is being made" (Donnellan, 1978, p. 56). That is, because the is often used to refer to a uniquely identifiable referent in a particular context, it may bias towards a referential interpretation. A (indefinite) should be more ambiguous, since it tends to be used when this requirement is not met (Strawson, 1950) or may even have a slightly attributive bias by implicature (i.e., by not using the, a nonreferential interpretation is implied; Grice, 1975). It is also possible that indefinites will not be interpreted referentially at all.

These three factors led to a context (3) × verb (2) × article (2) withinsubject design. Each subject heard two instances of each of these 12 scenario types, once as a fox race and once as a car race. The 24 scenarios were presented in a single blocked random order or its reverse, to which subjects were assigned randomly. Each subject received one occurrence of each of the 12 conditions before receiving the second instance of any condition. In addition, there were four filler scenarios (presented in positions 4, 11, 18, and 25) that were strongly biased towards a Property (attributive) reading (see Table 1). These were included because a pilot study found an overall referential bias in this task, and it was hoped that the fillers would make attributive responses seem acceptable. There were four practice items before the experiment to ensure that the subject understood the procedure.

Procedure and scoring. The Early picture was shown as the experimenter read the context and the target sentence aloud. Then the Later picture was shown, and subjects were asked which racer they thought the viewers were talking about. Choosing the racer that had fit the description in the Early picture but no longer fit the description at the time of the question was coded as a referential response, e.g., choosing car B in the Later picture of Figure 1 after hearing the referential context scenario about Jane and Rich (from Table 1). Choosing the racer that fit the description in the Later picture, but that had not fit in the Early picture was coded as an attributive response, because this suggested that the subject had not interpreted the description as applying to a specific entity but to whichever entity was fitting the description (i.e., choosing car A in the same example). Choosing the racer that never fit the description indicated

that the subject had not understood what the viewers were talking about and was coded as an error (i.e., choosing car C).

Instructions emphasised that each scenario was to be responded to on its own and that subjects should not feel that it was necessary to answer the questions consistently across scenarios. Subjects were told that there were no right or wrong answers, and no feedback was given during the experiment.

## Results

Analyses were conducted on the proportion of referential responses out of total correct responses, in each of the 12 conditions. Across the experiment, there were only three errors, each in a different condition.

There was no reliable main effect of presentation order (forward, backward), F(1, 22) = 2.06, p > .1, nor interactions with the order factor, Fs < 1.8, ps > .15 for a mixed within (context, verb, article) and between (order) analysis of variance (ANOVA), so results were collapsed over presentation order in the reported analyses. Reliability of all results reported in this article is with respect to an alpha level of .05 unless otherwise noted. Alpha levels for *t*-tests have been adjusted by the Bonferroni method to keep the alpha for the set of comparisons to .05 within each ANOVA. Note that because only two verbs were used, an item analysis was not possible.

The three experimental manipulations had effects in the expected directions (Figure 2). There was a main effect of context, F(2, 46) = 3.66,

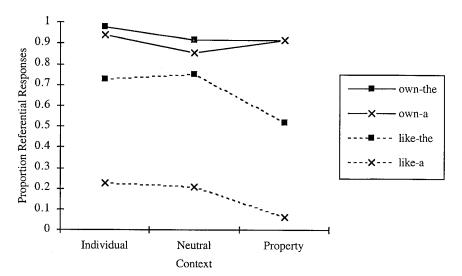


Figure 2. Experiment 1: Mean proportion of referential responses across subjects.

with a greater proportion of referential responses with the Individual context, M = .72, than with the Property context, M = .60; t(23) = 2.65. The Neutral context led to an intermediate proportion of referential responses, M = .68, which did not differ reliably from the Individual context, t(23) = .72, though there was a trend for a difference from the Property context, t(23) = 1.83, p < .08. There was a greater proportion of referential responses with the article the, M = .80, than with a, M = .54; F(1, 23) = 56.84, and many more referential responses with own, M = .92, than with like, M = .42; F(1, 23) = 242.88.

The predicate or verb had the largest effect. Verb was involved in two-way interactions with context and with article [context  $\times$  verb: F(2, 46) = 4.40; article  $\times$  verb: F(1, 23) = 67.26]. The remaining two-way interaction [context  $\times$  article: F(2, 46) = 1.06] and the 3-way interaction [F(2, 46) = 0.02] were not reliable.

Since there was an interaction with verb, separate analyses were conducted for each verb. When the verb was permanent (own), responses were nearly always referential (85-98% across conditions—see Figure 2) and the effects of other factors were eliminated (no main effects or interactions of context and article, all Fs < 1.72). However, if the verb was temporary (like), effects of both context and article were seen [reliable main effects in context  $\times$  article ANOVA, but no interaction; F(2, 46) = .40]. There were fewer referential responses with Property contexts, M = .29, than Individual or Neutral contexts, Ms = .48; t(23) = 2.57, t(23) = 2.58 respectively, and more referential choices with the, M = .67, than a, M = .17; F(1, 23) = 73.06.

A secondary question was whether a similar pattern of results would hold for both definite and indefinite descriptions. Both types of description were interpreted both referentially and attributively, so it is not the case that only definite descriptions can be understood referentially. Ignoring the ceiling effect with own (discussed above), the lack of interactions with article, Fs < 1, suggested that both context and verb exerted similar influences on both definite (the) and indefinite (a) descriptions, though as expected, indefinite descriptions had an overall lower proportion of referential responses.

# Discussion

One goal of the present was to demonstrate that the referential/attributive distinction has consequences for comprehension, and more specifically, for determining reference, and to discover whether a similar distinction applies to both definite and indefinite descriptions. Another goal was to examine some of the particular factors that might influence interpretations. Experiment 1 showed that the same definite description could be

interpreted in two different ways, leading to different choices for the understood referent. The choice of referred-to entity (and hence interpretation as referential or attributive) was influenced by both verb and context (shared knowledge and emphasis on individual). In particular, if the verb was *own*, a referential interpretation was strongly preferred. That is, when the verb indicated permanent properties of the entity, people tended to interpret the description as referring to a particular entity, regardless of the other variables. With *like*, this was not necessarily the case. Within the *like* condition, Property contexts (emphasis on qualities; intermediate level of shared knowledge) led to the fewest referential responses. So, when the verb indicated temporary properties, people were not constrained to interpret the description as referring to a particular entity. Instead, choice of referent depended more on surrounding context.

Also, consistent effects of context were found for both definite and indefinite descriptions, though as predicted by Donnellan, there were more referential responses overall with definites. The effect of verb was less clear. When the verb was *like*, there were more referential responses with *the* than with *a*. This difference was not seen with *own*, probably due to a ceiling effect. The similar effects of context on definite and indefinite descriptions, plus the fact that both descriptions were interpreted referentially and attributively suggests that the referential/attributive distinction may be usefully applied to both types of descriptions.

The primary goal of Experiment 1 was to examine some factors that listeners might use to determine whether the speaker was using a description referentially or attributively. Previous work showed that people can distinguish referential and attributive descriptions but did not systematically manipulate factors that may be providing addressees with cues. Using a forced-choice task, the present experiment provides evidence that referential versus attributive interpretation affects referent choice and that the manipulations that influence referential/attributive interpretations in definite descriptions also affect indefinite descriptions in a similar manner. So, as Donnellan (1966) suggested, to correctly interpret descriptions we must know if they were intended referentially or attributively. The experiment also demonstrated the influence of at least three factors that serve as cues to addressees: context (shared knowledge, emphasis on individuals), verb, and article.

### **EXPERIMENT 2**

Experiment 1 demonstrated that interpretation of the referent of a description depended in part on the verb. When the verb was *own*, people tended to interpret the sentence referentially, but when the verb was *like*, there were fewer referential interpretations. Experiment 2 used a wider

variety of verbs in order to confirm that referential versus attributive interpretations can be influenced by the particular predicate involved.

#### Method

Subjects. Thirty-two subjects were recruited from the Introductory Psychology subject pool and received class credit to participate in this and other unrelated studies.

Materials. There were 12 paragraph-long stories, each containing a target sentence that could be interpreted as being used referentially or attributively in the story. Each paragraph was followed by a question about which of two entities the subject thought a character in the story was referring to. One choice indicated that the target sentence had been interpreted referentially; the other was an attributive choice. Target sentences always contained a definite description, since similar patterns were seen for definites and indefinites in Experiment 1. Each sentence was constructed in two versions, differing only in the verb. One verb should make a referential choice more likely, while the other should make the attributive choice more likely.

For example, the target sentence, "My talk [verbs] the campaign of the most environmentally sensitive candidate," could contain the verb *helps* (as shown below with part of the story context):

At the Iowa caucuses, Elizabeth was campaigning on behalf of the environment... She had developed some of her recommendations with the help of a staff person on the Kim campaign. After the talk, she told a reporter, "My talk *helps* the campaign of the most environmentally sensitive candidate." The reporter pointed out to her, however, that Althea Kim had just received large contributions from construction companies ... Instead, a different candidate, Emily Garcia, had made proposals that were quite similar to what she had suggested in her talk.

Alternatively, the target sentence could contain the verb is based on. The questions assessing the subjects' interpretations as referential or attributive were, "Whose campaign was Elizabeth's talk helping/based on?" The choice indicating a referential interpretation of the target sentence would be Althea Kim (i.e., sticking to the same referent), whereas choosing Emily Garcia (i.e., switching referents) would indicate an attributive interpretation. The answer to this question would reveal whether the description (e.g., "the campaign of the most environmentally sound candidate") was interpreted referentially or attributively. The order of the choices was counterbalanced across booklets.

The only difference between the two versions of the paragraphs was the verb in the target sentence. Unlike Experiment 1, the verbs did not vary on

the permanent/temporary continuum. Instead, within the context of each story, verbs that were expected to lead to more attributive choices helped link the mentioned property to the referent, while verbs expected to lead to more referential choices did not encourage this link within that context. For example, it is reasonable to expect that an environmentally informed speech would cause the environmentally sound candidate's campaign to be helped. So, the verb *helps* (as in 5a) links what is being predicated (being the most environmentally sound) to the referent (the candidate), and an attributive interpretation is predicted.

5a. My talk helps the campaign of the most environmentally sensitive candidate.

On the other hand, that a speech is environmentally informed does not cause it to be based on the campaign of the environmentally sound candidate. So, the verb *is based on* (in 5b) does not link the predicate to the referent as well, and a referential interpretation is predicted.

5b. My talk is based on the campaign of the most environmentally sensitive candidate.

The two versions of each story were put into separate lists. Order of the two potential answers to the question was counterbalanced across stories. Subjects were randomly assigned to one of the two lists, so each subject was presented with one version of each of the 12 test stories as well as a single practice story. In summary, only the verb used in the critical sentence was manipulated. Even though this did not alter the description itself, it might still influence its interpretation. Using multiple verbs and scenarios will allow greater generalisability than did Experiment 1.

*Procedure.* Each of the two lists of 12 stories was put into 8 different random orders and the reverse of those orders, resulting in 32 booklets of 12 stories to which subjects were assigned randomly. Subjects were told they were participating in a study on reading in which they should read each story carefully and answer a question about it. They were told that the questions did not have right or wrong answers and that there was no time constraint on the task. They were asked not to turn back to previous stories after they had moved on. The task took about 10–15 minutes.

# Results and discussion

Analyses were conducted on the proportion of referential responses for the linking and nonlinking verbs separately. Across the 12 stories and 32 subjects, there was a single missing data point because one subject did not respond to one of the questions. Analyses were conducted with both subjects  $(F_1)$  and items  $(F_2)$  as the random factor.

There was no reliable main effect of order of presentation of the two choices,  $F_1(1, 31) = 1.34$ ;  $F_2(1, 11) = 1.09$ , nor interactions with the order factor,  $F_1(1, 31) = 1.34$ ;  $F_2(1, 11) = 1.05$ , all ps > .1, so results were collapsed over choice order in the reported analyses.

The type of verb did have an effect on the interpretation of the target sentence. There was a greater proportion of referential choices with the nonlinking verbs, M = .67, than with the linking verbs, M = .16;  $F_1(1, 31) = 161.71$ ;  $F_2(1, 11) = 64.23$ . For all 12 verb pairs, the nonlinking verb led to a greater proportion of referential responses than did the linking verb. Nonlinking verbs led to more referential interpretations for 31 of the 32 subjects (one subject had equal numbers of referential choices with each verb type). In short, within a particular context, the verb used had a very robust effect on the interpretation of the description.

The primary goal of Experiment 2 was to demonstrate the verb effects found in the first experiment were not specific to the verbs *like* and *own*. This was demonstrated to be the case. Interpretation of a sentence as referential or attributive can change based only on the verb. Verbs that made the link between the description and the referent clearer led to fewer referential interpretations than verbs that did not make this connection as clear.

### **EXPERIMENT 3**

# Discourse representation of referential and attributive uses

Experiments 1 and 2 demonstrated that subjects can distinguish between descriptions intended referentially and attributively, and that interpretation as referential or attributive can be affected by differences in context, verb, and article. But what is the discourse representation like when an attributive reference is understood? This question will be addressed in Experiments 3 and 4. Specifically, the second major goal of this paper is to investigate a specific hypothesis about a distinguishing feature of referential and attributive interpretations—whether they differ in that an attributive use of a description points to a set in the discourse representation, while a referential use points to a single entity.

A series of studies conducted by Mueller-Lust (Mueller, 1988; Mueller-Lust & Gibbs, 1991) examined the discourse representations of referential and attributive interpretations. She claimed that both uses reinstated their antecedents in a cross-modal priming paradigm, i.e., interpretation in either way tended to increase the accessibility of coreferential antecedents and that therefore both referential and attributive uses must refer. She then suggested that the difference between referential and attributive uses was that referential uses referred to a token, while attributive uses referred

to types or "to a general class of objects or persons and not a specific token" (Mueller, 1988, p. 70). Mueller-Lust could be understood as saying that attributive uses point to a set—the set of entities that could fit the description, or possibly the set of entities that fit the description over time, while referential uses point to specific single entities in the discourse model.

There are several potential concerns with Mueller-Lust's claims however. The first is that the evidence supporting her claim for reinstatement is somewhat unclear. For example, with attributive contexts there was not always a difference in decision times to probes between attributive introductions of entities and the baseline conditions (Mueller, 1988, Table 5, p. 41), which would seem to indicate no reinstatement. Also, when probe reaction times were reliably different, there appears to be the possibility of a speed/accuracy tradeoff (Mueller-Lust & Gibbs, 1991, Table 5, p. 120).

The second is that, contrary to Mueller-Lust, many writers have claimed that an attributive use points to a single, unique entity or token (Birner, 1991; Donnellan, 1966) or role (Boland & Dell, 1991), even if the particular individual is not specified (Roberts, 1993; Strawson, 1950). In fact, Webber (1979) states, "I do not feel it necessary to distinguish between 'attributive' and 'referential' definite noun phrases in order to derive appropriate IDs [identifications] for discourse entities and possible antecedents for definite pronouns" (p. 2-14). Her argument was that a sentence used attributively (6a) can be followed felicitously by a sentence with a definite pronoun (he) that refers to a unique discourse entity (6b). Hence even attributive uses put a token into the discourse model (also see Birner, 1991).

- 6a. The murderer of Smith, whoever he is, is insane.
- 6b. He ought to be locked up.

Third, it is difficult to say precisely whether Mueller-Lust's suggestion is compatible with existing models of discourse representation. In the Focus Memory model of discourse representation (Garrod, Freudenthal, & Boyle, 1994; Sanford & Garrod, 1981, especially chapter 8), referential and attributive uses may differ in the entities they point to. The model has four memory partitions that derive from two dichotomous distinctions: (1) active (dynamic) and long term (static) memory stores and (2) knowledge source—from the text or from general knowledge. Explicit focus contains tokens that are currently active due to textual information, while implicit focus contains both tokens and slots with default fillers that are currently active due to general knowledge. The static memory stores are long-term text memory and long-term semantic memory. When a definite description is encountered, a Retrieve process is initiated. In this situation, the focus or

active domain is searched. If an appropriate antecedent is found in explicit focus, the identity of that specific token is returned. If an appropriate match is found in implicit focus, the identity of that slot is returned. After the Retrieve, a Construct is executed. If a token was returned, the new information from the description is attached to that token. If a slot was returned, a token is added into explicit focus and linked to that slot. Indefinite descriptions call the processes in the reverse order (Construct then Retrieve).

Though not explicitly stated by Garrod et al. (1994), referential uses could be those in which a token is retrieved from explicit focus whereas an attributive use could be when a slot is returned from implicit focus. Note however, that in both cases a token, not a type or slot, is instantiated in explicit focus after the Retrieve and Construct processes are completed.

So Mueller-Lust's suggestion that referential uses point to a token while attributive uses point to a set can be questioned on several grounds. To test whether referential and attributive uses instantiate different discourse representations, an experimental method based on the linguistic argument set forth by Webber (1979) and Birner (1991) was used. Cohesiveness in text relies in part on devices that mark dependencies among expressions, including anaphora, which is a mechanism for pointing back in text. Pronouns are one type of anaphoric device, and they are not just substitutes for previously mentioned noun phrases (Donnellan, 1978; Gernsbacher, 1991) since it is possible for them to refer back to conceptual entities not overtly mentioned in the text (Carreiras & Gernsbacher, 1992; Hankamer & Sag, 1976).

In English, pronouns generally agree in number with their antecedents, so (7a) with a singular antecedent (*iron*) and singular pronoun (*it*) is fine, but (7b) with a singular antecedent (*iron*) and plural pronoun (*they*) sounds odd, even though it is semantically plausible in that irons in general are hot (examples from Gernsbacher, 1991).

- 7a. The iron is on the table. It is hot.
- 7b. ?The iron is on the table. They are hot.

However, if the referred-to entity was introduced with a grammatically singular NP that is conceptually plural, it is acceptable to refer back to that entity with a plural pronoun as in (7c) or (7d) (examples from Oakhill, Garnham, Gernsbacher, & Cain, 1992).

- 7c. Last night we went to hear a new jazz band. They played for nearly five hours.
- 7d. I need a plate. Where do you keep them?

This phenomenon is often called *conceptual anaphora* (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991; Oakhill et al., 1992) since pronoun

agreement is influenced by conceptual rather than grammatical number (Bock, 1995; Bock, Nicol, & Cutting, 1999). So if referential uses select a token, they should be understood as conceptually singular and better referred to by singular than plural pronouns (as in 7a). If attributive uses select a set, as suggested by Mueller-Lust, they should be understood as conceptually plural and plural pronoun reference should be acceptable (as in 7c).

Experiment 3 addressed whether an attributive description is conceptually singular (points to a single entity in the discourse representation) or is conceptually plural (points to a set), by referring to entities introduced attributively with a singular or plural pronoun and measuring the acceptability of the sentence with the pronoun. A preference for singular pronouns would indicate that an attributive use of an NP was interpreted as conceptually singular, while no preference for singular pronouns would indicate conceptual plurality. This asymmetrical prediction was made for two reasons. First, it follows previous findings (Carreiras & Gernsbacher, 1992; Oakhill et al., 1992). Second, pronoun agreement seems to be driven by conceptual plurality rather than grammatical plurality (Bock, 1995; Bock et al., 1999), so syntactically singular things that tend to appear in the world in collectives have two conceptual sources for pronoun agreement. For example, if *the band* is understood as referring to a single entity, the singular pronoun is appropriate. However, *the band* could also be understood as referring to multiple people, in which case the plural pronoun is appropriate.

A subsidiary aim of Experiment 3 was to replicate past studies of conceptual anaphora, in particular Oakhill et al. (1992) with true generics and a simpler task. The items Oakhill et al. called generic were not actually generic in the sense of referring to the whole category, as in (8a). Instead, they referred to an individual with prototypical properties, as in (8b) from Oakhill et al. (1992).

- 8a. A doberman makes a good guard dog.
- 8b. I was frightened by a doberman.

Also, the task they used involved collecting a reading time followed by a judgement as to the referent of a pronoun for each target sentence—tasks that may have interfered with one another.

In the present experiments subjects either made a judgement (Experiment 3) or their reading times were measured (Experiment 4), but not both. In the current experiment, sentences with pronouns were rated for how easy they were to understand after their antecedents had been introduced attributively, referentially, or generically. Referentially introduced entities should be treated as conceptually singular, since they are about a particular thing or person. Generically introduced entities should

be treated as conceptually plural (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991; Oakhill et al., 1992; Webber, 1980), meaning that both singular and plural reference should be considered appropriate. The question is whether attributively introduced entities are treated as more similar to referentially or to generally introduced entities.

#### Method

Subjects. Thirty subjects were recruited from the same population as Experiment 1 to participate in Experiment 3 for similar compensation. Twelve additional subjects participated in the norming of items. None of these subjects had participated in Experiments 1 or 2.

Materials. As in Experiment 2, only definite descriptions were used. For each of 24 definite NPs, three context sentences were created such that one was biased towards a referential, one towards an attributive, and one towards a generic reading of that target NP. Each of these sentences was followed by a sentence in which the target NP was used either referentially, attributively, or generically, as fit with the context sentence. Together these two sentences were the biasing contexts of the passage, as shown in Table 2. Half the noun phrases referred to people and half to nonpeople. This manipulation was included because there may have been a tendency to accept plural pronouns after references to people whose gender was unknown (the generic they as in: Each student should bring their [plural] own book to class; Miller & Swift, 1988). Thus, it seemed prudent to separate people and other types of antecedents. The effect of this manipulation was never reliable, nor did it interact with other variables, so it is not discussed further.

# TABLE 2 Example passages from Experiments 3 and 4

Referential biasing context

I went to the bookstore to see what was popular.

The best seller this month is by an author I dislike.

Attributive biasing context

The bookstore manager told me what to do with the new paperpacks.

The best seller goes in the center of the window.

Generic biasing context

There are a number of methods of determining the most popular paperback. The best seller is determined by the number of wholesale orders.

Pronoun sentence

It is/They are bound to get a lot of attention. [singular/plural]

To norm the biasing contexts, 12 subjects read them with the target NP underlined and decided whether the underlined phrase was intended to refer to a specific, known individual or not by selecting a paraphrase. Each subject received each of the 24 target noun phrases in either its referential or attributive biasing condition. This resulted in two lists of 24 items in which half the items were intended to be referentially biased and half attributively biased.

For most of the biasing contexts, subjects had no trouble choosing the expected paraphrase (as shown by Mueller, 1988; Mueller-Lust & Gibbs, 1991), but if 2 or more of the 6 subjects who received each target NP in a particular biasing condition (referential or attributive) did not choose the interpretation intended for that passage, the passage was reworded or changed to make it more strongly biased towards the intended interpretation. As a result, 8 of the 48 biasing contexts were changed.

The final passages used in the main experiment were composed of the

The final passages used in the main experiment were composed of the two context sentences: one that biased towards a referential, attributive, or generic interpretation followed by one that introduced an NP referentially, attributively, or generically. These were followed by a pronoun sentence which contained a singular (*he*, *she*, *it*) or plural (*they*) pronoun that was intended to refer to the same entity as the target NP.

One of the two versions (singular/plural) of the pronoun sentences followed each of the three biasing contexts, making six versions of each passage (see Table 2). Six lists were created that contained each of the 24 target NPs once. Each list contained four of the NPs in each of the six conditions in a semi-random order. Subjects were assigned to lists randomly. Stimuli were presented and ratings recorded by PsyScope 1.1 or 1.1b4 (Cohen, MacWhinney, Flatt, & Provost, 1993) running on Macintosh Quadra 630 computers.

Procedure. Subjects were told that they would read short passages on the computer and rate the goodness of the wording of the final sentence of each one. Before each passage, "PRESS SPACE BAR FOR NEXT STORY" was shown on the screen. When the subject pressed the space bar, the biasing context of the passage replaced the introductory message. When the subject had read and understood the biasing context and pressed the space bar, it was replaced immediately by the pronoun sentence. A 7-point rating scale was presented above this sentence. After a rating was entered, the introductory message reappeared, indicating the start of the next passage. Subjects were asked to rate, in relation to the first two (biasing) sentences, how well the third (pronoun) sentence was worded, and how difficult it was to understand what the sentence was about. The 7-point scale was labelled at only three points: BADLY WORDED (1),

SOMEWHAT CLEAR (4), WELL WORDED (7), though subjects were asked to use the whole scale.

### Results and discussion

The main effect of interest was whether the NP's introduction (referential, attributive, generic) affected conceptual number as indexed by relative acceptability of singular or plural pronoun reference. There was a main effect of pronoun type,  $F_1(1, 29) = 29.45$ ;  $F_2(1, 23) = 50.96$ , with higher ratings for singular, M = 4.8, than plural, M = 3.8, pronoun sentences in a context (3) × number (2) ANOVA (Table 3). Rating differences (singular – plural) will be reported as a measure of the strength of bias towards conceptual singularity.

The goal of Experiment 3 was to determine whether attributively introduced entities are conceptually singular or plural. There was a main effect of biasing context for the rating difference scores,  $F_1(2, 58) = 5.26$ ;  $F_2(2, 46) = 7.1$ . For generic contexts the difference scores did not differ from 0 [difference = .4;  $t_1(29) = 1.59$ , p > .12;  $t_2(23) = 1.86$ , p > .08] indicating conceptual plurality. However, rating differences for both referential [difference = 1.6;  $t_1(29) = 5.95$ ;  $t_2(23) = 6.7$ ] and attributive [difference = 1.3;  $t_1(29) = 3.72$ ;  $t_2(23) = 4.85$ ] contexts were reliably greater than 0, indicating conceptual singularity. That is, subjects dispreferred plural pronouns following attributively and referentially introduced NPs.

The results for generic references replicate the previous finding that generic uses are conceptually plural (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991; Oakhill et al., 1992) using true generics and a simpler task. Generics, like other conceptual plurals, can take either a singular or

TABLE 3
Mean ratings (Experiment 3) and reading times (Experiment 4) for pronoun sentence

		How NP was introduced				
		Referential	Attributive	Generic	Average	
Ratings (1–7)	plural	3.3	3.7	4.3	3.8	
	singular	4.8	5.0	4.7	4.8	
	difference (singular – plural)	1.6	1.3	.4		
Reading	plural	2572	2604	2319	2498	
Times (ms)	singular	2363	2273	2326	2321	
	difference (plural – singular)	208	331	<del>-7</del>		

plural pronoun. As noted above, our items were truly generic, in the sense of referring to the entire class, and so these results go beyond those reported by Oakhill et al. (1992), who tested specific references with prototypical properties. In contrast to the generics, singular pronouns were strongly preferred over plural pronouns to refer back to an attributively or referentially introduced entity, indicating conceptual singularity. So, a rating task found generically introduced entities to be conceptually plural and both referentially and attributively introduced entities to be conceptually singular. This contrasts with Mueller's (1988) proposal that attributive NPs refer to types, which predicted that plural pronouns should have been as acceptable as singular ones.

### **EXPERIMENT 4**

One possible problem with the rating task used in Experiment 3 is that it is a rather metalinguistic task. Subjects may have responded based on how they felt the sentences *ought* to be written rather than how easy the sentence was to understand. If so, the ratings would be telling us little about normal sentence processing. Experiment 4 used reading times, a more on-line task than ratings, of the materials from Experiment 3 to further examine whether attributively introduced entities are conceptually singular or plural. It is assumed that pronouns are resolved at the time of reading (Garrod & Sanford, 1994) and that reading a pronoun that is easier to understand is faster than reading one that is more difficult (Gernsbacher, 1991; Mueller, 1988). Thus, a mismatch between the discourse representation and a coreferential pronoun should affect reading time of sentences with the pronouns.

## Method

*Subjects.* Fifty-four additional subjects were recruited from the same population as Experiment 3. None had participated in the previous experiments.

Materials. The materials were the same as Experiment 3.

*Procedure.* Subjects were told that they were going to read short passages on the computer, sentence by sentence. Each passage began with a fixation cross at the left of the screen, and when the subject pressed a key on a PsyScope button box, it was replaced after 500 ms by the first sentence of the passage. After the subject read and understood the sentence and pressed the key, the sentence was immediately replaced by the next sentence. After the third (final) sentence of the passage, there was a blank

screen for 1000 ms, then the fixation cross appeared, indicating the start of the next passage. The time between presentation of each sentence and the subsequent key press was recorded.

# Results and discussion

One of us created for each subject a histogram of the reading times for the 24 pronoun sentences without condition information. The other examined the histograms to eliminate obvious outlier points. These were generally reading times more than 2 seconds longer than the next longest reaction time. This resulted in the elimination of 14 out of 1296 reaction times, or about 1% of the data. These points tended to be from the first or second trial of the experiment (7/14), confirming the assumption that they were outliers. They were approximately evenly distributed across conditions: 3 were referential, 6 attributive, 5 generic; 6 singular, 8 plural.

Reading a sentence with a singular pronoun more quickly than its plural counterpart was taken as an indication that the antecedent of the pronoun was conceptually singular. Overall, plural sentences, M=2498 ms, were read more slowly than singular ones, M=2321 ms;  $F_1(1,53)=10.13$ ;  $F_2(1,23)=12.67$ , which has been a general finding (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991; Murphy, 1984; Oakhill et al., 1992). Again, difference scores (plural – singular) for reading times of the pronoun sentences were analysed and are displayed in Table 3. There was a main effect of context on the difference scores [reliable by subject  $F_1(2,106)=5.86$ , but not by item,  $F_2(2,46)=1.40$ ]. For generic uses, the difference scores did not differ from 0, M=-7 ms;  $t_1(53)=.10$ ;  $t_2(23)=.77$ , indicating conceptual plurality. However, it took reliably longer to read plural sentences for both referential [difference = 208 ms;  $t_1(53)=2.57$ ;  $t_2(23)=2.57$ ] and attributive [difference = 331 ms;  $t_1(53)=3.68$ ;  $t_2(23)=3.22$ ] contexts, indicating conceptual singularity.

The results from this relatively on-line reading time measure corroborate those found in the possibly more metalinguistic rating task, suggesting that ratings did tap into processing and not prescriptive judgements. Though, as described above, the stimuli in this experiment were slightly different, these results agree with the conclusions of Carreiras and Gernsbacher (1992), Gernsbacher (1991), and Oakhill et al. (1992) in showing that entities introduced generically were conceptually plural. Entities introduced referentially and attributively were found to be conceptually singular as indicated by faster reading times for singular pronouns. The lack of a reliable overall item analysis weakens the strength of the conclusions slightly, but this is partially outweighed by the reliable (by both subjects and items) reading time advantage for singular pronouns after attributive and referential but not generic descriptions.

### GENERAL DISCUSSION

Four experiments explored the comprehension and discourse representation of descriptions (1a, repeated below) used referentially (referring to a specified person) or attributively (referring to whoever did the killing).

### 1a. The murderer of Smith is insane.

The first goal was to explore some factors that might cue the addressee as to the speaker's referential versus attributive intent. Subsidiary aims were to demonstrate that the referential/attributive distinction can affect which potential referent is thought to be referred to, and that a similar distinction holds for both definite and indefinite descriptions. The second goal was to examine whether one difference between referential and attributive uses is in the discourse representations they instantiate.

With regard to the first set of questions, Experiment 1 used a forced choice paradigm in which subjects could choose one of three racers, and their choice indicated whether the target sentence had been interpreted referentially. The experiment demonstrated that several factors influenced which entity was understood as the referent. Mueller-Lust (Mueller, 1988; Mueller-Lust & Gibbs, 1991) had shown that people could distinguish referential and attributive interpretations by choosing appropriate paraphrases, but Experiment 1 went further by demonstrating that the referential/attributive distinction has consequences for what is understood as the referent of a description. It also provided experimental evidence that this distinction could be made without overt syntactic cues such as whoever (Kripke, 1979) or that (Roberts, 1993). Also, by demonstrating experimentally that the same sentences were interpreted differently depending on surrounding context, the results indicate that predicate type cannot be the sole determinant of the referential/attributive distinction as suggested by Ortony and Anderson's (1977) stimuli (though it was one of the factors that influenced interpretations).

Furthermore, although the use of *the* led to more referential responses than *a*, as was expected since *the* is often used to pick out an entity uniquely identifiable in the context, the experiment showed that both definite and indefinite descriptions were subject to influences by the same factors. This provides evidence that there may be a similar distinction for both types of descriptions (Roberts, 1993).

In particular, three factors were shown to systematically affect whether the target sentences were interpreted referentially or attributively. Verb type had the strongest effect: when the verb was *own*, there was a strong bias towards referential interpretations; when *like* was used, the proportion of referential responses decreased, depending more on other factors. Experiment 2 confirmed this effect of predicate using a variety of verbs.

This experiment also suggested that the strength of the link between a predicate and its referent may contribute to its interpretation as referential or attributive. A stronger link may increase the likelihood of an attributive interpretation by focusing on the meaning of the description.

Context also changed the proportions of referential responses. Experiment 1 varied the amount of shared knowledge and the degree of emphasis on individuals versus properties. Property contexts that indicated an intermediate degree of shared knowledge among scenario participants and emphasised the importance of the qualities of racers led to the lowest proportion of referential responses. Individual contexts that indicated a high degree of shared knowledge and emphasised the importance of a particular racer, and Neutral contexts that indicated little shared knowledge and provided no information about the relative importance of a particular racer versus qualities of racers in general, led to more referential responses. Our experiment confounded the variables of shared knowledge and individuality. It would be interesting to examine these factors separately as well as to look at other aspects of context that might be important, such as the authority of the speaker (Ackerman, 1979) and the familiarity of the referent (Boland & Dell, 1991; Mueller, 1988; Mueller-Lust & Gibbs, 1991). However, the results of both the verb and context manipulations provide clear evidence that factors outside the description itself can have very strong effects on whether the NP is understood as referential or attributive.

The second major question of this paper was whether referential and attributive uses differed in the entities they instantiate in the discourse model. Experiment 3 measured ratings for a sentence with a singular or a plural pronoun that referred to an entity that had been introduced by an NP that had been used referentially, attributively, or generically. As found by Gernsbacher and colleagues (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991; Oakhill et al., 1992), the experiment showed equivalent ratings for both singular and plural reference back to generically introduced entities (conceptual plurality). Referentially introduced and attributively introduced entities were shown to be conceptually singular (higher ratings for singular than plural references). Experiment 4 used the same passages but measured reading times as a more on-line measure of comprehension. Again, generically introduced entities were understood as conceptually plural, while both referentially and attributively introduced entities were understood as conceptually singular. Therefore, the findings of this paper do not indicate that there is a discourse representation difference between referential and attributive uses of the sort suggested by Mueller-Lust, viz. that referential uses instantiate tokens while attribute uses instantiate sets, which are conceptually plural. Rather, the experiments are in line with suggestions that both referential and attributive uses

refer to single entities in the discourse representation (Birner, 1991; Webber, 1979).

The data from Experiments 3 and 4 suggest that Mueller-Lust's (1988, 1991) proposed discourse model for referential and attributive interpretations is incorrect since neither use seems to refer to a set of entities. An intuitive way of incorporating the referential/attributive distinction into existing discourse representation models (e.g., the Focus Memory Model; Sanford & Garrod, 1981) would be to hypothesise that referential uses pick out a specific token in explicit focus and either link it to a known slot from implicit focus or add the new information into explicit focus linked to that token. Attributive uses would pick out a slot (murderer) in implicit focus and link that to an underspecified token in explicit focus. The attributive reference would be interpreted as singular due to the slot's link to the underspecified token which would not stand for the set of murderers but only for a single contextually relevant but unknown murderer.

How do referential and attributive descriptions achieve different interpretations, then, if they do not differ in their discourse representation? One possible explanation is that they share the reference to an individual but bring differing presuppositions to that reference (Donnellan, 1966). As mentioned previously, the attributive interpretation of (1a) presupposes that there is some X that is the murderer of Smith. The referential use has the presupposition of the attributive use, plus the presupposition that there is some particular X, in this case the person who was just convicted, who can be identified as the murderer.

Together, these four experiments demonstrate that people can distinguish referential and attributive uses, but that both are represented by conceptually singular discourse referents. The results will help to constrain models of discourse representation that attempt to capture this important distinction.

Manuscript received September 1998 Revised manuscript received October 2000

### REFERENCES

Ackerman, B.P. (1979). Children's understanding of definite descriptions: Pragmatic inferences to the speaker's intent. *Journal of Experimental Child Psychology*, 28, 1–15.

Birner, B.J. (1991). Discourse entities and the referential/attributive distinction. 65th Meeting of the Linguistic Society of America, Chicago 4.

Bock, K. (1995). Producing agreement. Current Directions in Psychological Science, 4, 56–61.
Bock, K., Nicol, J., & Cutting, J.C. (1999). The ties that bind: Creating number agreement in speech. Journal of Memory and Language, 40, 330–346.

Boland, J., & Dell, G. (1991, March). When isn't the pope Polish?: Establishing reference for NPs in discourse. Poster presented at the CUNY Sentence Processing Conference, New York.

- Carreiras, M., & Gernsbacher, M.A. (1992). Comprehending conceptual anaphors in Spanish. Language and Cognitive Processes, 7, 281–299.
- Clark, H.H. (1979). Responding to indirect speech acts. Cognitive Psychology, 11, 430-477.
- Cohen, J.D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavior Research Methods, Instruments, & Computers*, 25, 257–271.
- Cole, P. (1978). On the origins of referential opacity. In P. Cole (Ed.), *Syntax and semantics: Pragmatics*, pp. 1–22. New York: Academic Press.
- Donnellan, K. (1966). Reference and definite descriptions. *Philosophical Review*, 75, 281–304.
  Donnellan, K.S. (1978). Speaker reference, descriptions and anaphora. In P. Cole (Ed.), *Syntax and semantics: Pragmatics*, pp. 47–68. New York: Academic Press.
- Garrod, S., Freudenthal, D., & Boyle, E. (1994). The role of different types of anaphor in the on-line resolution of sentences in a discourse. *Journal of Memory and Language*, 33, 39–68.
- Garrod, S.C., & Sanford, A.J. (1994). Resolving sentences in a discourse context: How discourse representation affects language understanding. In M.A. Gernsbacher (Ed.), *Handbook of psycholinguistics*, pp. 675–698. San Diego: Academic Press, Inc.
- Gernsbacher, M.A. (1991). Comprehending conceptual anaphors. *Language and Cognitive Processes*, 6, 81–105.
- Gibbs, R.W., Jr. (1984). Literal meaning and psychological theory. Cognitive Science, 8, 275–304.
- Grice, P. (1975). Logic and conversation. In S. Davis (Ed.), *Pragmatics*, pp. 305–315. New York: Oxford University Press.
- Hankamer, J., & Sag, I.A. (1976). Deep and surface anaphora. *Linguistic Inquiry*, 7, 391–428. Johnson-Laird, P.N., & Garnham, A. (1980). Descriptions and discourse models. *Linguistics and Philosophy*, 3, 371–393.
- Kripke, S. (1979). Speaker's reference and semantic reference. In S. Davis (Ed.), *Pragmatics*, pp. 77–96. New York: Oxford University Press.
- Miller, C., & Swift, K. (1988). The handbook of nonsexist writing (2nd ed.). New York: Harper and Row.
- Mueller, R.A.G. (1988). Anaphora resolution of attributive and referential definite descriptions. Doctoral dissertation, University of California, Santa Cruz. *Dissertation Abstracts International*, 49, 3667.
- Mueller-Lust, R.A.G., & Gibbs, R.W. (1991). Inferring the interpretation of attributive and referential definite descriptions. *Discourse Processes*, 14, 107–131.
- Murphy, G.L. (1984). Establishing and accessing referents in discourse. *Memory & Cognition*, 12, 489–497.
- Oakhill, J., Garnham, A., Gernsbacher, M.A., & Cain, K. (1992). How natural are conceptual anaphors? *Language and Cognitive Processes*, 7, 257–280.
- Ortony, A., & Anderson, R.C. (1977). Definite descriptions and semantic memory. *Cognitive Science*, 1, 74–83.
- Recanati, F. (1993). Direct reference: From language to thought. Oxford: Blackwell.
- Roberts, L.D. (1993). How reference works: Explanatory models for indexicals, descriptions, and opacity. Albany: State University of New York Press.
- Russell, B. (1905). On denoting. *Mind*. Reprinted in R.C. Marsh (Ed.), *Logic and knowledge: Essays 1901-50*, pp. 41–56. New York: Macmillan Co.
- Sanford, A.J., & Garrod, S.C. (1981). Understanding written language: Explorations of comprehension beyond the sentence. Chichester: John Wiley & Sons.
- Searle, J.R. (1979). Referential and attributive. Monist, 62, 190-208.
- Strawson, P.F. (1950). On referring. Mind, 59, 320-344.
- Webber, B.L. (1979). A formal approach to discourse anaphora. New York: Garland.
- Webber, B.L. (1980). Syntax beyond the sentence: Anaphora. In R.C. Spiro, B.C. Bruce, & W.F. Brewer (Eds.). Theoretical issues in the study of reading, pp. 141–164. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.