Small Worlds of Discourse and the Spectrum of Accommodation:

A Psycholinguistic Investigation of the Accommodation Entailed by the Definite Article when its Determinant is Discourse-New

William Evans
Honors Thesis
Linguistics and Psychology
Umass Amherst
Spring 2005

Abstract

I conducted three experiments in which I investigated how people use the definite article with entities that have never been explicitly introduced. I claim that using the definite article to determine a particular entity presupposes that that entity has a pre-established discourse referent. If an explicit discourse referent has not been supplied by previous context, additional presuppositions must be accommodated to infer one instead.

In experiment 1, participants were given a self-paced reading task to determine if less felicitous accommodations took longer to process than more felicitous accommodations. I compared both supported and un-supported inferences of a target discourse-new entity to contexts in which the entity had already been explicitly mentioned. I found that prior mention of the target entity made subsequent processing times of the same entity significantly faster, and well-supported accommodations were marginally faster to process than poorly-supported accommodations.

In experiment 2, I investigated the requirements of the definite article by asking participants to rate the felicity of two-sentence pairs in which one entity was singled out with the definite article, from a set of equally plausible entities. In one condition, context made one entity clearly more salient than others of its type. In the second condition, all entities were equally salient. I found that when the target entity was singled out over others of its type, items were rated significantly higher than when there was no reason to single out one entity. This indicates that in order to use the definite article, context must supply the designated entity with a single, unambiguous discourse referent with which it can link. This means that the entity must be both relevant and more salient than all others of its type. I call this view the Small World hypothesis. It resolves the traditional distinction between uniqueness and identifiability with the concept of discourse referents, because in order for the entity to have a single unambiguous discourse referent candidate, the target entity must both be identifiable and locally unique.

In experiment 3, I conducted a forced-production study. I had participants supply either the definite or indefinite article to contexts that met the requirements of the Small World hypothesis by varying degrees. I found that in contexts that met the requirements of the Small World, participants consistently chose the definite article. In contexts where the requirements were only partially met, participants chose either the definite or the indefinite. In contexts where the requirements were not met, participants consistently preferred the indefinite article.

I combined the findings of experiments 1, 2 and 3 to motivate the Spectrum of Accommodation, in which I claim that when using the definite article, the presuppositions necessary to treat discourse-new entities as discourse-old ones are accommodated with relative ease or difficulty, depending upon how well context manages to make a particular entity relevant and more salient than all others of its type, by linking it up to an unambiguous discourse referent.

Table of Contents

Section I: Introduction

When constructing discourse, speakers constantly introduce new information. Not surprising, since this seems to be its primary purpose. The surprising thing is that speakers often treat this new information as if it had already been introduced. One way speakers go about this in English is by introducing new information within a DP, headed by the definite article. The definite article is generally used with discourse referents that have already been previously introduced by other means, such as with the indefinite article. Let us look at some examples to gain some intuitions of what I am talking about:

1) Ronny found a kitten on his doorstep this morning. **The kitten** was mewling piteously.
2) Loretta was writing with a pen. **The pen** kept leaking.
3) Loretta was writing with a pen. **The ink** was running low.
4) Harris went to eat at an old diner. **The menu** hadn't changed in years.
5) Harris went to find a new apartment. #**The garden** hadn't changed in years.

In (1) and (2), the new entity is first introduced with the indefinite article, and then referred to with the definite article once it is an established part of discourse. Ho-hum. But what about examples (3) and (4)? How are they different from (5)? In none of these examples is the entity first explicitly introduced with an indefinite, and yet (3) and (4) seem to be perfectly felicitous. There may be an intuition about the relationship between Pens and Inks, or Diners and Menus, a relationship not evidenced as clearly between Apartments and Gardens, but the way that these relationships make (3) and (4) both acceptable and common bears investigation.

So what processes allow us to treat discourse-new entities as discourse-old ones? When does this happen and what makes it possible? In the following pages, I give a simple account of what occurs when we introduce new information with the definite article in English, both in terms of semantics and in terms of processing. I begin with an introduction to definiteness, followed by a brief explanation of presupposition and accommodation, setting the scope of my investigation from a semantics perspective. I continue with some background on discourse processing, with an overview on mental models, then discuss previous work involving discourse referents, bridging anaphora, and the processing of given versus new information.

These sources will motivate my own research on the processing of accommodation involving the definite article. I have by no means attempted a complete treatment of either accommodation or definiteness, but I hope to bring all these disparate sources together to create a more unified account of what occurs when we accommodate the presuppositions that allow us to introduce new information with the definite article. I propose a new way of framing the issue called the **Spectrum of Accommodation**, a hypothesis claiming that our ability to accommodate new information with the definite article is context-dependent, and that for each instance, processing time and felicity can be mapped along a continuum of relative difficulty, depending upon the interaction of various criteria I will discuss. Furthermore, I will bring together data from different studies that used self-paced reading, eye-tacking, and felicity judgments to support my

claims. But before I can get into the details, I need to give a more general overview of the nature of this topic. And since I wish to discuss how speakers process the accommodation of presupposition when using the definite article, a good place to start is with the properties of the definite article itself.

Section II: Definiteness

In the book *Definiteness* (1999) Christopher Lyons gives a comprehensive and accessible view of the topic of the same name. According to Lyons, a noun phrase headed by the definite article is an example of what he calls a **Simple Definite**. Other examples of Simple Definites in English are noun phrases headed by other determiners such as *this,* or possessive pronouns like *his*. Although Lyons merely discusses simple definites as a means of setting the stage for more complex work, I believe that for my purposes, an informal account of simple definites will suffice.

Lyons (1999) discusses two general ways of accounting for definiteness, **familiarity** and **uniqueness** . Familiarity is a hypothesis stating that when the definite article is used, it indicates that the noun phrase it determines must be known by both speaker and hearer. Another way to describe this is to say that the definite article is used to denote **hearer-old** information (Clark, 1975). Take two examples from *Definiteness*:

6) I bought **a car** this morning.
7) I bought **the car** this morning.

In both cases, the speaker is referring to a specific, unique car. But while in the first case, the *speaker* may know the information that makes it unique and specific, that information is probably not known by the hearer. In the second sentence, however, there is the assumption that the *hearer* also knows exactly which car I am talking about. If that assumption is not validated by prevailing context, then (7) would be an infelicitous utterance. This seems simple enough, and familiarity is an adequate criterion for many uses of the definite article, but it fails to account for the felicity of examples like (4).

Intuition tells us that we should not consider *the menu* to be known by the hearer in the same sense as *the kitten* is known in (1). Because of examples like (4), another take on familiarity is a reformulation known as **identifiability**. Instead of claiming that the referent has to be known by both the speaker and the hearer, identifiability claims that using the definite article indicates that the hearer has been given enough context to figure out what the DP refers to. It indicates that the hearer has the capacity to identify the referent in existing discourse (Lyons, 1999).

Although identifiability is a descriptive improvement when compared to familiarity, it is far from a complete account. Take another example from *Definiteness*.

8) I've just been to a wedding. **The bride** wore blue.

Lyons (1999) claims that identifiability cannot account for this example, because there is very little to identify the referent. The only thing we can conclude is that it is very likely that *the bride* is a unique entity, given the context. This observation spawned **uniqueness** as a formulation of definiteness; Namely, that the definite article is used to

denote an entity that is unique within a certain context. Like familiarity, uniqueness requires reformulation upon closer investigation due to examples that involve plurals:

9) We've just been to see John race. The Queen gave out **the prizes** (Lyons, 1999).

The critical DP above contains a plural set of entities, and its behavior has been explained by Hawkins (1978). He asserts that "the reference is the totality of the objects or mass in the context which satisfy the description" (Lyons, 1999). Hawkins calls this **inclusiveness**. He means that using the definite article must denote all appropriate members of a given type within a given context. uniqueness is merely an extension of this because in contexts where there is only one member of a certain type, inclusiveness and uniqueness are the same thing.

Lyons asserts that uniqueness is a type of inclusiveness, and not the other way around. He argues that if you try to apply uniqueness to plurals, simply by looking for a unique set instead of a unique individual, then the criterion would apply to subsets of that set as well. In (9), *the prizes* could refer to a subset of two prizes, even if three were handed out overall, unless inclusiveness tells you that it must be the entirety of the implied set.

I believe that this distinction can be argued either way, and that uniqueness need not be a special case of inclusiveness. As Lyons himself said, uniqueness requires that, in order for the definite article to refer to a unique set, it has to be the only delineable set of its type. The only way for that to be ensured is for inclusiveness to set the largest appropriate scope for the article, given the context. So just as uniqueness could be seen as a special example of inclusiveness applied to singleton sets, inclusiveness could be seen as a natural repair for uniqueness when dealing with plural sets. In a later section, one of my own studies will lend support to the latter relationship between uniqueness and inclusiveness, due to the properties of singling out one entity from a plural set. In addition, I will argue that while there are obvious distinctions between identifiability and uniqueness from a semantic point of view, the two criteria can be unified into something very close to the identifiability position when interpreted in terms of discourse representation, with uniqueness acting merely as a way to satisfy identifiability.

Before addressing these claims, however, I am going to give a brief introduction to presuppositions and accommodation, since they are necessary elements of my future discussion.

Section III: Accommodation and Presupposition

One of the first published works on accommodation was a 1979 paper by David Lewis, *Scorekeeping and the Language Game*, in which he tried to describe the rules involved in governing acceptable conversation. He believed that when people engage in conversation, many statements require the listener to accept certain assumptions about the world that are never explicitly stated. To use Lewis' famous example:

10) The King of France is bald

A listener who accepts this as a true statement has to **presuppose** that France has a king. A presupposition is basically an assumption a hearer must make at a certain point in order for him to accept the statement he is interpreting as true. Lewis called the process of incorporating these assumptions **accommodation**. Accommodating presuppositions such as the one in (10) allows conversations to progress quickly and easily, and keeps people from having to digress to the most basic epistemology every time they open their mouths.

Lewis formulated specific rules of accommodation that deal with definite and indefinite articles. He claimed that the definite article can be used to describe only the most salient entity in a domain of discourse. He did not believe that the definite article designates only unique entities, because of examples like the following:

11) The dog chased another dog.

The dog is not a unique entity within the context, but if we have designated the first dog with the definite article, it must be the more salient than the other dog. This ties in well with Lyons' description of identifiability because a high level of salience indicates that context should make it easy to figure out the reference.

Now that we have informal accounts of presupposition and accommodation, it is possible to restate identifiability and familiarity in these terms, making the distinctions between the two accounts more specific. Using the definite article entails a presupposition; The hearer has to assume that the entity can be specified by the context.

Familiarity claims that the referent has to be hearer-old. In order for this account to be supported, every time we came across the presupposition entailed by the definite article, it has to be supported by previous explicit context. Identifiability, on the other hand, takes further advantage of accommodating presuppositions. When the definite article forces us to assume that our referent is identifiable, then under certain conditions we can accommodate further presuppositions to help disambiguate its identity. Take, for example:

12) Harold turned on a computer at work. **The keyboard** wouldn't work.

Familiarity cannot account for this example because *keyboard* has never been explicitly introduced. Identifiability predicts that this example is acceptable because context allows us to accommodate a presupposition that identifies the referent. In this case, we accommodate the presupposition, "the computer had a keyboard." Since familiarity needs the reference to be previously known to the hear via explicit introduction, it would not predict (12) to be acceptable. If a speaker wanted to utter the second sentence in (12), he would have to say something like:

13) Harold turned on a computer at work *that had a keyboard*. **The keyboard** wouldn't work.

There is a strong intuition that "had a keyboard" is redundant in (13); It is something we can easily infer, given the knowledge we possess about computers. Why is

this the case? In a later section, I will discuss in further detail the types of inference that seem to license the accommodation of definiteness

It is easy to see how beneficial accommodating presuppositions is in creating discourse, since it allows speakers to introduce new discourse referents and new information about existing discourse referents with being required to draw every connection and association explicitly (even if this were possible, conversations would have to progress on a geological time scale). Accommodating presuppositions is a process that allows us to produce efficient, flowing speech, but how exactly do we go about it? Given a certain context, which presuppositions are acceptable and which are unacceptable? How do the mechanisms that produce these discourse effects show up in processing time? In the following section, I will begin to answer these questions by giving an overview of some of the processing literature, starting with mental discourse models.

Section IV: Gist & Mental Representation

When discussing the relationship between language, thought, and mental representation of the world, people have claimed that the terms language and thought are interchangeable, and that our mental representations of phenomena and the words we used to describe them are one in the same. It turns out that there is very little evidence to support this claim.

Words do not seem to form our underlying mental representations in and of themselves. As Steven Pinker argues in *The Language Instinct*, how can words be the true underlying form of our thoughts when everyone has experienced the feeling of failing to articulate what they are trying to say? If words and thought are one in the same, then where does this feeling come from (Pinker 1994)? The intent to articulate must exist in some form separate from the product of that intent. There must be some underlying mental representation of meaning that we sometimes fail to put into words. People have hypothesized that language users rely on a mental discourse model to translate a more abstract system of mental representation into words and words back into a more abstract system of representation.

What exactly is a mental discourse model? In one of the classic works on the topic, *Mental Models as Representations of Discourse and Text,* Alan Garnham gives a basic account. The idea is that explaining the mechanisms of speech production and comprehension simply in terms of surface form is inadequate, and that even a semantic account must go beyond an explanation based on direct relations between words. Garnham believes that in order to account for the observed psycholinguistic data, language users must use a superordinate, abstract system of discourse representation that can both generate, and be generated from language.

In support of this claim, Garnham reviews of a large body of early research that used recall studies to give evidence for an underlying representation of meaning. This studies showed that people often only remember the **Gist** of an utterance, and not its exact form. When we understand the basic meaning of something without remembering the exact words, we have formed a gist (Garnham 1987). Early researchers (Bartlett, 1932), realized that gist formation was more than an interesting memory flaw; It had

interesting implications in terms of mental representation. These implications were first explored using recall studies.

The basic process in these studies was for a participant to be given a section of discourse, either written or spoken, which he was later asked to recall, word-for-word. In recognition versions of the task, they were given an initial passage and then a second passage at a later date, which they had to determine whether or not it was identical to the first, or they were asked to identify the sentence they had heard or read from a number of options.

In one such study, Johnson-Laird and Stevens (1970) looked at the form of our mental representation. They wanted to see whether or not we encode discourse in syntactic form, using Deep Structures proposed by Chomsky (1965), or if it was encoded in terms of semantic meaning. Chomsky's idea was that there was is an underlying level of syntax in which certain relationships such as subject and object are assigned. Any one deep structure is capable of spawning various surface forms (Garnham, 1987).

Johnson-Laird and Stevenson tried to see if deep structure could account for gist, or if the underlying representation had to be semantic in some form. They created test items in which sentences had basically the same content, but differing deep structures due to altered subject-object relations. Participants were read a sentence like (14) as part of a passage.

14) John liked the painting and he bought it from the duchess.

Less than a minute later, participants were given a forced-recall test, in which they had to pick the sentence they had just heard as one of 8 options, seen below:

15)
a. John liked the painting and he bought it from the duchess.
b. John liked the painting and the duchess sold it to him.
c. The painting pleased John and he bought it from the duchess.
d. The painting pleased John and the duchess sold it to him.
e. The duchess liked the painting and she bought it from John.
f. The duchess liked the painting and John sold it to her.
g. The painting pleased The duchess and she bought it from John.
h. The painting pleased The duchess and John sold it to her.

The first four versions all had the same basic semantic content as the sentence from the passage, while the second four had differing semantic content, due to the fact that John and the Duchess had switched roles. Each version had a different deep structure. Half the participants had been warned that they were going to be asked to recall specific texts, half had not. Both warned and unwarned participants consistently chose options that agreed with the target sentence semantically, but only those participants that has been warned were regularly chose the option that agreed with the target sentence word-for-word (Garnham, 1987).

The results of this study indicate that people only encode syntactic structure into long-term mental representation when they know they are expected to recall it. If people used deep syntax to encode meaning, then all participants would have been able to

distinguish between each of the eight options. This implies that information from discourse must be represented and retained in some other form, perhaps by using a discourse model that keeps track of information about discourse referents.

Section V: Discourse Referents

We have seen evidence that we represent information in some underlying form, since our capability to recall the exact wording of a discourse is diminished except when we expect a recall test. In most cases, we only remember the gist. Can a discourse model explain this? What exactly do we encode when we understand discourse, and what is the nature of the mental model that lets us produce it? Garnham says the best way to capture discourse representation is to use the concept of **discourse referents**. Informally, we can think of a discourse referent as a mental slot for a specific, unique entity within a given discourse. When a new entity is introduced to discourse, it is given a mental token to refer to it in the underlying discourse model. New information about the entity may be added to the same discourse referent, and once a relationship is established, semantically different descriptors may refer back to the same discourse referent. This multiple descriptors all referring back to the same entity is known as **coreference** (Garnham, 1987).

Evidence for the concept of discourse referents can be seen in a recall study done by Anderson and Bower (1979). Participants were given lists of sentences where they used world knowledge to set up coreference of individuals between names and descriptions (Garnham, 1987). For example, a given list might contain the following:

16)
a. George Washington had good health.
b. The first president of the United states was a bad husband.
c. Abraham Lincoln was a good husband.
d. The president who freed the slaves had good health.

After hearing the sentences participants would later be asked to recall the correct version of each sentence. When asked to recall (a), their options would consist of the following:

17)
a. George Washington had good health.
b. George Washington had bad health.
c. The first president of the United States had good health.
d. The first president of the United States had bad health.

Participants were able to recall the correct sentence about 62% of the time, but most of the errors the made were due to a confusion of the referent, and not the descriptor. When asked to recall, “George Washington had good health,” they picked “The first president of the United States had good health” 22% of the time (Garnham, 1987). This tendency to confuse the individual over the predicate indicates that

participants were using coreference to have both name and descriptor refer back to the same discourse referent.

In addition, both name and descriptor occurred in separate sentences, something that earlier semantic theories, such as the one proposed by Katz and Fodor (Katz, 1972) have difficulty accounting for. Their theory only treated sentences and their referents in isolation, without a means to connect the same referent through a series of sentences, as well as lacking a means of discussing why the same entity can be referred to in different ways (Garnham, 1987). Using the concept of discourse referents is a way of getting beyond this difficulty, and can account for examples like (18):

18) A man arrived at church late. The man was getting married. The priest saw the groom arrive, and sighed in relief.

In this example, when we see "a man," we give the entity its own slot in our discourse representation. Information that refers to this entity is then linked to this DR as it appears, even across sentences. One of the easiest ways to represent a discourse such as (16) is to use the box notation developed for **Discourse Representation Theory**, or **DRT** (Hans and Kamp, 1993). In this system, discourse is represented as a file card, and each entity is given a variable (x,y...) as its unique discourse referent. Information about that discourse referent is listed with that variable. The discourse referents of this system seems to share similarities with those in the mental models discussed by Garnham, but whether or not they are the same is hard to say. Nonetheless, DRT notation is a useful way to think about how a discourse referent might function. If we were to use DRT to describe (18), the first sentence would be represented as follows:

19)

x
man (x) arrived at church late (x)

As we can see, the discourse referent is placed on the top of the card, and designated with the variable X. Information about the referent is listed below. We could add the second sentence to the discourse box as follows:

20)

X
man (x) arrived at church late (x) was getting married (x)

And once we add the third sentence we get:

21)

x, y
Man (x) Arrived at church late (x) Was getting married (x) Groom (x) Priest (y) Saw x (y) Sighed in relief (y)

As we can see, when a new entity, *the priest,* is introduced into discourse, we see it is given its own discourse referent, and information about it is listed under the new heading. This notation clearly shows the concept of discourse referents in graphic form, but can this concept be used to better describe accommodation in terms of definiteness? Why is it possible to coreference (x) as *a man*, *the man*, and *the groom*? And why can we refer to (y) as *the priest* when (y) has not be previously introduced?

As I discussed previously, the definite article entails the presupposition that you can identify the discourse referent of the DP it heads (per identifiability). If you cannot identify a unique discourse reference given the context, then you should not be using the definite article. When (x) is first introduced, there is no existing context in which its discourse referent can be identified, so it is introduced with the indefinite article.

When the second sentence begins with *the man*, we assume we are still referring to (x) because at this point in the discourse, (x) is still the only identifiable referent. So why does *the groom* refer to (x) as well? Again, the definite article presupposes that the referent is identifiable, and (x) is already there. Additionally, we have been supplied with the information that (x) is getting married, so when we are given *groom*, we use its lexical content to link it to (x).

So why is *the priest* not identified with (x) as well at the time it is introduced? We already know (x) is getting married, so in order for *the priest* to serve as an antecedent, we would have to accommodate the presupposition, “Priests are likely

candidates for marriage," a presupposition not supported by our world knowledge (if the priest is catholic, at least). There are additional complications in this case because repeated use of "the" creates an overspecification problem that creates processing difficulties.

It is clear at this point why *the priest* does not refer to (x), but how can it both be definite and refer to the new discourse referent (y) at the same time? Again, the definite article tells us we can identify its designated referent, but in this example there are no appropriate pre-existing antecedents, since we have already discounted (x) as a candidate. At this point, we must accommodate a presupposition that singles out *the priest* as a specific, identifiable entity, and create the discourse referent (y) at the point of its introduction. Fortunately, a presupposition such as, "weddings frequently have one (and only one) priest in attendance," is easy to accommodate given what we know about the world.

In this section, I have reviewed evidence that indicates the need for a mental model of discourse to account for the observed recall data. I have described one such model that uses discourse referents to represent information in a form separate from that of surface utterances, and I have described how such a system handles presupposition and accommodation. However, I have glossed over the way in which internal mechanisms actually processes accommodation within a discourse model. In the following section, I will review some previous literature on sentence processing. I will explain what self-paced reading, eye-tracking, and ERP studies tell us about accommodating new information into discourse.

Section VI: Processing with Discourse Models

Although the data gathered in recall studies is useful for concluding that we use some form of mental representation to form gists, it would be dangerous to conclude from these studies anything about the real-time encoding procedure at the actual time of utterance. Encoding results may be sullied by the retrieval processes (Garnham, 1987). It is necessary to initial encoding procedures by some other means such as eye-tracking, phoneme monitoring, or self-paced reading.

The basic procedure in a self-paced reading study is for a participant to be given a body of text on a computer screen, broken up into segments. The participant presses a button, and the first critical region appears. When he finishes reading that segment, he presses the button, and a new segment of text appears. The time it takes him to read each segment of text is recorded. The methodology behind this process is that, all things being equal, sentences of the same length and complexity should take about the same amount of time to read. Thus by manipulating the specific content of a given text across a number of participants and items, it can be inferred that differences in reading times are caused by differences in processing difficulty.

One study by Garrod and Sanford (1977, Expt 1), used self-paced reading to investigate the relationship between processing a category and a member of that category. They gave participants two sentences pairs like the following:

22)

a. A bird would sometimes wander into the house. The robin was attracted to the larder.
b. A robin would sometimes wander into the house. The bird was attracted to the larder.

In the first version, category preceded instance, and in the second, instance preceded category. They tested the time it took for participants to read the second sentence of each version. They found that reading times were significantly faster in (b), when instance preceded category (Garnham, 71). These results can be explained in terms of definiteness and accommodating presupposition.

In (a), we add a new discourse referent to our model when we read *a bird* in the first sentence. Then in the second sentence, when we read *the robin*, the definite article tells us that we should have context to identify *robin*. We fulfill this requirement by hooking it up to the existing discourse referent created by *bird*, by accommodating the presupposition, "The bird we have been talking about is in fact a robin." This accommodation is unnecessary in the second sentence of (b), because we know robins to be birds by definition. There is no new information to accommodate, which means there is no processing deficit when we try to link *bird* to the existing discourse referent created by *robin*.

These results go against the **instantiation hypothesis** proposed by Anderson et al (1976), in which they claimed that word meaning was based entirely on context, and that words were automatically "fleshed out" by their context (Garnham, 1987). If the hypothesis had been supported, there would have been no difference in processing time between the second sentences of (a) and (b) because readers would already have assumed that the bird was a robin in the first sentence, robins being the most typical type that fit the context. This does not seem to be the case, and additional evidence seemed to show that when introducing new information with the definite article, all the processing occurred at the point of introduction, and not before.

In another self-paced reading study, Haviland and Clark (1974) found a processing-time penalty when incorporating new information into discourse (Garnham, 72). They also used two-sentence pairs, but altered the content of the first sentence while leaving the second the same. Participants read either (a) or (b) of the following:

23)
a. We got some beer out of the trunk. The beer was warm.
b. We checked the picnic supplies. The beer was warm.

They found that reading times were faster when *beer* had been explicitly mentioned in the first sentence, as in (a), compared to beer has not been mentioned, in (b). They called the difference in information states between both target *beers* **Given vs. New**, and explained that new information was slower to process because it required a **bridging inference** to create an antecedent for *beer* (Clark, 1975). These observations are very similar to ones I have already discussed. Given vs. new information is the basis for the identifiability requirement of definiteness, and bridging inferences are the way in which Clark describes the process of accommodating the presuppositions that allow the creation of a new discourse referent, when context serves to make it identifiable.

In both previous studies, we have seen that accommodating new information about either existing discourse referents or inferrable discourse referents comes with a processing penalty, but is this always the case? O'Brien, Shank, Meyers, and Rayner (1988) published a study in which they investigated whether an elaborative inference could occur before an entity was explicitly introduced. They defined an **elaborative inference** as one that expands or further defines explicitly introduced information without being necessary for comprehension. In terms of accommodation, this means that an elaborative inference is a presupposition that is accommodated before it is necessary to do so. To see if such elaborative inferences occur, the conducted a number experiments using both self-paced reading and eye-tracking.

In one such experiment, they used eye-tracking to determine the processing difficulty of various sentences. In this form of eye-tracking, participants are made stable by a bite-bar, and an invisible laser that tracks movement is shone into their eye while they read text on a computer screen. This allows the experimenter to know exactly where the participant is looking in a sentence at a certain time. This design avoids much of the experimental ambiguity experienced in self-paced reading studies because it allows you to track precise expenditures of processing-time allocation within the text.

To test for elaborative inferencing, O'Brien et al created a number of passages in which they manipulated the variables of inferential context and explicit mention to see when participants would create an elaborative inference:

24)

> All the mugger wanted was to steal the woman's money. But when she screamed, [stabbed][assaulted] her with his (knife/weapon) in an attempt to quiet her. He looked to see if anyone had seen him. He threw the knife in the bushes and ran away.

Each passage had four basic versions:

25)

a. High-context/ explicit, "stabbed her with a knife."
b. High-context/ implicit, "stabbed her with a weapon."
c. Low-context/ explicit, "assaulted her with a knife."
d. Low-context/ implicit, "he assaulted her with a weapon."

They found no significant differences in processing times between explicit and implicit mention when context was high (*stabbed*). However, reading times were significantly longer in the low-context condition (*assaulted*) when the antecedent had not been explicitly mentioned. These results indicate that readers do in fact generate elaborative inferences under highly constrained contexts, such as those caused by highly prototypical verb-instrument relations. Surprisingly, these elaborative inferences seem to allow new, unmentioned entities to act as if they have explicit antecedents.

In terms of accommodation in the high-context/ implicit condition, each anaphor still has an explicitly mentioned discourse referent to link up to, created by a more general antecedent. In the example above, the introduction of *weapon* forces us to create a new referent at the time of its introduction, so that all we need to do is to accommodate

a presupposition like, "The man's weapon is in fact a knife." If we follow the argument of O'Brien et al, we would conclude that this accommodation occurs in the form of an elaborative inference when the discourse referent first comes into being. In this way, the authors claim that certain highly-constrained contexts can force instantiation.

The Garrod and Sanford (1977) study finds a trend consistent with the low-context condition of O'Brien et al, where *knife* took longer to process when it had not been explicitly mentioned, meaning that although a possible antecedent exists in the low-context/implicit condition, linking accommodation must occur when the anaphor is introduced, slowing processing time at that point.

The findings of O'Brien et al seem at first to go against those of Haviland and Clark (1974), but these differences can be accounted for. In the study conducted by Haviland and Clark, the bridging anaphor had no previously introduced discourse referent, so that the penalty in processing time could be accounted for by the need to create a new one. In the study done by O'Brien et al, a discourse referent already exists. The effect information status has on processing times should be clear by this point, but where exactly do these effects come from?

Hess, Foss, and Carroll (1995) conducted a study in which they looked at the sources of context-effects on discourse processing. Early work, (Meyer & Schvaneveldt, 1971; Foss, 1982) showed a clear facilitation effect in reading times when a particular word was preceded by a related word or words. This was shown in lexical decision tasks, where the time it took to process and name a particular word was timed. When a certain word (such as *fish*) was preceded by a related word (such as *gills*), processing time of the target word improved, when compared to when the target word was preceded by an unrelated word (such as *color*), or no word at all. This facilitation effect is known as **semantic word priming**.

Hess, Foss, and Carroll discuss three major models that account for these facilitation effects. The first focuses on the word-to-word facilitation effects from within the lexicon. It predicts that priming effects should occur regardless of context. The second type of model focuses on discourse effects on facilitation. It predicts that priming should only occur with associations that fit the context. The third type of model is a hybrid model, that claims facilitation occurs from some combination of both. The authors claim that a discourse model of priming can account for all priming effects.

Hess, Foss, and Carroll conducted a number of studies in which they tested their hypothesis. They presented their material auditorily. At some point during the discourse, the audio ceased, and they presented the word that finished the passage they were just listening to on a computer screen and timed the amount of time it took participants to name the word. They concluded that faster naming times in one context over another were due to facilitation effects.

In order to determine the nature of facilitation effects, the authors manipulated the variables of globally related/ unrelated vs. locally related/unrelated. They then went through a number of studies to determine if facilitation in word-naming was found if only local context supported inferencing the word, if only global context supported inferencing, or if both global and local context were required.

They found that global context was the only reliable source of facilitation effects, and that local context had little-to-no effect if it disagreed with global context. What do these findings say about accommodation of presupposition with the definite article?

They indicate that in order for a certain presupposition to be deemed worthy of accommodation, the global discourse context must make it a member of a limited inferrable set. If membership in this set is well-established, then facilitation effects will be seen.

To summarize this section, the way in which we accommodate the presuppositions that allow us to treat new information as given shows up in the time it takes us to process discourse. We have seen different processing effects in accommodating both new discourse referents and new information about existing discourse referents. Incorporating a new discourse referent appears to create a processing penalty, as does accommodating new information about an existing discourse referent unless there is a highly constrained context, when elaborative inferences are possible. Additionally, we have seen that these differences in processing times are the result of global context created by a discourse model, and that acceptable accommodations are judged dependent upon that model.

In the next section, I will give an overview of my own research, which investigates the boundaries of facilitation effects for accommodation. Is accommodation all-or-none, on-off procedure? Or does the discourse model create gradations of acceptable inference that will show up in processing time?

Section VII: Experiments

Experiment 1

In the research in the last section, we saw that context effects create differences in processing time between discourse old and discourse new information. Previously unmentioned discourse referents take longer to process than ones that have already been mentioned. The exception to this occurs in highly constrained contexts in which specifying information can be elaborately inferred about a pre-existing discourse referent. This is because the elaborative inference primes the subsequent anaphor. The entity is presupposed, so there is no need for accommodation at the point of its introduction.

Let us assume that a given linguistic context produces a set of potentially inferrable entities. It is these entities, and only these entities, that we will consider potentially available through accommodation in that context. If there is mention of an entity that is relevant but previously un-introduced, then we incorporate it into discourse by accommodating a presupposition about it at the time of its introduction. Look at the context created in the sentence below:

26) Joe swept the office.

Entities that we might infer from this context include things we might associate with a person named Joe, the action of sweeping, or the location of an office. So we could infer new entities like *executive* or *dustpan*, but probably not entities like *Apollo* or *chinchilla*. I chose to investigate processing times in terms of relative ease of inferrability. Just as some category members are more highly associated with a category than others (Garnham, 1987), I believe that with a given context, some entities are more highly inferrable than others. Typicality effects in processing time have been shown

within categories; I predicted that similar effects would show up with inferrable entities and context. The harder an entity is to infer, the longer it should take to accommodate the relevant presuppositions.

This perspective on inference is similar to one proposed by Schwanenflugel & LaCount (1988) that uses **feature-generation mechanisms** to explain facilitation effects. They claimed that when processing a sentence, the features of upcoming words are restricted by the context created by previous words in the sentence, meaning the more specific the sentence, the smaller the scope of the set of highly inferrable entities.

Instead of trying to study inferencing in every situation new information is added, I chose to look specifically at instances where the definite article was used to introduce new information. The identifiability required for definiteness makes it easy to determine what inferences have occurred within a given context; If the definite article is used to designate an entity that has not been explicitly introduced by previous context, then we must infer that the entity has a specific discourse referent, and accommodate presuppositions to match this inference. The harder this is to do, the longer it should take to process.

I decided to test this hypothesis in a controlled setting by generating a list of two-sentence pairs in which the target inferred entity remained the same, and preceding context varied. I predicted that some contexts would allow participants to process the target entity more quickly than others. Take for example:

27) When he finally got up the nerve, Joe proposed to Sally. **The ring** had been his grandmother's.

28) When he finally got up the nerve, Joe spoke to Sally. The ring had been his grandmother's.

I predicted that *the ring* would be easier to incorporate into discourse in (27) compared to (28), because the overall facilitation effect was greater due to the fact that *ring* was more closely related to *proposed* than to *spoke*, which would make it easier to accommodate the necessary presupposition "Joe gave Sally a ring when he proposed." I predicted that this difference would show up in self-paced reading.

Method

Participants

A total of 48 people participated in my study (the results of four participants were discarded due to a programming error). The self-paced reading task was given at two locations on the Umass Amherst campus, in South College 401 and Tobin 203. Participants were psychology and linguistics majors who received extra credit in their classes.

Materials

Online testing occurred on a Windows computer running E-prime. Twenty-four items were used in the study. Each item had four different versions, each a combination of the two variables I manipulated; Inferrability (supported vs. unsupported) and Discourse Status (previously mentioned vs. unmentioned). In condition A, context made

the target entity easy to infer though it had not been previously mentioned. In condition B, context made the target entity easy to infer and it was also previously mentioned. In condition C, the entity was harder to infer and it had not been previously mentioned. In condition D, the entity was harder to infer, but it had been previously mentioned. Here is an example item (a full list of the items is included in appendix A):

28)
a. Clark swept the carpet thoroughly after the office closed. The broom was getting old.
b. Clark swept the carpet with a broom after the office closed. The broom was getting old.
c. Clark cleaned the carpet thoroughly after the office closed. The broom was getting old.
d. Clark cleaned the carpet with a broom after the office closed. The broom was getting old.

I predicted that the final sentence of condition A would take less time to process than that of condition C, and that B and D would be about the same, due to the differing facilitation effects acting on *the broom*.

Each item consisted of a setup sentence that differed in each condition, followed by a simple sentence introducing an entity with the definite article that was the same in each. Items were of two different types, depending on how the inference was created by the context. In one type, a verb-instrument relationship was used (swept>broom vs. cleaned>broom). In the other, the relationship was based on a stereotypical situation and the typical contents of that situation (Christmas dinner> turkey vs. dinner> turkey).

For the Inferrability variable, unsupported context was created by supplanting the verb or situation used in the supported context with a more general verb or situation. This was important because it maintained plausibility. The set of possibly inferable entities created by the altered sentence still contained the target entity as a member, so if there was no gradation in the facilitation effect for different entities, then both supported and unsupported versions would have been processed the same.

The conditions of each item were sorted so that each participant read one condition at random for each item, and six of each type overall. Items were combined with 86 test items from three other studies, to act as controls and to avoid participant bias.

I normed my materials using the judgments of several cohorts, as well as having conditions evaluated by four acquaintances using a simple rating scale of 1-5 (1 being worst and 5 best). In all cases, they agreed that condition A was easier to comprehend than condition C (average ratings given of 4.1 and 2.5, respectively).

Procedure

Participants were sat in front of a computer and told to begin the study at any time by pressing the space bar. They were read instructions and given several practice trials. They were asked to read the sentences as quickly and accurately as possible. Each sentence was broken up into segments, and the time it took participants to read each segment was recorded. For each item, first an asterisk was flashed on a screen, then a series of dashes appeared, each dash corresponding to a letter of the sentence pair. When the participant pressed the spacebar, the first segment of the item would appear in place of the dashes that held its place on the screen. After reading the first segment, the participant would press the spacebar to receive the next segment. Segments within items

were always of the same approximate length, and the last sentence was always its own segment.

Results

After collecting results, I conducted two separate repeated-measure ANOVAs, one within-participants and one within-items. For within-participants, Discourse Status was significant, (1,44)=10.622; p=.002. Neither Inferrability, (1, 44)= .402; p= .529, nor its interaction with Discourse Status, (1,44)=1.527; p=.223 were significant. For within-items, Discourse Status (1, 24)=9.158; p=.006 was significant, and the interaction was marginally significant, (1,24)=3.254; p=.084. Inferrability remained insignificant (1,24)=.377; p=.545. Overall, unmentioned conditions took significantly longer to read than mentioned conditions (1468 ms to 1356 ms, respectively), and there was some evidence of an interaction. The pattern of means between conditions reflected this; Condition A (supported-unmentioned)=1436ms, Condition B (supported-mentioned)=1372ms, Condition C (unsupported-unmentioned)=1501ms, and Condition D (unsupported-mentioned)=1340.

Discussion

Contrary to my predictions, only a marginal effect was found in the interaction between Inferrability and Discourse Status, although a highly significant difference in reading time was found between sentences containing a previously mentioned entity and those that contained an unmentioned one. The difference in Discourse Status might at first seem to contradict the results of O'Brien et al, where they found that discourse-new entities could be processed as quickly as discourse-old entities when context could license an elaborative inference. On closer examination, however, it can be seen that the majority of my items lacked pre-established discourse referents.

In the O'Brien et al studies, in each case context went from a general category to a specific member of that category, but the introduction of the general category created a discourse referent which the more specific category member was hooked up to. In my study, however, only five out of my twenty-four items had such pre-established discourse referents (items 10, 13, 15, 18, and 22 in appendix A). So when compared against explicit mention in conditions B and D as a base-line, not only was the participant forced to accommodate presuppositions in A and C, he was also forced to create new discourse referents at the point of introduction as well. This could account for the highly significant difference between mentioned and unmentioned context.

Although I predicted a larger interaction between my variables, there was a major flaw in my study that may explain my weak results; Questions about the text being read are often included in self-paced reading studies to make sure participants are actually reading carefully, but I failed to supply questions in my study. This factor paired with the fact that my study was conducted in the last several weeks of the fall 2004 semester (participants are notoriously hard to garner good results from at the end of a semester) can probably account for my marginal results. Another factor I identified was the fact that my target sentences were all relatively simple and easy to process, which may have allowed my participants to overcome some of the difficulty they would have otherwise have experienced in condition C. Nonetheless, I would still expect that if I re-ran this

study in the middle of a semester with comprehension questions, my results would be much clearer.

However, the fact that there was some indication of interaction is still heartening, and correlates well with the results of my norming study. My results seem to indicate, albeit weakly, that when accommodating new entities with the definite article, context seems to facilitate some entities more than others, even when all the options are still considered felicitous. In terms of the Spectrum of Accommodation, it appears that while discourse-new entities may take the same amount of time to process as discourse-old entities when context supports an elaborative inference, entities that lack that context may still show varying degrees of facilitation from previous context. It does seem that context allows some entities to be processed more quickly than others, based on how easy it is to accommodate the necessary presuppositions. I have tied these results together to the point where I feel I can formalize the Spectrum of Accommodation:

> When incorporating a previously unmentioned entity designated by the definite article, the amount of time it takes to accommodate the presupposition required to assign definiteness is dependent upon the typicality of the inference, given the context.

The harder it is to assign definiteness to an entity potentially inferrable by a given discourse, the longer it will take to process. I will attempt to trace the implications of this hypothesis with the results of Experiment 3, but in the following experiment, I have taken a step back and looked more closely at the requirements of the definite article in terms of identifiability and uniqueness, to better understand the scope within which my Spectrum might apply.

Experiment 2

.

In section II, I described uniqueness, inclusiveness, and identifiability. I claimed that, in opposition to the claim made by Lyons, inclusiveness should be seen as a natural repair for uniqueness when dealing with plural sets, and that definiteness could be given a more unified account by using a discourse model. Uniqueness is merely a way in which identifiability can be satisfied.

Lyons tells us that the need to think of uniqueness separate from identifiability comes from dealing with examples like (8), where we do not seem to know anything specific about *the bride.* How could it be that she is identifiable in this context? This is only a problem if we expect *bride* to link directly to the world, instead of to a discourse model. It is actually a perfect example of identifiability if we think in terms of discourse referents.

Indentifiability says the definite article indicates that the hearer can identify the referent from the context, meaning that the hearer can either connect the entity to an existing discourse referent or has reason to infer a new discourse referent. In (8), *wedding* makes *bride* an easily inferrable entity. When she is introduced in the following sentence, it is easy to accommodate her and to create a new discourse referent at the point

of introduction. When an entity is unique, makes it easy for the hearer to identify it by linking it to a specific discourse referent.

From this point of view, it matters very little that no "real" information has been supplied for the sake of identifiability. As long as you can hook an entity up to a discourse referent, it is identifiable. Using a discourse model of identifiability is also useful in describing a common observation; People are actually very good at talking in great detail about things they know nothing about. Building a linguistic system in this way makes it easy to acquire new knowledge. Eavesdropping on a conversation is a good example of when this might apply. Just keep building up information about some abstract discourse referent, until at some point you receive a piece of information that hooks your abstract discourse model up to a real-world entity. Everyone has probably had the experience of listening to an unfamiliar group of people who all know one another, talking around you about people you know nothing about. If you are interested in eventually understanding what's going, you keep track of mystery referents you know nothing about.

I claimed earlier that we should think about uniqueness and not inclusiveness. The reason for this is the way in which we can treat definite entities that are part of a plural set. While inclusiveness claims you must use the definite article to refer to the entirety of a relevant set of entities, it is in fact possible to single out one entity from a plural set, if context supports this narrowing of context. This claim is a reformulation of both identifiability and uniqueness called the **Small World hypothesis**, which states:

> In order for a new single entity to be introduced with the definite article, the scope of its frame of reference must be small enough to single that entity out as more relevant than all others of its type. In this way, an entity designated by the definite article must be *locally unique* within its frame of reference.

Inclusiveness, then, need only come into play if this singling-out procedure is not used, and then it is sufficient to say that the speaker then has to refer to the entire set of fully relevant entities, in order to still use the definite article.

In contrast to the Small World hypothesis, another possibility for the requirements of the definite article is an alternate formulation of Identifiability. It claims that instead of needing to link to a specific discourse referent, an entity merely needs to be able to link to *some* previously introduced/inferred discourse referent, given the frame of reference, and that uniqueness is unnecessary. This view is summarized by the **Inferrability Hypothesis**, which predicts that as long as an entity is relevant to the frame of reference, its presuppositions can be accommodated, even if similar entities are equally relevant within that frame. This hypothesis is more closely tied to the original concept of familiarity, which stated that the definite article merely required that the entity be hearer-old, without any requirement of specificity or relative salience.

In a pilot study, I attempted to test the Small World hypothesis against the Inferrability hypothesis by creating scenarios in which the ongoing frame of reference made multiple entities of the same type relevant. I created a number of 2-sentence pairs, and had participants judge the quality of the continuation from the first sentence to the second, where a new entity was introduced with the definite article. Judgments were rated on a five-point scale.

In my small world experimental examples, context was given to single out one entity as more relevant than others. Context was created by manipulating either the initial sentence or the predicate that followed the critical DP. All my Small World examples were compared against corresponding Inferrability controls that did not create reason to pick out a single entity, but were the same in all other aspects (the critical DP referred to a relevant entity). I predicted that participants would consider these bad continuations and give them worse ratings than they gave my Small World examples, even though Inferrability would predict these to be good. Several examples follow:

(29)
Myrtle went to the dry cleaners to pick up her outfit. The shirt was stained.
*Myrtle went to the dry cleaners. The shirt was stained.

(30)
Todd finally found the bathroom stalls. The toilet was dirty.
*Todd finally painted the bathroom stalls. The toilet was dirty.

(31)
Lester looked at the coat rack. The peg was broken.
*Lester looked at the coat rack. The peg was wood.

(32)
Horace walked into the house. The window was open.
*Horace walked into the house. The window was glass.

Although my experiment garnered significant results, I realized that my design was flawed in several different ways. Participants did rate Small World items significantly higher than Inferrability items, but this may have been because many of the Inferrability items I created were infelicitous. Not only did they not contain context to single one entity out over others, but they may have failed to create otherwise coherent continuances. I decided to re-run the experiment with better items.

Method

Participants

A total of 20 people participated in my study. All were native speakers of English and students at Umass Amherst.

Materials

In order to corroborate my previous findings, I reran the experiment with new items that did not contain the problems of the original. My Small World items still gave reason to make one entity more salient than others of its type in a given context, but I changed my Inferrability items so that in each case, introducing multiple entities of the same type would have been perfectly felicitous. The only problem then, was singling out one entity when there were multiple equally salient ones. If all that is required to assign definiteness is Inferrability, then there should have been no difficulty singling out one

entity from the introduced set. Several of the rewritten items follow (a full list may be found in Appendix B):

(33)
Juan drove up to the busy tollbooths. The toll taker was rude.
*Juan looked at the busy tollbooths. The toll taker was rude.

(34)
Todd finally found the bathroom stalls. The toilet smelled horrible.
*Todd finally painted the bathroom stalls. The toilet smelled horrible.

(35)
Adam looked at his birthday card. The front was smudged.
*Adam looked at his birthday cards. The front was smudged.

(36)
Lyla opened the book. The page was ripped.
*Lyla flipped through the book. The page was ripped.

My items were of two types. In the first type, context for the plural set of equally plausible entities was created by the explicit introduction of a plural set of associated entities, as in (34). In this example, multiple bathroom stalls have been explicitly introduced, meaning that we can infer the existence of multiple toilets. In the second type, context for the set of equally plausible entities was created by the explicit introduction of a single entity that had a plural set of inferrable parts, as (36). In this example, a book has been explicitly introduced, allowing us to infer the existence of multiple pages.

In both types of items, context to single out one entity as more relevant that others of it type was dependent upon the actions taken by the subject of the first sentence. In (36), while opening a book may allow us to imagine a single, salient page (the page Lyla opens to), flipping through a book fails to make a single page salient.

Procedure

Each participant was given a questionnaire in which they were asked to judge the quality of continuation between a pair of sentences (see appendix B). Their instructions were as follows:

> In the following questions, you will be given pairs of sentences. Please indicate whether or not the second sentence of the pair is a good continuation of the scenario created by the first sentence. For each question, please circle a number from one to five, where one indicates a highly unnatural or awkward continuation, and five indicates a very natural continuation (the better the continuation sounds, the higher you should rate it).

Each participant was given only one version of each item, either the Small World or the Inferrability version. Items were counterbalanced between two versions of the questionnaire. Each version was taken by ten participants.

Results

After tabulating data, I conducted two paired-subject t-tests to determine the significance of my results. I compared the scores from my Small World condition against those of my Inferrability condition, both within participants and within-items. Results were highly significant: Items, $t(15)=8.06$; $p= 0.000$, Participants, $t(19)=7.17$; $p= 0.000$. Participants gave significantly higher ratings to Small World items than to Inferrability items (overall means 3.6 to 2.4, respectively). For each item and for each participant, the Small World condition showed a higher mean score than its Inferrability counterpart.

Discussion

These results show that native speakers of English consistently rate a continuation better when the frame of reference serves to make one entity more salient than others of its type. This suggests that Inferrability is insufficient when accommodating with the definite article, and supports my Small World hypothesis, because participants liked "the x" better when it was the most relevant x, and liked "the x" worse when other x's were equally relevant.

For inclusiveness, these results show that it is not necessary for the definite article to refer to the entirety of a set if context serves to single out one entity over others. They also indicate that in order for identifiability to give an accurate account of the data, it must include some treatment of specificity. It is insufficient for an entity designated with the definite article to merely be inferrable from the context; It must also be clear which discourse referent it attaches to, and if more than one candidate exists in the set of possibly inferable entities, additional context must serve to single out one candidate over others.

These results also indicate that we really use some form of situation model in our head to compute context and assign discourse referents. In (33), for example, the fact that we can infer a single, salient toll taker when we have already explicitly introduced multiple tollbooths means that the specific actions allow us to constrain the explicitly established frame of reference. In this case, we know that if Juan is driving up to a set of tollbooths, he will most likely stay in his lane and only interact with a single toll taker, making that toll taker salient. If we were unable to use situational context to zoom in on this smaller frame of reference, we would unable to talk about a single toll taker, because one would be inferrable from each tollbooth, and they would all be equally salient. This is exactly what occurs when in the Inferrability version, when Juan is merely looking at the tollbooths. The action of looking fails to single out one toll taker, since multiple toll takers remain equally inferrable. This need to create local uniqueness by pairing-down context until only one entity of a single type is salient is the requirement placed on the deifinite article by the Small World.

The Small World can be directly connected to the Spectrum of Accommodation. These results tell us that when assigning definiteness to a single entity, processing times should be dependent upon both degree of relevance to the context, and relative degree of

salience, especially when there are multiple entities of the same type within the initial context. The harder it is to determine either relevance or local uniqueness, the harder it should be to process the accommodation. If either of these criteria are too hard to meet, using the definite article to single out one entity over others will be infelicitous, and the speaker will have to introduce the referent by some other means, such as with the indefinite article. Determining the cut-off point at which it is no longer possible to accommodate an entity with the definite article was the purpose of the experiment that follows.

Experiment 3

The purpose of this study was to map the boundaries of the Spectrum of Accommodation. Given the Small World hypothesis, and the fact that all native speakers of English have clear intuitions about when and when not to use the definite article, I felt that I could draw some inferences about the general nature of those boundaries. I chose to test this using a fill-in-the-blank questionnaire, in which participants determined a discourse-new entity with either "a" or "the." The hope was to link this production data to my previous processing work and with the experimental literature to explain how different factors contribute to the processing of discourse-new information designated by the definite article.

The methodology of my questionnaire relied not only on my Small World account of definiteness, but on an account of indefiniteness, as well. I assumed that when context is unable to fulfill the local uniqueness required by the Small World, then discourse-new entities will be introduced by some other means. In either case, the new entity must be relevant (a member of the set of potentially inferable entities created by the context), but the definite article can only be used when context makes a single discourse referent available.

The indefinite article then, is used when the new entity is relevant, but cannot be matched up to any previously established discourse referent. A new discourse referent must be created at this time. So what is the difference between this use of the indefinite, and instances in which the definite article licenses the creation of a new discourse referent? While using the definite article indicates to the hearer that he *already* possess context to disambiguate the new discourse referent from others, using the indefinite article entails no such presuppositions, and instead *creates* the disambiguating context.

I used this distinction to determine what English speakers consider adequate context for inferring definiteness. I predicted that given the right setup, virtually all participants would consider a target entity definite from unmentioned prior context, and would designate it with the definite article. Given a different setup, I predicted that there would be some differences between participants, depending on the nature of their personal discourse model, and that they would be more stratified in their use of the definite versus the indefinite. Given another setup, I predicted that no-one would consider the target entity definite, and everyone would use the indefinite article instead. These predicted differences in production were due to a manipulation of local uniqueness. The details will be discussed below in materials.

Method

Participants

The same 20 people who participated in experiment 2 participated in this study as well (both studies were part of the same questionnaire).

Materials

Materials consisted of sixteen two-sentence pairs. The first sentence of each pair set up the context, while the second introduced a discourse-new entity which participants had to designate, either with the definite or indefinite article, by filling in a blank. Items were broken up into four different conditions, with four items in each condition. An example of each follows:

37)

a. When Joe proposed, he gave Sally a ring set with a large gem. __ diamond was nearly flawless.

b. Ron swept the office. __ broom was starting to wear out.

c. Helga lounged at the small private beach. __ lifeguard was yelling at some kids.

d. Camilla was in the library. __ book fell off the shelf as she walked past.

In Condition A, the target entity is both relevant and locally unique. In addition, the target entity refers back to a pre-established discourse referent, and would be expected to generate an elaborative inference if part of a processing study. Uniqueness was supplied by category-membership (a large gem/diamond) relations, in which a single entity of a more general category had already been introduced. Because the original category discourse referent was unique, its corresponding anaphor was unique as well. I predicted that for these items, participants would consistently choose the definite over the indefinite.

In Condition B, the target entity is again both relevant and locally unique, but there is no pre-established discourse referent, so one must be supplied at the point of its introduction. Uniqueness was fulfilled either by verb-instrument relations, in which the action of the verb implied a unique referent (swept- broom), or part- whole relationships, in which knowledge of the whole implied a unique part (car-steering wheel). Although Condition A and Condition B have been shown to have differences in processing time, when compared against explicitly discourse-old entities, I predicted that there would be no significant difference in terms of definiteness, and that participants would consistently choose the definite article in these items.

In Condition C, the target entity was always relevant to the context, but uniqueness was only satisfied if participants viewed the discourse in a certain way. In (c) above, *lifeguard* is a potentially inferable entity given what we know from *small private beach*. Even though the beach is small and private might bias the reader towards putting only one lifeguard in their mental model, nothing in the context actually forces them to do so; They could easily imagine multiple lifeguards. In this way, uniqueness could either be supported or unsupported. Because of this ambiguity, I predicted that

participants would be split on their use of the definite, although due to the nature of the items, I expected the indefinite article to be preferred overall.

In Condition D, the target entity was still relevant, but I intentionally manipulated context so that meeting local uniqueness for the target entity was practically impossible. I did this by making the target entity the member of a set in which multiple members were equally salient. I predicted that participants would consistently choose the indefinite article, because the only way to get a definite description out of the context was to use inclusiveness to apply to the entire set, and the target entity was only a member of that set. A full list of the items used may be seen in Appendix C.

Procedure

Each participant was given a questionnaire in which they were asked to fill in the blank with the appropriate discourse referent. It was a forced production task because they were asked to only use the definite or indefinite article, the reason for this being participants may have otherwise used possessive pronouns in some of the test items. Since there was no need for counterbalancing, each participant was given all sixteen items. Questionnaires were scored by tabulating number of definite article responses out of total responses for each condition of my spectrum, both by participant and by item.

Results

I conducted two 1-way repeated subject tests, one for participants and one for items. Both were highly significant: Within-participants, $(3,19)=364.217$; $p=.000$, within-items, $(3,15)=41.644$; $p=.000$. Within-participants, every condition differed significantly from each other condition ($p<.001$ after Bonferoni correction), except for conditions A and B ($p=.337$ after Bonferoni correction). Within-items, both conditions A and B differed significantly from conditions C and D, but there was no significant difference between A and B ($p=1.0$ after Bonferoni correction), or between C and D ($p=.169$ after Bonferoni correction). The pattern of means was consistent with these results (average number of definite article uses out of twenty); Condition A= 18.5, Condition B=20, Condition C=5.25, Condition D=.025.

Discussion

The items of each condition went almost exactly as I predicted. There was no significant difference between conditions A and B, even though a difference showed up in processing time between these two types of examples in experiment 1. In condition C, participants chose either the indefinite or the definite article, although the indefinite was strongly preferred. In condition D, where I predicted that the definite article was completely unsupported, only one participant in one condition chose the definite article. The only item that did not go the way that I predicted was one from condition A:

38) After stabbing the victim, the assailant threw away his weapon and fled. __ knife was found in the bushes.

For this item, only fifteen out of the twenty participants chose the definite article, while almost everyone chose the definite article for each item in both conditions A and B.

This was at first surprising, until I realized that the item could be read as a factual account of a crime scene. In the case of crime scenes, there is a conscious attempt to avoid casual inference, since it could be thrown out of court as "grounds for speculation." Some of my participants may have interpreted the sentence in those terms, and intentionally cancelled out the natural inference that would otherwise occur.

In summery, when asked to supply a determiner for a discourse-new entity, people's willingness to use the definite article is directly dependent upon how well the context manages to single out the entity from the set of possibly inferable entities in the discourse model. These production results match the felicity judgments of experiment 2, and tie in with the processing effects seen in experiment 1. Given the gathered evidence, I can introduce the Spectrum of Accommodation in its full form.

The Spectrum of Accommodation is a model claiming that with the definite article, the presuppositions necessary to treat discourse-new entities as discourse-old ones are accommodated with relative ease or difficulty, depending upon how well context meets the requirements of the Small World hypothesis. To put it simply, we have a hard time incorporating discourse-new information with the definite article when discourse context fails to single out an unambiguous discourse referent, either pre-established or inferrable, to link it to within our discourse model.

Good context to infer and accommodate the necessary presuppositions for possibly inferrable entities is created by strong verb-instrument relations, part-whole relationships, general-specific category relations, and the relation between typical members and their stereotypical scenarios. In each of these cases, if multiple entities of the same type are possibly inferrable from the context, then additional sub-context must exist to single one out over others of its type, or all relevant entities of the same type will be treated as definite as per inclusiveness.

These claims are supported by both felicity judgments and forced production, where harder accommodations are either rated poorly, or introduced with the indefinite article. The processing data seems to support this spectrum as well, evidence suggesting that harder accommodations take more time to process. The relationship between felicity and processing time is not in a one-to-one correspondence, however, because in some contexts, the need to create a new discourse referent creates an additional penalty in processing time.

Section VIII: General Discussion

In this thesis, I have attempted to give a history of the topic both from a semantic and psycholinguistic perspective. The traditional distinctions between uniqueness/ inclusiveness and familiarity/identifiability, combined with the concept of discourse referents, allowed me to motivate the Small World hypothesis, in which I claim that for a new entity to be introduced with the definite article, the scope of its frame of reference must be small enough to single the entity out as more relevant than all others of its type.

I used this simplified account of the definite article to set the boundaries of the Spectrum of Accommodation, in which I claim that our ability to accommodate new information is dependent upon our ability to supply the target entity with an unambiguous discourse referent, which is in turn dependent upon the nature of our context. I have kept the scope of this hypothesis intentionally small, but if correct, these results have many

interesting implications in terms of discourse structure, both for other English referring devices and cross-linguistically.

It appears that the way in which we treat given and new information is fluid, with certain relationships between words (verb-instrument, etc…) allowing us to talk about new entities like old ones as long as we can match them up to preexisting discourse referents or infer new ones. I have shown that with the definite article, the system of accommodation that allows this to occur meets the Small World requirements of definiteness by varying degrees of success, depending upon the context.

Although I only talked about the Spectrum of Accommodation in terms of the definite article, it is likely that the same type of graded acceptability should apply to other referring devices that rely on the given-new distinction for their use. For example, it is sometimes possible for a pronoun to single out an implicit antecedent in English. Look at the following:

39) He picked out his donut and ate **it**.
40) She picked her nose and ate **it**.
41) #She picked her teeth and ate **it**.

In (39), the pronoun clearly refers to *donut*, but in (40), it does not refer to *nose*. Why is this the case? Context allows us to link the pronoun to *booger*, an antecedent that has never been explicitly introduced. The fact that this implicit antecedent is acceptable in (40), but not in (41) suggests that the requirements that make (40) felicitous may be satisfied by varying degrees, given the context. This parallels the Spectrum of Accommodation laid out for the definite article. The requirements may be more constrained for pronouns, since context must serve not only to single the implicit antecedent out other possibly inferrable entities, but must cancel out any possible explicit antecedents, as well. It is likely that if investigated, that acceptability for different implicit antecedents would be rated on a graded spectrum, and that these differences might show up in processing time.

Another area where the Spectrum of Accommodation might be useful is in describing certain cases of scrambling in Japanese, where speakers rearrange the order of their sentences to put given information before new information. Preliminary data collected last semester seems to indicate that Japanese speakers are willing to treat highly inferrable entities as given information by fronting their corresponding constituents, even though these constituents have never been explicitly introduced:

42) Joe proposed to Mary and gave her **the ring**.
43) Joe asked Mary to be his valentine and gave her **the flowers**.

In translations of the above examples, Japanese speakers will treat *ring* as given, but not *flowers*. Again, this indicates that since discourse-new information may be treated as old in some cases but not in others, Japanese scrambling might best be described by some form of the Spectrum of Accommodation. Although Japanese speakers are unwilling to front *flowers* as given information, English speakers are perfectly willing to treat them as discourse-old by determining them with the definite article This suggests that even though accommodation may occur in Japanese, the scope

of acceptable accommodations may be more constrained than that of English. The definite article may allow more liberal accommodations than scrambling.

Overall, this thesis is intended merely as a sketch of these issues, a basic outline of a perspective that has the potential to make sense of a large amount of traditionally unrelated data. With more research, it should be possible to determine the nature of the precise relationship between felicity and processing when accommodating discourse-new entities determined by the definite article. Much of my data is inferential, but many of these connections could be made explicitly. One possible future study would be to put all four conditions of experiment 3 together in a self-paced reading study, to determine in what way processing times correlate with felicity judgments on a broader range of the Spectrum than was tested in experiment 1.

One of the most interesting things about the system outlined by my Spectrum of Accommodation is that it not only makes it *possible* to use the definite article when its referent has not been explicitly introduced, but that it makes it *preferable*. Using the definite article with discourse-new information is even necessary within certain contexts, as conditions A and B showed in experiment 3. This observation has interesting implications about the way we construct our linguistic descriptions of the world.

If we tie this in with the studies of Hess, Foss and Carroll (1995), where overall discourse context had an immediate, online effect on facilitation, we get a picture of discourse in which our mental model keeps constant track of each discourse referent. In addition, the model seems to track potentially inferrable entities generated by the associations within a particular context, so that when a new constituent is added, it can be compared against this set of possible inferences to determine whether it should be treated as given or new.. As context is added to discourse, this set of possibly inferrable entities becomes constrained, as described by Schwanenflugel & LaCount (1988) with their feature generation mechanisms. The Small World functions in this way, because context must be narrowed to the point where a single entity is salient over others of its type.

The definite article, then, is a marker that allows the speaker indicate to the hearer how he should navigate the morass of association created by discourse content. By entailing the presupposition that its referent is identifiable within previous context, the definite article allows the hearer to sort back through all the possibly inferrable entities created in the discourse model, and pick out one both relevant and unique within the frame of reference. This process is necessary to satisfy the requirements of the Small World, and the ease with which it can occur within a specific context is the Spectrum of Accommodation.

Using the Small World and the Spectrum of Accommodation to look at how the definite article interacts with discourse-now information resolves the traditional discontinuity between uniqueness and familiarity, makes sense of existing processing data in terms of graded acceptability, and creates a powerful framework with which to analyze the impact of specific situations on information status.

References:

Garnham, A. (1987). Mental Models as representations of Discourse and Text. West Sussex: Ellis Horwood Limited.

Haviland, S., Clark, H. (1974) What's new? Acquiring new information as a process in comprehension. Journal of verbal Learning and Behavior, 13, 5, 512-521

Hess, J., Foss, D., &Carroll, P. (1995). Effects of Global and Local Context on Lexical Processing During Language Comprehension. Journal of experimental Psychology, 1, 62-82.

Lewis, D. (1979). Scorekeeping in a Language Game. Journal of Philosophical Logic, 8, 339-359.

Lyons, C. (1999). Definiteness. Cambridge: Cambridge University Press.

O'Brien, E., Shank, D., Meyers, J., & Rayner, K. (1988). Elaborative Inferences During Reading: Do They Occur On-Line? Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(3), 410-420.

Pinker, S. (1994). The Language Instinct. New York. Harper Collins Publishers.

Appendix A: Experiment 1 Test Items

1.

a. When he finally got up the nerve,/ Joe proposed to Sally./ The ring had been his grandmother's.

b. When Joe proposed to Sally,/ he gave her a diamond ring./ The ring had been his grandmother's.

c. When he finally got up the nerve,/ Joe spoke to Sally./ The ring had been his grandmother's.

d. When Joe spoke to Sally,/ he gave her a diamond ring./ The ring had been his grandmother's.

2.

a. John's hockey team returned home on the bus/ after winning the finals./ The trophy was made of silver and intricately engraved.

b. John's hockey team returned home with a trophy/ after winning the finals./ The trophy was made of silver and intricately engraved.

c. John's hockey team returned home on the bus/ after playing the game./ The trophy was made of silver and intricately engraved.

d. John's hockey team returned home with a trophy/ after playing the game./ The trophy was made of silver and intricately engraved.

3.

a. The anxious golfer searched/ by a stream in the deep woods./ The golf ball was nowhere to be found.

b. The anxious golfer searched/ for a golf ball in the deep woods./ The golf ball was nowhere to be found.

c. The anxious detective searched/ by a stream in the deep woods./ The golf ball was nowhere to be found.

d. The anxious detective searched/ for a golf ball in the deep woods./ The golf ball was nowhere to be found.

4.

a. Linda tried to read/ while she waited in the dentist's office./ The magazine was really old.

b. Linda read a magazine/ while she waited in the dentist's office./ The magazine was really old.

c. Linda tried to read/ while she waited to take her Psych exam./ The magazine was really old.

d. Linda read a magazine/ while she waited to take her Psych exam./ The magazine was really old.

5.

a. Cameron was hungry,/ so she put money/ into the vending machine./ The candy bar cost a dollar.

b. Cameron wanted a candy bar,/ so she put money/ into the vending machine./ The candy bar cost a dollar.

c. Cameron was hungry,/ so she gave money/ to her friend./ The candy bar cost a dollar.

d. Cameron wanted a candy bar,/ so she gave money/ to her friend./ The candy bar cost a dollar.

6.
a. They always cooked as a family/ when Harry ate Christmas dinner/ at his mother's house./ The turkey was a little dry this time.
b. They always had a turkey/ when Harry ate Christmas dinner/ at his mother's house./ The turkey was a little dry this time.
c. They always cooked as a family/ when Harry ate dinner/ at his mother's house./ The turkey was a little dry this time.
d. They always had a turkey/ when Harry ate dinner/ at his mother's house./ The turkey was a little dry this time.
7.
a. Nicole gave Phil a big hug/ when he won the race./ The gold medal was his first.
b. Nicole gave Phil a gold medal/ when he won the race./ The gold medal was his first.
c. Nicole gave Phil a big hug/ when he won the scholarship./ The gold medal was his first.
d. Nicole gave Phil a gold medal/ when he won the scholarship./ The gold medal was his first.
8.
a. Kate asked for a table/ when she went to the restaurant/ where Adam worked./ He brought her the menu.
b. Kate asked for a menu/ when she went to the restaurant/ where Adam worked./ He brought her the menu.
c. Kate asked for a table/ when she went to the pool hall/ where Adam worked./ He brought her the menu.
d. Kate asked for a menu/ when she went to the pool hall/ where Adam worked./ He brought her the menu.
9.
a. Mike was hungry,/ so he called Dominos./ A new driver delivered the pizza to Mike.
b. Mike wanted a pizza,/ so he called Dominos./ A new driver delivered the pizza to Mike.
c. Mike was hungry,/ so he called a restaurant./ A new driver delivered the pizza to Mike.
d. Mike wanted a pizza,/ so he called a restaurant./ A new driver delivered the pizza to Mike.
10.
a. It was Halloween,/ and the Jones household put decorations/ on their porch./ The pumpkin had not been carved.
b. It was Halloween,/ and the Jones household put a pumpkin/ on their porch./ The pumpkin had not been carved.
c. It was fall,/ and the Jones household put decorations/ on their porch./ The pumpkin had not been carved.
d. It was fall,/ and the Jones household put a pumpkin/ on their porch./ The pumpkin had not been carved.
11.
a. Clark swept the carpet thoroughly/ after the office closed./ The broom was getting old.
b. Clark swept the carpet with a broom/ after the office closed./ The broom was getting old.

c. Clark cleaned the carpet thoroughly/ after the office closed./ The broom was getting old.
d. Clark cleaned the carpet with a broom/ after the office closed./ The broom was getting old.
12.
a. Cindy stopped going to the class/ of a subject she didn't like./ The professor wouldn't answer her questions.
b. Cindy stopped going to the class/ of a professor she didn't like./ The professor wouldn't answer her questions.
c. Cindy stopped going to the website/ of a subject she didn't like./ The professor wouldn't answer her questions.
d. Cindy stopped going to the website/ of a professor she didn't like./ The professor wouldn't answer her questions.
13.
a. Horace forgot/ he was cooking food for breakfast./ The omelet was burnt.
b. Horace forgot/ he was cooking an omelet for breakfast./ The omelet was burnt.
c. Horace forgot/ he was cooking food for dinner. / The omelet was burnt.
d. Horace forgot/ he was cooking an omelet for dinner./ The omelet was burnt.
14.
a. The lumberjack whistled/ as he chopped away cheerfully/ in the afternoon./ The ax had just been sharpened.
b. The lumberjack whistled/ as he chopped with an ax/ in the afternoon./ The ax had just been sharpened.
c. The chef whistled/ as he chopped away cheerfully/ in the afternoon./ The ax had just been sharpened.
d. The chef whistled/ as he chopped with an ax/ in the afternoon./ The ax had just been sharpened.
15.
a. Carl saw a motionless shape/ when he went to refill the bird feeder./ The sparrow looked injured.
b. Carl saw a sparrow/ when he went to refill the bird feeder./ The sparrow looked injured.
c. Carl saw a motionless shape/ when he went to refill the ice-cube tray./ The sparrow looked injured.
d. Carl saw a sparrow/ when he went to refill the ice-cube tray./ The sparrow looked injured.
16**.**
a. A mishap/ got the carpenter off work/ early today./ The tablesaw had broken.
b. A tablesaw/ got the carpenter off work/ early today./ The tablesaw had broken.
c. A mishap/ got the man off work/ early today./ The tablesaw had broken.
d. A tablesaw/ got the man off work/ early today./ The tablesaw had broken.
17.
a. Janine scrubbed the pan/ with a scowl after finishing dinner./ The brush was wearing out.
b. Janine scrubbed the pan/ with a brush after finishing dinner./ The brush was wearing out.

c. Janine washed the pan/ with a scowl after finishing dinner./ The brush was wearing out.
d. Janine washed the pan/ with a brush after finishing dinner./ The brush was wearing out.
18.
a. After that last game,/ the baseball team became fed up/ with their player./ They traded the pitcher to another team.
b. After that last game,/ the baseball team became fed up/ with their pitcher./ They traded the pitcher to another team.
c. After that last game,/ the team became fed up/ with their player./ They traded the pitcher to another team.
d. After that last game,/ the team became fed up/ with their pitcher./ They traded the pitcher to another team.
19.
a. Bess went to a frat party/ with a friend./ The keg ran out early.
b. Bess went to a frat party/ with a keg./ The keg ran out early.
c. Bess went to a dinner party/ with a friend./ The keg ran out early.
d. Bess went to a dinner party/ with a keg./ The keg ran out early.
20.
a. Fred went/ to Kathleen's birthday party with friends./ He gave the card everyone had signed to her.
b. Fred went/ to Kathleen's birthday party with a card./ He gave the card everyone had signed to her.
c. Fred went/ to Kathleen's house party with friends./ He gave the card everyone had signed to her.
d. Fred went/ to Kathleen's house party with a card./ He gave the card everyone had signed to her.
21.
a. Gunter was shopping/ in a supermarket with his son./ The shopping cart had a squeaky wheel.
b. Gunter was shopping/ in a supermarket with a shopping cart./ The shopping cart had a squeaky wheel.
c. Gunter was shopping/ at a flea market with his son./ The shopping cart had a squeaky wheel.
d. Gunter was shopping/ at a flea market with a shopping cart./ The shopping cart had a squeaky wheel.
22.
a. Always a romantic,/ Carter picked a bouquet/ for his girlfriend, Eunice./ He gave the flowers to her.
b. Always a romantic,/ Carter picked flowers/ for his girlfriend, Eunice./ He gave the flowers to her.
c. Always a romantic,/ Carter bought a present/ for his girlfriend, Eunice./ He gave the flowers to her.
d. Always a romantic,/ Carter bought flowers/ for his girlfriend, Eunice./ He gave the flowers to her.
23.

a. Ralph found a sweet deal/ on furniture.com and went for it./ UPS delivered the chair to him.
b. Ralph found a sweet chair/ on furniture.com and went for it./ UPS delivered the chair to him.
c. Ralph found a sweet deal/ at a yard sale and went for it./ UPS delivered the chair to him.
d. Ralph found a sweet chair/ on furniture.com and went for it./ UPS delivered the chair to him.
24.
a. The soccer player chased/ an opponent across the field./ The soccer ball went out of bounds.
b. The soccer player chased/ a soccer ball across the field./ The soccer ball went out of bounds.
c. The defense player chased/ an opponent across the field./ The soccer ball went out of bounds.
d. The defense player chased/ a soccer ball across the field./ The soccer ball went out of bounds.

Appendix B: Experiment 2 Test Items

1.
Myrtle went to the dry cleaners to pick up an outfit. The shirt was torn.
*Myrtle went to the salvation army to pick out an outfit. The shirt was torn.
2.
Todd finally found the bathroom stalls. The toilet smelled horrible.
*Todd finally painted the bathroom stalls. The toilet smelled horrible.
3.
Juan drove up to the busy tollbooths. The toll taker was rude.
*Juan looked at the busy tollbooths. The toll taker was rude.
4.
Lester put his jacket on the coat rack. The peg broke under the weight.
*The whole class put their jackets on the coat rack. The peg broke under the weight.
5.
Meredith stepped up onto the ladder. The rung broke.
*Meredith was climbing the ladder. The rung broke.
6.
Lyla opened the book. The page was ripped.
*Lyla flipped through the book. The page was ripped.
7.
Looking forward to reading his favorite columnist, Marc went to the newsstand. The magazine was sold out.
*Looking forward to reading something new, Marc went to the newsstand. The magazine was sold out.
8.

Ronny put a nickel in the gum machine. The gum ball was red.
*Ronny refilled the gum machine. The gum ball was red.
9.
Ben sat down in the classroom. The desk wobbled.
*Everyone sat down in the classroom. The desk wobbled.
10.
He ordered a happy meal for his daughter. The hamburger was missing.
*He ordered happy meals for his children. The hamburger was missing.
11.
Jeremy was sitting at his college graduation. The cap was too tight.
*They were sitting at their college graduation. The cap was too tight.
12.
Trina got in the car. The seat belt was broken.
*Trina and Sam got in the car. The seat belt was broken.
13.
Jill started up her computer. The hard-drive crashed.
*The programming class started up their computers. The hard-drive crashed.
14.
Adam looked at his birthday card. The front was smudged.
Adam looked at his birthday cards. The front was smudged.
15.
The figure-skater was lacing up her skate. The lace broke.
*The figure-skaters were lacing up their skates. The lace broke.
16.
Laura got her bike out of the garage. The front tire was flat.
*Laura and Stephen got their bikes out of the garage. The front tire was flat.

Appendix C: Experiment 3 Spectrum Examples

Condition A ("the" should always be preferred)
1.
After stabbing the victim, the assailant threw away his weapon and fled. __ knife was found in the bushes.
2.
When Joe proposed, he gave Sally a ring set with a large gem. __ diamond was nearly flawless.
3.
Fred was terrified when he found a large bug in his bathroom sink. __ spider looked venomous.
4.
When the cut of meat was perfectly seared, Sean took it off the grill. __ steak was going to taste great.

Condition B ("the" should always be preferred)
5.

Ron swept the office. __ broom was starting to wear out.
6.
Roberta was pounding nails. __ hammer was brand new.
7.
Henrich started the car and let it warm up. __ steering wheel was cold.
8.
Candice could see that her pine tree wasn't doing well. __ trunk looked damaged.

Condition C (Should be somewhat ambiguous, "a" preferred)
9.
Clarence looked around the parking lot. __ porche was parked next to his car.
10.
Lilly walked into the classroom. __ overhead projector was sitting next to the door.
11.
Helga lounged at the small private beach. __ lifeguard was yelling at some kids.
12.
Bernie parked and walked towards the large central bank. __ teller was sitting outside on a cigarette break.

Condition D ("the" is infelicitous)
13.
The professor looked at the front row of her crowded classroom. __ student was asleep.
14.
The factory manager walked into the factory's bustling cafeteria. __ worker smiled and waved to him.
15.
Ethel was reading in the park. __ squirrel scampered up the tree she was reading under.
16.
Camilla was in the library. __ book fell off the shelf as she walked past.