Weak Definite Noun Phrases*

Greg Carlson, Rachel Sussman, Natalie Klein, Michael Tanenhaus

University of Rochester, Rochester, NY

1. Introduction

Semantic accounts of the definite article have tended to fall into two general classes: 'uniqueness' accounts (Russell, 1905), and 'familiarity' accounts (Heim, 1982; see also Abbott, 2004, *inter alia*). It is also generally known that there are instances of definite descriptions that are resistant to such analyses, and are difficult to subsume under either treatment, or even a combination of both. Birner and Ward (1994), Poesio (1994), and Barker (in press) explicitly deal with what we are going to call "weak" interpretations of certain definite NP's. Roberts (2003) is just one recent example of an attempt to integrate instances of this sort into a more general account.

In this paper we are going to argue that so-called "weak" definites should in fact not be subsumed under a more general semantic treatment, but form a distinguished class of (apparent) definite descriptions on their own that shares a semantics with (at least) bare count singulars, and probably not with definites more generally. We are going to present the argument in two forms. First, we will review the distributional and interpretive arguments for a separate class of items, and then substantiate these results with psycholinguistic experimentation, which not only supports the distinction but also provides some insight into the nature of the distinction. Much of this material can be found, in a more abbreviated form, in Carlson and Sussman (2005), but here we provide additional distributional and experimental evidence. We will be discussing instances that usually differ from the examples examined by Poesio and Barker, and will leave unresolved the question of whether their examples should be subsumed under our analysis.

2. Bare Singulars

* This material is based upon work supported by the NSF under Grant No. 0328849 and NIH Grant HD 27206. We thank Kevin Murphy, whose Hidden Markov Model toolbox for Matlab was used to run the simulations we discuss.

Our argument will be that weak definites and bare singulars form a natural class together in having the same constraints on interpretation and distribution. However, the term "bare singular"—intending determinerless non-plural count nouns—is probably a varied class. Examples of the type we are interested in for the moment are exemplified in (1a). Our claim is going to be that there is an interpretation of the examples in (1b) that are parallel.

(1) a. Sue took her nephew *to college/to prison/to class*
b. Sue took her nephew to the hospital/to the store/to the beach

However, "bare singulars" in English may occur in a number of different types of constructions, which we wish to set aside here for the time being (see Stvan, 1998 for an extensive overview of classes and types). For example, we are setting aside the coordinated cases examined by Heycock and Zamparelli (2003), which are argued to be instances of (anaphoric) definites:

(2) …a high degree of interplay between *artist and patron* is not necessarily so important

We also set aside bare singulars that can plausibly be argued to be truly indexical expressions, as in (3):

(3) When I'm in *town*/Bob is at *home*

And, to the extent that it is possible to set aside the participation of bare singulars in idioms, we intend to do so:

(4) Take *heart!*/ Mary was with *child*/ No books are on *hand.*

The examples we are interested in are generally found as objects of prepositions and objects of verbs. For example:

(5) a. They found him in *bed.*
b. The ship is at *sea*/at *port*.
c. He's in *jail*/in *prison*/in *church*.
d. Mimi attended *college/class/school*.

In many instances, they must be the object of a specific preposition. So if we change the propositions or verbs in the examples of (5), even to ones that in some sense ought to make sense, the examples nonetheless are ungrammatical.

(6) a *They found him on bed.
b. *The ship is in sea (OK in port)
c. *He's next to jail/prison/church
d. *Mimi destroyed college/class/school

We should note a couple of qualifications, however. In English, it appears that while prepositional co-occurrences are lexical-item specific, it would appear that Verb-Noun co-occurrences are mediated by the semantic class of the verb. So, for instance, the phrase "attend class" does not depend on the specific verb "attend," but any verb of presence or absence will be acceptable in this construction:

(7) a. Be present in class.
b. Go to class.
c. Miss class/skip class.
d. Have class.

Similarly, it's possible for entire classes of nouns to be licensed by a co-occurrence. For instance, one construction examined in Baldwin et al (2003) is where the preposition *by* takes the names of vehicles as its object:

(8) He went by *train/bus/car/airplane/motorcycle*

The generalization would appear to be that any noun whatsoever that designates a vehicle is acceptable here. So, for instance, skates and skis are not vehicles, but bulldozers and forklifts are (though hardly canonical vehicles for travel), and so it's fine to say:

(9) Max went to Philadelphia from New York by *bulldozers/forklift*.

It is interesting to speculate whether in the case of the verbs the co-occurring feature is a 'light verb" or something of that character, as outlined in Harley (200X), for instance, but this would take us well beyond our present line of thought. For the time being, there must be a co-occurrence in order for bare singulars to appear. This distinguishes English bare singulars from bare plurals, for example, and from bare singulars in some other languages, such as Brazilian Portuguese (Munn and Schmidt, 2000) or Albanian (Kallulli, 1998).

Another salient characteristic of bare singulars is that the lexical identity of the noun itself determines whether it can participate in the construction. Substitution of synonyms (except in cases such as those already noted) does not automatically yield grammaticality.

(10) a. *They found him in couch/cot/hammock (even if he sleeps there all the time)
b. *The ship is at ocean/lake
c. *He's in penitentiary/brig/mosque
d. *Mimi attended university(AmE)/seminar/institution

A further characteristic is that bare singulars cannot be modified by any restrictive modifiers, whether pre- or post-nominally. The following are not grammatical.

(11) a. *She traveled on sore foot
b. *He was found in silk-sheeted bed.

c. *Mimi attended class taught by Prof. Linskowski.
d. *The ship is now in port that's being dredged.

Finally, bare singulars are largely absent from subject position.

(12) a. *Car took the students to the conference.
b. ??College was attended by Bob.
c. *Port was entered by the oil tankers.

However, in certain cases such as copular sentences and other instances, they seem pretty acceptable:

(13) a. Prison is no place to make friends.
b. (?) Class should be attended regularly to get the best grades.
c. but: *Foot is one way to travel from Tibet to Ceylon.

In terms of their interpretations, one of the more salient characteristics of bare singulars is that they invoke "semantic enrichment." That is, they seem to call to mind more than a straightforward composition of parts would appear to invoke.

(14) a. Being *at school* is not simply being at *a* school, but that and more…
b. Being *in prison* is not simply being in *a* prison, but that and more…

Being at school is not simply a spatial relation but also means participating in the institution in the intended way, e.g. by getting an education. One merely visiting prison might be in a prison in order to do so, but being *in prison* means being incarcerated as well. For this reason, such semantically enriched readings are occasionally referred to in the literature as "activity readings." Also, for this reason, as well as the co-occurrences involved, these are occasionally called idioms, or characterized as having "idiomatic" readings. However, since there are a lot of bare singulars, and their "idiomatic" readings are systematically based on and include the spatial readings, we shy away from the label "idiom" as not so much wrong as misleading or not entirely appropriate, preferring to reserve this term for the canonical semantically opaque instances, insofar as possible.

3. Weak Definites

Having reviewed some basic properties of bare singulars, we now move to a consideration of definite noun phrases, the present point being that there is a subclass of definite NP's, the weak definites, which have just this same set of characteristics. The difficult part in this instance is identifying the subclass, as the form is that of a definite description and hence, in the usual instance, is systematically ambiguous. Our claim is that there is a very natural reading of the definite NP's in examples such as those in (15) which are parallel in interpretation to bare singulars.

(15) a. Mary went to *the store*.

b. I'll read *the newspaper* when I get home.
c. Open *the window*, will you please?
d. Fred listened to the Red Sox on *the radio*.

One general characteristic of definite descriptions is that they have a reference that carries over in anaphora. So, for instance, in (16), both people need to have heard about the same riot:

(16) Mary heard about *the riot*, and Bob did, too.

However, if one has a weak definite, then there is no such requirement of referential identity. In (17) the participants need not have listened to the same radio:

(17) Mary heard about the riot on *the radio*, and Bob did, too.

We are not going to go into an extended argument here about whether "the radio" in such examples is a device, or a medium, a station, or whatnot, as we will see that such observations carry over to a wide variety of other cases. What we're going to focus on here is the claim that such apparently non-referential readings correlate with a lot of other properties, namely, those properties also associated with bare singulars. Note that bare singulars do not induce identity of reference, either:

(18) Bob is *in jail*, and Fred is, too.

There is no necessity in Fred and Bob being in the same jail (though of course they may just happen to be.

We are going to use the VP-deletion structure as a rough-and-ready indicator of whether one has a strong, or regular interpretation for a definite NP, as in (20), or whether a weak on is possible, as in (19).

(19) Bob went to *the store*, and Mary did too. (different stores OK)
(20) Bob went to *the desk*, and Mary did too. (must be same desk)

In the following we implicitly invoke a rough-and-ready VP-deletion test in making our assertions. In nearly all cases, the definite is ambiguous between a regular reading and a weak reading where a weak reading is possible. At the moment, we know of some colloquial terms that are unambiguously only weak (e.g. words for prison that include expressions like "the pen," "the pokey," "the slammer," "the cooler," and possibly examples such as "the boondocks" or (on) "the briny" -- doubtless there are many more). Most, however, are ambiguous.

Once we tune into this distinction, we note that weak readings are only possible if the definite NP is governed by, or co-occurs with, another lexical item, normally a preposition or a verb.

(21) a. Kenneth is *at the store* vs. *behind the store*
b. They took the crash victims *to the hospital* vs. *past the hospital*
c. Sally *checked the calendar* vs. *tore the calendar*

They are also lexically restricted by the identity of the noun used. Substitution of a synonym or a definite expression with potentially the same reference does not necessarily guarantee a weak reading:

(22) a. He went to *the hospital.* vs He went to *the building.*
b. Scarface is in *the pen*. vs. Scarface is in *the cage*.
c. They listened to *the radio.* vs. They listened to *the tape recorder*.
d. You should see *the doctor*. vs. You should see *the nurse.*

If there is any modification, the weak reading is destroyed and the results are unambiguous regular or strong readings. The following lack weak readings:

(23) a. He went to *the 5-story hospital*.
b. They both checked *the calendar that was hanging upside down*.
c. Each man listened to *the red radio on the picnic table*.
d. Fred went to *the big store*.

One small difference between weak definites and bare singulars is that weak definites may be modified by affective expressions:

(24) a. I hear Bob is back in the ol'/doggone hospital again.
b. *Bob went to ol'/doggone college when he was just 16.

Like bare singulars, weak readings are largely absent when in subject position:

(25) a. (*)The store looked old.
b. (*) The calendar had a picture of a puppy on it.
c. (*) The radio fell off the picnic table.

But, like bare singulars, there are some examples (particularly copular sentences) that sound quite fine:

(26) a. The pen is no place to make friends.
b. The hospital is where you should go when very ill.
c. But: *The slammer was renovated in 1994.

And, finally, as with bare singulars, semantic enrichment is often involved, lending the sentences an "activity" reading.

(27) a. Going to the store is going to *a* store and more…(shopping)
b. Being in the hospital is being in *a* hospital, and more…(healing)

c. Looking at the calendar is looking at *a* calendar, and more… (gathering information of the type calendars are designed for)

There are a few instances in which standard American and British English differ in terms of whether the definite article is used. In American English, one goes to *the* hospital, or is in *the hospital*; in British English one is in hospital, or goes to hospital. British English uses the term *university* as a bare singular, whereas in American English *the university* has no weak reading, and in order to express the same thought one uses *college* even where one is in fact attending a university. Another relevant observation is that in German, there do appear to be instances that correspond to the weak definite readings of English, but are mostly limited to those instances where the definite article and a preposition form a portmanteau, e.g. *zum*, *ans*, *ins*, etc. The extent to which weak definites are common in languages with definite articles remains to be investigated.

3. Semantics of Weak Definites

One would think offhand that weak definites have the semantics of definites noun phrases in general. But this is the assumption that we are questioning in this paper. In fact, weak definites appear to have most of the properties of indefinite rather than definite noun phrases. As with the discussion of whether bare singulars are definite or indefinite (Stvan, 1998), we are hampered by the distributional limitations of these constructions. So, for example, a ready test for whether something has an indefinite reading is whether it may appear in the existential *there* construction in English, or as the object of inalienable *have*. When we attempt this with weak definites and bare singulars, though, we find the results unacceptable:

(28) a. *There was port on the coast.
b. *Bob has hand.
c. (*)There is the church over by the plaza.
d. (*)This town has the hospital

Note that the existentials are acceptable for weak readings provided one gives the so-called "List-there" reading which is compatible with proper names, definites, and strong quantification (and hence no test for indefiniteness). So the question is whether the examples of (28) are unacceptable because the noun phrases are not indefinite, or whether the appropriate co-occurrences fail (i.e. *have* and *there be* just do not 'govern' weak readings or bare singulars).

We can see, however, that like bare singulars, the interpretations of weak definites (if thought of as quantificational) can be scoped over in the same way as narrow-scope indefinites; the examples in (29c) and (30c) show that they may not only be scoped over, but also distributed over as well:

(29) a. Each mobster went to prison.
b. Most of them are in class.
c. My children attended college.

(30) a. Each mobster was sent to *the pen* for nine to twelve.
b. Most of the consumers went to *the store* for bargain-priced milk.
c. My children need to see *the doctor*.

The mobsters could clearly have been sentenced to different penitentiaries, the consumers have gone to different stores, and the children could all go to different doctors; the readings are the same possibilities as in the case of the bare singulars in (29).

It is true that definites can be "scoped over" if understood functionally or as definites enabled by "bridging."

(31) Each car in the parking lot had had *the steering wheel* removed.

We also know that "bridging" is possible from parts of (token) events:

(32) Sandra visited Yellowstone last summer, and Michelle visited the Grand Canyon. *The walking* made them sore.

So, we ask whether the phenomenon we are examining can be analyzed in terms of bridging from events, e.g. if Jill checked the calendar, then there was some one unique calendar she checked in that event, regardless of how many other calendars might have been around at the time. This would certainly promise a definition of uniqueness for the definite article to extend to even the scoped examples, and an assimilation to uniqueness and/or familiarity accounts.

There are some difficulties with this suggestion, however, which appear to offer little insight into the phenomena discussed above. For one thing, in a given event expressed by a weak definite, there is not necessarily any unique object picked out. For instances, if one says:

(33) Jacqueline took *the train* from Paris to Moscow.

She easily could have boarded a series of trains, so in the event there is no unique train at all. This is further evidenced by the fact that either of the following are entirely plausible continuations of (33):

(34) a. …it as clean and ran right on time.
b. …*they* were all clean and ran right on time

The other major problem is that, even where there is uniqueness, it badly over-predicts. So for instance, in an event of painting a desk, there is some unique desk that gets painted. Still, if we apply the VP-deletion test, or examine the scoping possibilities, there is no weak reading available:

(35) a. John painted the desk, and Bill did too. (Same desk only)

b. Each of the workers painted the desk. (Same desk only)

Further, it is not at all clear how "bridging" offers us any account of lexical restrictions, the co-occurrences, the a lack of modification, or semantic enrichment. Much the same thinking can be applied to functional readings of definites, which are for instance typically made more salient, and not less, by modification, especially by relative clauses.

The analysis of weak definites would seem to require according them the truth-conditions of indefinites, i.e. as equivalent to existentially quantified noun phrases. Consider the following scenarios:

Jim-Bob and Bubba at the scene of the accident.

Scenario 1: Jim-Bob and Bubba live in a remote small town that has two equally capable hospitals, Barton and Reed. They're out in Bubba's truck and witness a car accident. Two ambulances arrive, each picks up one victim, and departs. Bubba says: "They're taking the victims to the hospital." He believes they're both going to Barton, but in fact the ambulances take the victims to Reed. Did Bubba utter an untruth/something inappropriate? (no)

Scenario 2: Same thing except the ambulances decide the local hospitals cannot handle the injuries properly, so one goes to Kansas City General (for head trauma), and the other goes to Tulsa Mercy (for burn victims). Neither Bubba nor Jim-Bob have heard of these hospitals, and, in fact, live in a place so remote they have never even heard of Tulsa or Kansas City. Did Bubba utter an untruth/something inappropriate? (no)

Scenario 3: The EMT's of Ambulance 1 decide to take Victim 1 to the Strand Theater for a showing of an Adam Sandler movie, to cheer him up. They buy him a ticket, wheel him in and leave him there. The EMT's of Ambulance 2 instead decide to take Victim 2 to a McDonald's for a Happy Meal, to cheer her up. They wheel her in, pay for the meal, and leave her there. Did Bubba utter an untruth/something inappropriate? (yes)

No matter how we construct the scenarios, the truth-conditions are simply those of an existential. It is of course implied, in some fashion or other, that the (e.g.) hospitals are familiar and known, but it is not truth-conditionally required. If for instance Sam takes his family on a vacation in an entirely unfamiliar place, he can still go out to "the store" to get his family some goodies for the motel room, even when no one party to the conversation has any idea whatsoever what stores, if any, might be around. If he's at home, then there are typical, expected places he will frequent, but even there it is not required that he shop at any one of those places. Any store, even unfamiliar ones in distant and unlikely places, will do.

Are we to conclude then that weak definites are, in fact, indefinites in some kind of disguise? We don't believe so. For one thing, preliminary experimental evidence, to be described briefly below, makes it appear that indefinites are quite different from weak definites. Further, the lack of scoping is unlike what we find with indefinites, and the number neutrality is much more reminiscent of incorporated nominals than of indefinites (Farkas and de Swart, 2004). Thus it would appear that weak readings are more like

property expressions than real indefinites (McNally, 1995). The semantic enrichment is a very typical feature of incorporated structures, as well. And to restate the obvious, there is that definite article that requires some account.

Finally, and this is a vague intuition at the moment, it appears that weak definites, like bare singulars (and bare plurals) but unlike clearly marked indefinites, do not require a novel discourse referent (Heim, 1982). The first two examples found in (35), in contrast to the third, allow for the same discourse referent to be assigned the object expressions:

(36) a. Fred went to jail, and Bob did, too. (same jail seems fine)
b. Fred went to the store, and Max did, too. (same store seems fine)
c. Fred voted for a candidate, and Max voted for a candidate, too. (same candidate seems possible, but not quite so likely)

Whether this is a plausible argument remains unclear, but it would seem to call for further investigation.

Thus far, we have argued that there is a subclass of apparent definite NP's that are distributionally and semantically distinct from the regular definite noun phrases. In the following section, we seek to subject this hypothesis to experimental test. As we will see, there is substantial experimental evidence in favor of the distinction we have made from both judgment tasks and the less intrusive and more sensitive eye-tracking measures. We conclude with a brief report on modeling of the eye-tracking results, which shows that the distinction we have made is a genuine contingent variable in the processing of these constructions.

4. Experimental Evaluation

In this section, we present three experiments verifying the existence of a class of weak definites. The first presents an offline survey testing speaker judgments of weak definites in context. The second tests weak definites in a visual-world based eye-tracking task, based on earlier work by Spivey, Eberhard, Sedivy, and Tanenhaus (2002). The third experiment presents an eye-tracked visual-world adaptation of the judgment task, allowing us to probe not only speaker judgment, but also referents considered during the selection process.

4.1 Experiment 1

Experiment 1 tests speaker judgments about the possible referents of weak definites under ellipsis. We have previously observed that if an elided VP was referentially linked to a VP antecedent containing a regular definite, the elided form must reference the same entity that was indicated by the regular definite in the overt VP. However, for VPs involving weak definites, this condition need not hold. In example 37, Sally and Jane must have read the same book; in example 38, they need not have read the same newspaper.

(37) Sally read the book, and Jane did too.
(38) Sally read the newspaper, and Jane did too.

Experiment 1 presented 16 native speakers of English with brief written descriptions of two events. The events were of the same general type (i.e., could be referred to by the same verb), but involved two clearly distinct patient participants. Each description was followed by a brief summary sentence that always referred to the second event in the description via an elided VP. Participants were asked to decide whether this summary sentence accurately described the event they had just read about. Examples of both descriptions and summary sentences can be found in Table 1. We predicted that if the events in question involved a patient that was referred to by a weak definite, this summary sentence would be judged as accurate more often than when the events had involved patients that were referred to by regular definites.

4.1.1 Experiment 1 Results

For weak definites, participants accepted the elided sentence as an accurate description 73% of the time, while for regular definites, the sentence was accepted only 24% of the time, which constituted a statistically significant difference between the two definite

	context sentence	**Target sentence**
Regular Definites	Bill read Jane Austen's *Pride and Prejudice*, and Joe read *The Hitchhiker's Guide to the Galaxy*, by Douglas Adams	Bill read the book, and Joe did too.
Indefinite Definites	At breakfast, Samantha read the New York Times. Across the table from her, Frances was reading the Democrat and Chronicle.	Samatha read the newspaper, and Frances did too.

Table 1: Sample materials from Experiment 1

types ($t_1(15)=5.93$, $p<.001$, $t_2(5)=6.14$, $p<.001$)[i]. Thus, speaker judgments reflect a reliable and robust difference in the availability of the weak definite reading.

4.2 Experiment 2

Experiment 1 provides evidence that speakers reliably judge regular and weak definites to have different referential properties. Experiment 2 extends a result found in Spivey, Eberhard, Sedivy, and Tanenhaus (2002) to test the online interpretation of weak definites. Spivey et al. found that when instructions to manipulate objects in a scene involved a definite determiner and a noun, participant attention was naturally drawn to the singleton instances of that noun in the scene, despite the presence of other instances of the noun type that were presented as part of a cohesive group. The results of this experiment confirm the function of the definite article as put forth in Roberts, (2003) – that is, that definite articles serve to pick out some sort of unique entity in the context. In

light of their results, the Spivey et. al. work can be seen as evidence that a regular definite article is automatically interpreted as referring to a "unique" entity in the context. Thus we should expect to see looks to items that can easily be isolated as having some unique property (in the Spivey et al. experiment, separation from the cohesive group). In this way the Spivey et al. experiment has provided us with a well-defined set of expectations of how regular definites will be processed on-line with respect to a certain context. However, given the consistent failure of indefinite definites to pick out a unique referent outlined in the previous section, we might expect them to behave differently in on-line tasks. Namely, we would expect that for indefinite definites, the tendency for a noun-phrase of the form "the [noun]" to draw attention to singleton referents in the context should be lessened.

Our experiment was designed to determine whether indefinite definites would behave differently from regular definites during online referential processing. Specifically, we hypothesize that regular definites should draw participant attention to singleton targets, while indefinite definites will not.

We selected six nouns that often function as indefinite definites and matched them with comparable nouns that were obligatorily regular definites. The noun pairs were matched with verbs that were known to support the indefinite definite reading (as verified in our off-line judgment survey, described above) and placed into a sentential frame. This yielded a set of 6 pairs of matched experimental sentences: one version that contained an indefinite definite, and one version that contained a regular definite, but was otherwise identical.

For each sentence pair, we constructed a visual context meant to depict the scene just before the action depicted in the sentence is carried out. The scene showed a human actor, and three tokens of the object that was to serve as the patient of the action. Two of these were clustered near each other in a group, while the third was alone and isolated from the group by some distance. Additionally the scenes contained 3 distracter items that were not mentioned in the sentence, also presented in the form of a singleton and a small group of two (see Figure 1).

Regular definites

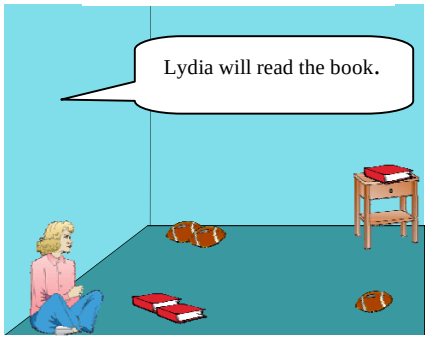

Indefinite definites

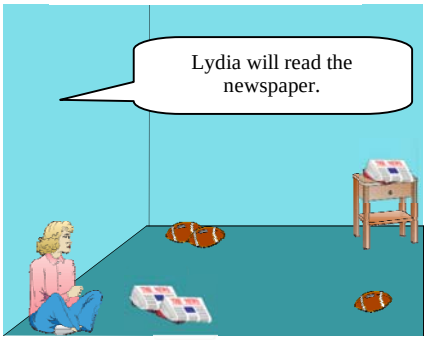

Figure 1: Sample scenes from Experiment 2

The arrangement of the items in the scenes was counterbalanced across items so as to avoid the possibility of participants coming to expect targets to appear in a particular location. This also served to avoid building participant expectations based on some interaction of object arrangement and actor eye gaze, body posture, etc. Crucially, however, the position of the actor and objects in the scenes remained constant across noun type conditions; that is, the indefinite

definite version of an instruction was presented with exactly the same scene configuration as for the regular definite version. In this way we insure that any differences observed in the processing and comprehension of indefinite vs. regular definites cannot be due to variations across conditions in target salience or proximity to the depicted actor.

Participants saw the visual displays on a computer screen while they heard a pre-recorded spoken version of the sentence matched to the display from a nearby speaker. They had been instructed that after hearing the sentence, their task was to choose the item in the display that they thought was most likely to be involved in the upcoming action. By involving them in a task that forced them to referentially link the spoken materials to the provided visual context, we hoped to get an idea of the particular item that participants interpreted our nouns of interest as referring to. Given the referential properties of indefinite definites as well as the results of the Spivey et. al. work, we expected that while regular definites would result in more participants choosing the singleton target as most likely to be involved in the action, indefinite definites should exhibit less of a tendency to be interpreted as referring to the singleton item. Included within the experiment were 12 filler trials that used an indefinite NP to refer to the object involved in the upcoming activity.

As participants were performing this task, we monitored their eye movements. A large body of work has established that eye movements are closely time-locked to spoken language comprehension and thus provides a useful tool for observing processes of reference resolution (Tanenhaus, Spivey-Knowlton, Eberhard and Sedivy, 1995, Eberhard et. al., 1995, Arnold et al, 2000, Runner, Sussman, and Tanenhaus, 2003, inter alia). By analyzing the time-course of eye-movements participants make as the spoken instructions unfold, we can get an idea of which items in the display are being considered as referents for our target noun at any given moment.

A total of sixteen members of the University of Rochester community took part in the experiment. All had normal or corrected-to-normal vision. None of the participants had taken part in any of the earlier pilot versions of this experiment, or in the pen-and-paper survey reported above.

4.2.1 Experiment 2 Results

This experiment yielded two types of results: target choice (member of the group target or the singleton target) and eye-movements.

For the indefinite definites, participants were much more likely to guess that a member of the group target would be involved in the upcoming action, choosing one of these items on 61% of trials. For regular definites, participants chose a member of the group target as involved in the action on only 33% of trials. This result illustrates two important aspects of indefinite definites. Firstly, target choice for indefinite definites and regular definites was significantly different, with indefinite definites eliciting more choices of group targets ($t_1(15)=4.66$, $p<.001$, $t_2(5)=3.45$, $p=009$). The second aspect to note is that for indefinite definites, choice of target item was equally distributed among the three available compatible targets, with each individual target being selected on a third of trials. For regular definites, a target item that was a member of the group target had only a 17% chance of being selected as the item most likely to be involved in the

action described, while the singleton target had a 66% of being selected. In this way, regular definites exhibit a marked preference for the singleton target item, while indefinite definites fail to give rise to any specific expectation of which target item will be involved in the action.

The analysis of the eye-movement data revealed a similar story. Here, during the window of time when the participant would be hearing the target noun of the spoken materials, they were (marginally) significantly more likely to fixate the group target if they were hearing an indefinite definite noun phrase than if they were hearing a regular definite ($t_1(15)=1.09$, $p=.14$, $t_2(5)=2.15$, $p<.05$).

In addition to these main analyses also performed a post-hoc analysis of the filler trials, where the spoken materials used true indefinites. Interestingly, we found that these indefinite trials patterned differently than both the regular and weak definite conditions. Here, participants were far more likely both to fixate and choose the singleton item in the display. Since these trials were not included as part of the original design of the experiment, we cannot draw any firm conclusions from these observations. However, the result does begin to suggest the importance of number neutrality to the weak definite interpretation, as well as isolate some critical differences between weak definites and true indefinites.

4.3 Experiment 3

Experiment 3 presents an online, eye-tracked version of the judgment task described in Experiment 1. Subjects were seated before a computer display while their eye-movements were monitored and recorded. On the left side of the computer screen, participants would see a picture of a person engaged in some activity involving an object or place. Next to the picture was a center column containing three items. One of these was the same object or place that had been involved in the activity depicted on the left. Also in the column was another object or place of the same type as the one depicted on the left, as well as a final item that was unrelated to the picture on the left. As participants viewed this picture, they heard a sentence describing the activity. Next, the picture on the left side of the screen disappeared, and a new character appeared in the rightmost column of the screen. Participants then heard a sentence stating that this new character was about to engage in the same activity that the person on the left had performed. To accomplish this, the sentence employed either a full VP including pronominal reference to the item involved in the earlier action, or an elided VP construction. Participants were asked to click on the item from the center column that they felt was most likely to be involved in this upcoming event.

A. Weak Definites

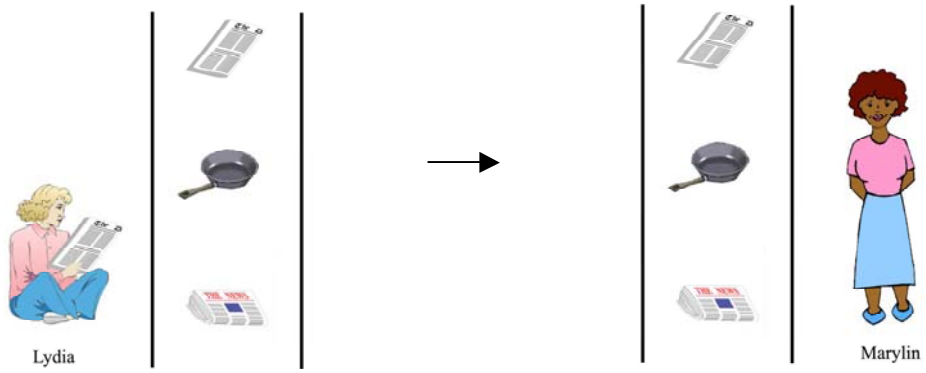

Lydia read the newspaper

Elided: Marylin will too.
Non-Elided: Marylin will read it too.

B. Regular Definites

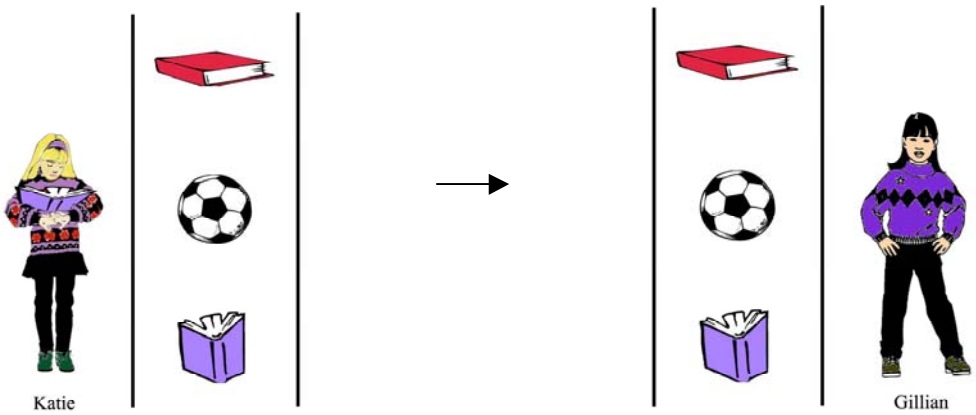

Katie read the book

Elided: Gillian will too.
Non-Elided: Gillian will read it too.

Figure 2: Sample materials from Experiment 3

Experimental trials varied as to whether the events described involved NPs that were known to have weak definite readings, or whether they included a different noun that could not support a weak definite interpretation. We took tendency to choose/look at the alternate item of the same type as the item involved in the original action as our dependent measure. Within the design of the experiment, the non-elided versions of the materials were treated as a baseline measure; in these cases, since the VP contains overt pronominal reference to the object or place involved in the first event, we expected that participants would not be likely to consider the different item as involved in the upcoming action. For elided VP constructions however, we predicted that the likelihood of participants to choose or look at the different item would vary with the availability of the weak definite reading of the noun. For VP anaphora that referenced an action

involving a noun with a weak definite reading, we expected that participants would be more likely to choose and look at the alternate item of the same type as the item involved in the original action than in cases where a full, pronoun-containing VP had been used. For sentences containing regular definites, we predicted that participants would not be any more likely to choose or look at the different item in the elided condition than in the non-elided condition.

Nineteen native speakers of English were recruited from the University of Rochester community to participate in this experiment. None of them had participated in Experiments 1 or 2.

4.3.1 Experiment 3 Results

Like Experiment 2, Experiment 3 yielded two kinds of data: participant choices and eye-movements. Participant choices of the item most likely to be involved in the upcoming action yielded a pattern of results consistent with our predictions about weak definites and their interpretation under VP ellipsis. Materials involving weak definites were more likely to elicit the choice of the "different" item in the display when presented in the VP-elided construction than when the same material had been presented in non-elided form ($t_1(17)=-1.9$, $p<.05$, $t_1(7)=-1.9$, $p=.05$). Regular definites however, were equally likely to prompt the participants to choose the "different" item from the center column for both the elided and non-elided conditions ($ts<1$). Additionally, the likelihood to choose the "different" item did not differ for the non-elided constructions across both regular and weak definite conditions ($ts<1$).

Early analyses of eye-movements reveal a similar pattern of results. During elided constructions involving a weak definite, participants are more likely to fixate the "different" item in the center column than when a regular definite had been used. Looks to the different vs. the same item in the center column do not seem to vary with noun type across the non-elided conditions. Further analyses are planned to determine how these patterns of looks depend on the participant's ultimate choice of referent. For these analyses, we will divide trials into groups based on the participant's ultimate choice of referent for the construction ("same" vs. "different" item). This will allow us to examine in detail the referent set that was considered as participants interpreted the materials. At issue is whether weak definites reliably serve to elicit more looks to the different item in the center column--even in cases where the participant ultimately chose the "same" item as the referent of the construction. Since in these cases, looks to the "different" item in the display cannot be linked to the participant's upcoming need to click on that item (and hence, the correlated series of movements that must be coordinated in order to move the mouse to this location) this finding would establish that weak definites reliably activate a different set of candidate referent than their regular definite counterparts.

In a preliminary examination of our eye-movement data that is contingent on which referent a participant picked during a particular trial, we have used trials from the first ten participants to train two separate Hidden Markov Models on sequences of eye-movements that correspond to either the selection of the identical referent or the new,

different token of the same noun. We chose Hidden Markov Models because of their applications to sequences of observations in time and their relationship to underlying "hidden" states, which is appropriate for our eye-movement data, as well as applicable to other areas of language research (Rabiner, 1989). Importantly, we chose to assign eight observational states to these models. These correspond to possible objects each participant could be looking at, plus a "loss of track" code for when the eye-tracker momentarily failed to capture the position of the eye, and a "beginning" and "ending" code for when the trials transitioned. Because of the way the experiment was designed, we intuited that the models would be most effective with four underlying states: choosing the same token of the critical noun, choosing the new or different token, choosing the unrelated object, or transitioning between trials.

One model was trained only on randomly ordered sequences of eye-movements made during trials in which the identical referent was selected, while the other model was trained on the sequences from trials where the new referent was ultimately chosen. Then both models were tested on an unseen data set to see whether the models could correctly predict which new data corresponded to a "different" choice and which corresponded to a "same" referent choice. Preliminary results are very promising: the models were fairly able to take a new sequence of eye-movements and predict which referent a subject would ultimately choose, reflecting which reading of the definite article that participant was giving the critical noun during that trial. This is with a small training set, when the model must also somehow account for observations such as "looking off the screen," which, while necessary for coding eye-movements, add noise to the training data.

Most important perhaps is the way the models organized the observation and transition matrices, which reflect how they "conceptualize" the data. (The observation matrix represents probabilities of seeing a particular observation given an underlying mental state; the transition matrix shows the probabilities of moving from one type of state to another.) Similar models built with more or fewer states did not make as accurate predictions, nor were their matrices as interpretable and compatible with the behavioral data. Our results thus far suggest choice is predicted in a meaningful way by sequences of eye-movements during a trial, and that choice should reflect whether an individual is giving a noun phrase a weak or regular definite reading.

5. References

Abbott, B. 2004. Definiteness and indefiniteness. In Handbook of Pragmatics, eds. Laurence R. Horn & Gregory Ward, 122-149. Oxford: Blackwell.

Arnold, Jennifer, J. Eisenband, S. Brown-Schmidt, and J. Trueswell. 2000. The rapid use of gender information: Evidence of the time course of pronoun resolution from eyetracking. Cognition 76:247-264.

Baldwin, T., J Beavers, L. van der Beek, F. Bond, D Flickenger, and I. Sag. 2003. In search of a systematic treatment of determinerless PPs. In Proceedings of the ACL-SIGSEM Workshop on the Linguistic Dimensions of Prepositions and their Use in Computational Linguistics Formalisms and Applications, Toulouse, France.

Barker, C. 2005. Possessive weak definites. In Possessives and Beyond: Semantics and Syntax, eds. Kim, Ji-yung, Lander, Yury, and Partee, Barbara H. 89-113. Amherst, MA: GLSA Publications.

Birner, B. and G. Ward. 1994. A Unified Account of English Fronting Constructions. In Penn Working Papers in Linguistics, Vol. 1, 159-165. Department of Linguistics, University of Pennsylvania.

Carlson, G and R. Sussman. 2005. Seemingly indefinite definites. In Linguistic Evidence, eds., S. Kepsar and M. Reis. Berlin: de Gruyter.

Eberhard, K., M. Spivey-Knowlton, J. Sedivy, and M. Tanenhaus 1995. Eye movements as a window into real-time spoken language comprehension in natural contexts. Journal of Psycholinguistic Research 24:409-436.

Farkas, D. and H. de Swart. 2004. The semantics of incorporation: from syntax to discourse transparency. Stanford, CA: CSLI.

Heim, I. 1982. The semantics of definite and indefinite noun phrases. UMass Dissertation.

Heycock, C. and R. Zamparelli. 2003. Coordinated Bare Definites. Linguistic Inquiry 34(3).

Kallulli, D. 1999. The comparative syntax of Albanian: On the contribution of syntactic types to propositional interpretation. University of Durham dissertation.

McNally, L. 1995. Bare plurals in Spanish are interpreted as properties. In Formal Grammar, eds. Glyn Morrill and Richard Oehrle, 197-212. Polytechnic University of Catalonia, Barcelona.

Munn, Alan, and Cristina Schmitt.1999. Bare nouns and the morpho-syntax of number. In Proceedings of the Linguistic Symposium on Romance Languages 1999.

Poesio, M. 1994. Weak definites. In Proceedings of Salt IV, eds. M. Harvey and L. Santelmann 282-299. DMLL, Cornell.

Rabiner, L. R. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. Proceedings of the IEEE, 77, 257-286.

Roberts, C. 2003. Uniqueness in definite noun phrases. Linguistics and Philosophy 26:287-350.

Runner, Jeffrey, R. Sussman, and M. Tanenhaus. 2003. Assignment of reference to reflexives and pronouns in picture noun phrases: evidence from eye movements. Cognition 89:B1-B13.

Russell, B. 1905. On denoting. Mind 14:479-93.

Spivey, Michael, M. Tanenhaus, K. Eberhard, and J. Sedivy. 2001. Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution. Cognitive Psychology 45:447-481.

Stvan, Laurel. 1998. The semantics and pragmatics of bare singular noun phrases. Ph.D. dissertation, Northwestern University.

Tanenhaus, Michael, M. Spivey-Knowlton, K. Eberhard, and J. Sedivy, 1995. Integration of visual and linguistic information during spoken language comprehension. Science 268:1632-1634.