



2014

What Does It Mean for an Implicit Object to be Recoverable?

Leila Glass
Stanford University

Follow this and additional works at: <https://repository.upenn.edu/pwpl>

Recommended Citation

Glass, Leila (2014) "What Does It Mean for an Implicit Object to be Recoverable?," *University of Pennsylvania Working Papers in Linguistics*: Vol. 20 : Iss. 1 , Article 14.
Available at: <https://repository.upenn.edu/pwpl/vol20/iss1/14>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/pwpl/vol20/iss1/14>
For more information, please contact repository@pobox.upenn.edu.

What Does It Mean for an Implicit Object to be Recoverable?

Abstract

Traditionally, researchers have claimed (e.g., Fillmore 198, Gillon 2006) that some verbs (such as *eat* and *bake*) lexically allow Implicit Objects (IOs) whereas other verbs (*devour*, *kill*) do not. This lexical idiosyncrasy is thought to explain why *I at* __ and *I baked* __ sound more natural than *I devoured* __ and *I killed* __. Other authors (e.g., Resnik 1993, Goldberg 2001, Scott 2006) have tried to explain such contrasts in a more principled way by exploring how IOs interact with discourse and information structure. In particular, they have observed that a verb's object must be "recoverable" in the discourse in order to be omitted. In this paper, I explore two prongs of this "recoverability" criterion. As the first prong, I argue that "recoverability" is a matter of degree; some objects can be recovered with more precision than others. I show that a given context's standard of "recoverability" is pegged to the speakers' goals and interests, so that an IO can be only loosely recoverable when it does not bear on speakers' goals, but must be more precisely recoverable when it is important. Turning to the second prong, I argue that an IO's "recoverability" depends on the common ground of a particular community. For example, since athletes routinely lift weights, it's part of their common ground that *I lifted* __ tends to mean *I lifted weights*. I report a simple corpus study showing that in communities where the action denoted by a given verb is associated with a routine action with a predictable object, as with *lifting weights* for athletes, the verb is more likely to appear with an IO. In both of these ways, I show that speakers' interests and shared knowledge can help to explain the apparent idiosyncrasy surrounding English IOs.

What Does It Mean for an Implicit Object to be Recoverable?

Lelia Glass *

1 Introduction

“Implicit object” (IO) is a name for what happens when a verb we normally consider transitive appears without an object:

- (1) a. Teresa baked ___; Phil ate ___; Sharese read ___.

In English, researchers have long observed that some verbs allow IOs whereas others do not. The most well-known minimal pair is is:

- (2) a. Phil ate.
b. *Phil devoured.

Traditionally, such contrasts have been analyzed as lexical idiosyncrasy, so that the lexical entry of *eat* specifies that it allows IOs while that of *devour* specifies that it does not. More recent work, however, has attempted to explain such data in terms of pragmatics and discourse information. For example, it’s often argued that the content of an IO need to be “recoverable” in the discourse context (e.g., Resnik, 1993; Goldberg, 2001) In this paper, I examine two factors that contribute to this criterion of recoverability.

First, I argue that recoverability is a matter of degree. At a minimum, one simply knows that an object exists; at a maximum, one knows exactly what it is; and there is plenty of middle ground in between. For an object to be omissible, it must be sufficiently recoverable along this continuum for speakers to pursue their communicative goals in the context. Second, I argue that a given object may be easier to recover, and thus easier to omit, against the common ground of a particular community of practice in the sense of Eckert and McConell-Ginet (1992): A group of people who unite to pursue a common goal, such as a swim team or a law firm. I perform a few preliminary corpus studies to show that the same verb may appear with IOs in one community of practice more than others; for example, *lift* is more often used with an IO (understood as *weights*) in fitness magazines compared to a general corpus. By considering both speakers’ goals and the community’s common ground, I try to elaborate the criterion that an IO must be “recoverable,” because I think this criterion can explain a great deal of the messy data surrounding English IOs.

Of course, some data will remain messy. Presumably Chinese speakers and English speakers are just as good at recovering information from the context, yet only Chinese speakers can omit the object of *want*; *Bu yao* (NEG WANT) means *I don’t want it* in Chinese, but the English *Don’t want* is quite unnatural. However, while discourse factors cannot explain every constraint on IOs, I will show that these factors can explain at least some of the data.

I’ll sketch the scope of these data before I turn to the analysis. Although the literature (e.g., Fillmore, 1986) distinguishes between “indefinite” and “definite” IOs, I have some qualms about this distinction, which I elaborate in Section 4 (see also AnderBois, 2012; Scott, 2006). Therefore, I consider data from both sides of this distinction. However, I limit myself to IOs that seem to stand in for DP’s (rather than CP’s). I also don’t take a stand on whether IOs are represented in the semantics, as in e.g., AnderBois 2012, or whether the verbs are simply intransitive ($\lambda x.ate(x)$) and a patient argument (the thing eaten) is pragmatically inferred (as in Recanati, 2007); the term “IO” is descriptive only.

1.1 Literature: Lexical Views vs. Discourse Views

Some researchers argue that whether a verb allows an IO or not hinges on its lexical entry, whereas others argue that it depends on how important or informative a verb’s object is in the discourse

*I am grateful to the PLC reviewers and PLC audience for inspiring comments; to the Stanford SemFest audience; and to Beth Levin and Chris Potts for their thoughts on a draft.

context. Here, I briefly introduce both views, suggesting that the more flexible discourse-oriented view can better handle the data.

Since *eat* and *devour* both describe consumption, it may appear difficult to explain why *eat* allows an IO while *devour* does not on semantic grounds alone. Instead, one may have to posit idiosyncrasy in the lexicon. Many other pairs of verbs pattern in the same way. Gillon (2006) provides a list building on that of Fillmore (1986):

Allows IOs	No IOs	Allows IOs	No IOs
<i>follow</i>	<i>pursue</i>	<i>suppose</i>	<i>maintain</i>
<i>leave</i>	<i>vacate</i>	<i>try</i>	<i>attempt</i>
<i>arrive</i>	<i>reach</i>	<i>promise</i>	<i>pledge</i>
<i>look</i>	<i>seek</i>	<i>find out</i>	<i>discover</i>
<i>insist</i>	<i>demand</i>	<i>object</i>	<i>oppose</i>

Like *eat* and *devour*, these minimal pairs have also been analyzed as lexically idiosyncratic (e.g., Fillmore, 1986; Gillon, 2006) and thus unpredictable. In this paper, however, I try to see how far we can get if we assume that these apparent idiosyncrasies have predictable sources.

A more predictive lexical analysis is proposed by Rappaport Hovav and Levin (1998) and Rappaport Hovav (2008). They distinguish *manner* verbs from *result* verbs and propose that result verbs do not allow IOs. Manner verbs describe the manner in which an action is carried out, but not its result; for example, *run*, *stab* and *fly* specify a type of movement but not a consequence thereof. Result verbs describe the result of an action but not the manner in which it takes place; for example, *arrive*, *kill* and *ascend* specify a change from not present to present, alive to dead, low to high, but not the manner in which this change was brought about. Since result verbs describe a change, Rappaport Hovav argues that the entity affected by this change must be present in the sentence, meaning that IOs are not allowed for result verbs (and may or may not be allowed for manner verbs). This analysis explains why, for example, *He hunted* does not need an object, but *He killed* sounds strange without one.

Such lexical analyses presuppose, however, that verbs can be clearly distinguished into those that allow IOs and those that do not. In fact, this distinction does not seem to be categorical. For almost every verb that is not supposed to allow IOs in the list above, I have found an example from the Web (indicated by [W]; via Google searches conducted between March and July 2013) in which it does appear with an IO. (To foreshadow, it's interesting that many of these surprising attestations, e.g., (4, 5, 8, 10, 12), seem to come from specific communities in which that verb might be especially significant.)

- (3) For dinner we let our little three year old choose. His request, chips and salsa. So off to Chili's we went. He **devoured**. Then we ordered a molten cake. ([W], mom's blog)
- (4) Officers **pursue** because they believe the person they are trying to stop must present a serious threat ([W], police forum)
- (5) Divorce raised its (not so) ugly head and they **vacated** at the end of the lease. ([W], investment blog)
- (6) But if you see what I see, if you feel as I feel, and if you would **seek** as I **seek** ... then I ask you to stand beside me (*V for Vendetta*)
- (7) Because I can't accept myself, I **demand** all the time. In all my relationships I **demand**, and no relationship is ever lasting... ([W], teen website)
- (8) Putting on weight now that I'm trying to **maintain** ... ([W], diet blog)
- (9) I **attempted** and failed miserably, only because I forgot that Sunday was the Sept. 2nd. ([W], site proposing a day where people don't go online)
- (10) Pledges belong in the same ensemble of church practices as prayer, study, worship, service... But unless you **pledge** ... none of this is possible. ([W], ellipsis in original; church website)
- (11) We are all human and we're always **discovering**. (self-help website)
- (12) Government teams will conclude, 'Therefore, we are proud to propose.'..., Opposition teams often say, 'Therefore, we are proud to **oppose**. ([W], Parliamentary Debate website)

Moreover, as Goldberg (2001) points out, even though result verbs are generally less acceptable with IOs than manner verbs are, result verbs productively allow IOs in certain semantic contexts. For example, many result verbs sound fine with IOs in generic statements, modal statements (which I add to Goldberg's list), repeated actions, infinitives, emphasis (which Goldberg calls "strong affective stance") and contrastive focus:

- | | |
|---|-----------------------------------|
| (13) Tigers only kill at night. | <i>generic</i> (p. 506) |
| (14) Dresses I would murder for | <i>modal</i> [W] |
| (15) Scarface killed again. | <i>repeated action</i> (p. 507) |
| (16) The singer always aimed to please/impress | <i>infinitive</i> (p. 506) |
| (17) Why would they give this creep a light prison term? He murdered!! | <i>emphasis</i> (p. 513) |
| (18) He burglarized, but she murdered! | <i>contrastive focus</i> (p. 514) |

From a lexical perspective, both the Web data and Goldberg's observations come as a surprise. If a verb is not supposed to allow IOs, whether because its lexical entry stipulates as much, or because it is a result verb, then it should not allow IOs under any circumstances. Thus, it seems that a lexical analysis cannot fully explain when and why verbs allow IOs or not. Instead, it is useful to consider the data from a more flexible discourse-oriented perspective.

One such discourse-centered analysis comes from Goldberg (2001), who argues that IOs are allowed when a verb's object is "low in discourse prominence." In the IO-friendly contexts Goldberg identifies, she suggests that the verb is placed in the informational foreground while the object is backgrounded. By omitting the objects of these verbs, the speaker emphasizes the important information and leaves out the rest. For example, in the contrastive focus and emphatic contexts, the speaker seems upset that the person in question murdered anyone at all, but may not care specifically who was murdered. Thus, the speaker can omit the object of *murder* to emphasize only the information she finds important.

A different sort of discourse-oriented analysis comes from Resnik (1993). He shows that the verbs that allow IOs are those whose objects are more statistically predictable than the verbs that do not. For example, *eat* provides more information about its object than *need* does, both intuitively, as one generally eats *food* of some sort, but one could need anything from a glass of water to brain surgery, and statistically, using Resnik's measurement of "selectional strength." Since *eat* provides more information about its object than *need* does, the object of *eat* is easier to recover than that of *need*, explaining why *eat* allows IOs whereas *need* does not. While Resnik's explanation does focus on qualities of particular lexical items, a given verb's selectional strength, his analysis hinges on how easily a hearer can recover an IO, a criterion that is inherently situated in a discourse.

In the data above, we see that verbs that are not supposed to allow IOs actually do allow them. In particular, surprising uses of IOs appear in Web data, especially from specific communities of practice, and in the semantic contexts noted by Goldberg. As I mentioned, these surprising uses of IOs are hard to capture on a purely lexical account which states that certain verbs should categorically disallow IOs. For example, since *kill* and *murder* are result verbs, it's surprising to find them with IO's in (13–18) if result verbs are not supposed to allow IOs. Instead, I have suggested that these data are better understood from a more flexible discourse perspective, in which IOs are allowed when the verb's object can be inferred in context and is not crucial to the discourse.

Next, I try to elaborate this perspective by exploring two issues related to recoverability. First, in Section 2, I discuss how the required degree of recoverability depends on interlocutors' goals. I suggest that this gradient view of recoverability helps to make sense of Goldberg's data. Second, in Section 3, I explore how an object may be more recoverable against the background of a particular community of practice. Essentially reproducing Resnik's finding within particular communities, I suggest that specific verbs "select" or predict their objects more strongly within a particular community, making it easier to recover that object. This idea, I argue, helps to make sense of some of the web data above, in which IOs appeared on websites for specific communities of practice for whom that verb has a predictable object.

2 Recoverability is a Matter of Degree

In this section, I argue that whether an IO is “recoverable” or not is not a yes-or-no question, but a matter of degree; and further, that how recoverable an object must be to license an IO is determined by the interlocutors’ goals.

2.1 Degrees of Recoverability and Speakers’ Goals

Although it’s commonly said (e.g., AnderBois, 2012; Goldberg, 2001) that an IO must be “recoverable,” it’s not always clear what that means. For every IO, we at least know that it must exist; we know that one can’t eat without eating something. We even know that the IO must be some sort of plausible object for the verb: if Mary ate, she probably ate something edible, not a democracy or a bookcase. We might make inferences based on the subject as well: if Mary is a tiger, she probably ate a different sort of food than if she were a human. Therefore, it seems that all IOs are recoverable in some minimal sense. Beyond these inferences, however, we may not know anything more specific. Except in some context where quinoa is the only food available or the only one Mary eats, we will not infer from *Mary ate* that *Mary ate quinoa*. Thus, all IOs are recoverable in a rough sense, but some are not recoverable in detail.

From this perspective, recoverability seems like a vacuous criterion. It does not explain, for example, why *Mary ate* sounds better than *Mary devoured*, because we can infer that Mary ate some existing, devour-able thing just as we can infer that Mary ate some existing, edible thing. In order to make recoverability meaningful, then, we need some way of fixing a standard for what counts as “recoverable” *enough*.

To find a motivated standard for recoverability, I suggest that we consider what speakers want to achieve in their conversation. Conversations, Roberts (1996) argues, are structured around “Questions Under Discussion” (QUD): questions that interlocutors are pursuing together. While Roberts argues that the ultimate goal of a conversation is to answer “the Big Question”: “What is the way things are?”, I prefer to think that the ultimate goal is to settle whatever QUDs the speaker cares about in order to pursue her broader, non-linguistic goals: to build a relationship, evaluate other people, convince someone of something, or decide what to do next. Whereas specifying a verb’s object would always help articulate “the way things are,” leaving it implicit can sometimes facilitate the speakers’ real-world goals while eliminating superfluous information.

Thus, considering speakers’ real-world goals can help determine how recoverable a verb’s object must be in order to be left implicit. For example, this reasoning can help explain why *eat* with an IO sometimes sounds fine, and other times strange. In the conversation below, *I ate* sounds natural because Phil is not interested in what Teresa ate for lunch, but whether she wants to go to lunch now. If Teresa adds detail, she may start a friendly chat about quinoa, but this information is superfluous if the only goal is to decide whether they should go eat.

(19) **Phil:** Want to have lunch?

Teresa: Thanks but I already ate. *vs.* Thanks but I already ate quinoa.

In contrast, in the conversation below, *I ate* sounds strange because Phil wants to evaluate the new Japanese restaurant. Here, Teresa is far more cooperative interlocutor if she specifies what she ate and how she felt about it:

(20) **Phil:** How was that new Japanese restaurant?

Teresa: It was great! I ate there. *vs.* It was great! I ate this amazing sashimi there.

Thus, by pegging “recoverability” to speakers’ goals, we can explain why the object of *eat* is recoverable enough in one context but not in another.

The same sort of logic can shed light on the well-known contrast between *eat* and *devour*. As Resnik shows, an object is more “recoverable” and more likely to be omitted when it is more predictable given the verb. Since he does not report the statistics on *eat* and *devour* in particular, I performed a simple corpus study to see how they differ. I collected the first 50 hits on CoCA

(Davies, 2008) for *eat* and *devour* when they appeared with *explicit* objects (*eat ice cream, devour autobiographies*). Then I categorized these explicit objects into five categories: food (*devour an entire pint of Ben and Jerry's*); metaphorical (*devour autobiographies*); human (*devour human flesh*); non-food substance (*devour both acetylene and hydrogen*, referring to a hypothetical life form on the moon Titan); and unclear (*devour it* where the antecedent was unclear). The results are given below:

Category	% of hits for <i>eat</i>	% of hits for <i>devour</i>
Food	80	16
Metaphorical	2	42
Human	0	28
Non-food substance	8	10
Unclear	10	4

This study shows that the object of *eat* is likely to be a sort of food, whereas the object of *devour* could be anything from an autobiography to human flesh. Since *eat* has a more consistent set of objects, its object is more recoverable than that of *devour*. To unite this finding with speakers' conversational goals, it is easy to imagine a context in which it is important that someone ate something-or-other without caring whether she ate ice cream or quinoa. In contrast, it is much harder to imagine a context in which it is important that someone devoured something-or-other without caring whether she devoured a person or an autobiography. Thus, *I devoured* is often too uninformative to suit most imaginable discourse goals.

To make these intuitions more precise, I propose that an object may be implicit when it is recoverable to a relevant degree of precision (RRDP), as defined below:

(21) **Proposal: Recoverable to a Relevant Degree of Precision (RRDP)**

Let V be a transitive verb and let $S1$ and $S2$ be sentences that are identical except . . .

- a. In $S1$, give V an IO and interpret $\llbracket V \rrbracket = \exists y \lambda x . xVy$
- b. In $S2$, give V an explicit object and interpret $\llbracket V \rrbracket = \lambda y \lambda x . xVy$
- c. The IO of $S1$ is *Recoverable to a Relevant Degree of Precision* if interlocutors' purposes are *not thwarted* when $S1$ is uttered instead of $S2$.

This proposal draws on two ideas from the philosophical literature. First, in the IO case, the verb's object slot is saturated using the "variadic function" proposed by Recanati (2007), which existentially quantifies over the object and turns the verb into an intransitive one. Second, the term "thwarting . . . purpose" is inspired by the work of Fara (2000). In her analysis of vagueness, she points out that some distinctions are relevant for human purposes and others are not. For example, if she is measuring coffee, a half scoop will make much weaker coffee than a whole scoop, so to use her words, her purpose would be "thwarted" if she ignored this distinction. In contrast, 12.01 grams of coffee grounds will make coffee that tastes just the same as if she'd used 12.03 grams of coffee grounds, so her goal of making coffee would *not* be thwarted if she ignored this distinction. While Graff Fara only applies this idea to vague predicates such as *tall*, I think it is true in general that people use language to highlight the distinctions that are important for their goals while ignoring the ones that are not. Thus, I suggest that this idea can also be used to understand whether IOs are recoverable enough for interlocutors' purposes. If Teresa just needs to explain that she is not hungry, her purpose is not thwarted if she uses an IO to obscure whether she ate a burger or quinoa.

2.2 Returning to Goldberg's Data

Next, I show how this proposal can shed light on why IOs are especially acceptable in the contexts identified by Goldberg. While Goldberg's data are surprising for a lexical view of IOs, I think these data make more sense from a discourse perspective like the one I have just outlined, in which IOs are allowed when they are recoverable enough for speakers' goals.

As I mentioned above, Goldberg shows that in certain semantic contexts, even verbs that are not generally thought to allow IOs do allow them quite productively (13–18). Goldberg argues that

these contexts are ones where the object is low in discourse prominence and can thus be omitted. While I agree with this analysis, I think it can be made clearer using the RRDP proposal.

First, in Goldberg's generic, repetitive, modal and infinitival contexts, the object would be low in informational content even if it were explicit (please see (13–18)). In all these cases, the IO does not pick out a single entity, but rather a series of entities across different situations (13–18, prey killed by different tigers, people killed in different worlds in which the speaker could murder someone to get a dress, people killed at different times by Scarface, and audience members in all the worlds where the singer achieves her aims). Since such a series of entities is so abstract and nonspecific, it would not be a very informative contribution to the sentence. In contrast, in a sentence that describes a single episode, the missing object could be a specific entity such as *the tiger killed the youngest calf*, which would be much more informative.

Therefore, when an IO appears in an episodic sentence, more information is lost than when an IO appears in a sentence describing some sort of iteration. In other words, in a sentence describing iteration, the opportunity cost (measured in terms of information) of using an IO is lower than in an episodic sentence. Thus, in these sentences describing iteration, it becomes less likely that interlocutors' communicative purposes would be thwarted if the IO version is uttered instead of an explicit object, perhaps explaining why IOs are more common in these contexts.

Turning next to Goldberg's emphatic context, I agree with Goldberg that emphasis foregrounds the verb to the exclusion of the object, but I think this analysis can be strengthened when we incorporate speakers' interests directly into the theory, as the RRDP proposal does. When a speaker uses emphasis to evaluate an event emotionally, she may use an IO if the verb alone is sufficient for this evaluation. For example, if the speaker's goal is to decry someone's morals, perhaps she only cares that he killed someone, but doesn't care who. Thus, the object of *kill* may be omissible in an emphatic construction when it's not relevant for the speaker's purposes. When we consider the speaker's goals in this way, we see that IO is not "low in discourse prominence" as an accidental fact about emphatic construction, but because such constructions are likely to be ones in which the speaker finds the object irrelevant. (Finally, I leave contrastive focus for future work.)

So far, I have shown that the RRDP proposal helps to derive Goldberg's IO-friendly contexts and clarify why a given IO might be "low in discourse prominence." Next, I suggest that it also helps to resolve what Goldberg calls a "paradox" (p. 515) that definite discourse references can be omitted when they are "highly topical", for example, when pushing a car, *You push, I'll steer*. It is not clear why such a salient referent would be "low in discourse prominence." However, on the current proposal, this data point makes a bit more sense. According to the RRDP proposal, an object can be omitted when it is recoverable enough for the speaker's goals. Thus, when the speaker's goal does not concern the object of a verb, the standard of recoverability is fairly low. For example, when Teresa just needs to decline Phil's invitation, it doesn't matter what she ate beforehand so she can omit the object of *ate*. Phil will not be able to infer from the context what exactly Teresa ate, but presumably he doesn't care. In contrast, when the speaker's goal *does* concern the object of the verb, the standard of recoverability is higher. For example, when trying to move a broken car, it does matter what the speaker wants the hearer to push. However, since the car is so salient in this context, the hearer will be able to infer what the speaker wants him to push. Thus, in both the *eat* case and the *push* case, the object is recoverable to a relevant degree of precision; the only difference is that in the *eat* case that degree is low, whereas in the *push* case, it is high. Thus, RRDP helps to resolve Goldberg's "paradox."

3 Recoverability Depends on the Community's Common Ground

So far, I have explored one of two issues related to recoverability: that it is a gradient notion keyed to interlocutors' goals. I have shown how this analysis sheds light on that are puzzling for lexical analyses of IOs. Next, I turn to the second issue related to recoverability: that it is shaped by the common ground of a particular community of practice. This observation helps to explain many of the surprising attestations of IOs found on the Web.

As I described above, Fillmore (1986) and Gillon (2006) propose a list of verbs that are sup-

posed to categorically disallow IOs, but these verbs are attested with IOs in Google searches. Many of these attestations, I mentioned, come from websites for a specific community of practice for whom the verb may have special significance. For example, *pursue* is used with an IO on a discussion board for police officers, who routinely pursue suspects, and *vacate* appears with an IO on a blog for people who invest in real estate.

I suggest that these surprising data make sense if we consider how the IO may be more recoverable within the common ground of the surrounding community of practice. Since police officers routinely pursue suspects, as opposed to art history degrees or snow leopards, *pursue* may work for them just as *eat* works for people in general. Just as people normally eat food, police officers generally pursue suspects, so the object of *pursue* can be implicit just as that of *eat* can. To generalize, I propose that a community is likely to use a verb without an overt/explicit object some of the time when the community routinely engages in an action that is usually described by a certain verb, and the action involves an object that can be inferred to a degree that is adequate for speakers' goals.

For the object to be sufficiently inferable, it could be the sort of thing that people generally only have one of; for example academics usually only have one dissertation each, and regular people usually have only one home each, so academics can interpret *Sven defended today* and people discussing homes understand *We're going to buy*. Alternatively, the object could be a unique thing upon which everyone in the community performs the action. For example, soccer players all interact with the same ball, so they understand *She shoots*; debaters in a tournament all interact with the same motion, so they understand *We are proud to oppose*; and retail associates who all work at the same store will understand *I closed last night*. Finally, the object could be a set of objects within which it is not always important to make distinctions. For example, when an athlete lifts weights, she lifts many different weights but doesn't need to explain which in order to describe the activity. Thus *I lifted* is common parlance among athletes. Financiers regularly buy and sell various stocks, so *It's a good time to sell* is interpretable to them. Soccer players and musicians will understand *I haven't played all week* to mean either *soccer* or *music*, since each group commonly plays this abstract mass noun. Finally, the object may be a large space in which it is not always important to make distinctions. For example, *I cleaned* or *I swept* may refer to my apartment or my floors.

Uniting all these cases is the idea of a conventional routine for that community. Within such a community, the object of these verbs is easier to infer than in the general community, and thus it makes sense that it is more frequently omitted within these communities.

This analysis is similar in some ways to Levin and Rappaport Hovav's (2013) discussion of *clean*. They show that *clean* can have an IO when it is used as a manner verb to describe the routine actions involved in cleaning (mopping the floor, wiping the counter and so on) but not when it is a result verb simply indicating that something was made clean in some way or other. But although both my analysis and theirs involve a routine, my analysis does not distinguish result verbs from manner verbs. For example, *sell* is a result verb (requiring that a transaction took place, without specifying how it happened) but it still allows IOs within the fields of finance and real estate. Similarly, *close* is a result verb, and yet if we all work at a store, it is natural for me to say *I closed last night*, leaving the object implicit. Thus, while the idea of a routine is central to L&RH's analysis and this one, this one predicts that result verbs should sometimes allow IOs, as the data seem to illustrate.

To test this analysis, I conducted a corpus study in which I compared the uses of IOs in the general linguistic community to particular sub-communities on the Corpus of Contemporary American English (Davies, 2008). I chose three verbs and three communities in which I thought that verb would have predictable object for that community: *lift* in fitness magazines, and *buy* and *sell* in finance magazines. First, I found 200 random hits for each verb on general CoCA (deleting any from the specific community I was interested in, for example deleting finance magazines for *buy* and *sell*) and counted (using a Python program, after tagging each hit manually) how many times the verb appeared with an IO. Second, I used CoCA's scroll menu to search only within "home/health" magazines (there is no specific fitness magazine option) for *lift* and finance magazines for *buy* and *sell*. (For *lift*, I deleted hits from home magazines such as *Southern Living* and kept only the results from *Men's Health* and *Shape* in order to create a corpus of fitness magazines only.) Among the first 200 hits for each verb within the specific magazine community, I tagged the data and used the same

Python program to count how many times the verb appeared with an IO. Next, I used a Binomial Exact Test in R to check whether the general corpus/specific corpus difference was significant. The Binomial Exact Test is designed for tests with Boolean results, which is appropriate for analyzing explicit vs. implicit objects. It is also designed to handle data that are not necessarily normally distributed, like these. For all three verbs, I found that they appeared with IOs far more frequently in the specific community than in the general corpus ($p < 0.01$):

Lift	Explicit Objects	Implicit Objects
General CoCA	196	4
<i>Men's Health/Shape</i>	179	21

Buy	Explicit Objects	Implicit Objects
General CoCA	193	7
Finance Mag CoCA	180	20

Sell	Explicit Objects	Implicit Objects
General CoCA	187	13
Finance Mag CoCA	173	27

These data are consistent with my hypothesis: objects are easier to omit in communities of practice in which these objects are easier to infer. This finding is not predicted if IOs are simply listed in the lexicon as available or not, but the finding is expected if IOs are licensed by more flexible discourse factors.

4 Definite vs. Indefinite IOs

So far, I have explored two factors that affect whether an IO is sufficiently recoverable that it can be omitted: how much the IO matters for speakers' goals, and how an IO can become more recoverable within a particular community of practice. But before I conclude, I return to an issue that I mentioned above: the literature's distinction between "indefinite" and "definite" IOs (Fillmore), or equivalently between "existential" vs. "anaphoric" IOs (Condoravdi and Gawron, 1996, henceforth C&G). These researchers argue that the particular verb determines whether its IO will be in one category or another (although see Recanati, 2007, for a different view), and they propose diagnostics to determine whether a verb selects indefinite (existential) or definite (anaphoric) IOs. However, none of these diagnostics seems as clear as it should be (Recanati, 2007; Scott, 2006).

Fillmore proposes that a verb takes "indefinite" IOs when one can assert a verb while denying knowledge of the object, whereas a verb takes "definite" IOs when one cannot deny knowledge of the object. This can be tested with a sluice, as in *s/he VERB-ed, but I don't know what*:

- (22) a. He **ate**, but I don't know what.
 b. ??She **noticed**, but I don't know what.

However, sometimes the IOs that are supposed to be "definite" do license sluicing, contrary to this diagnostic. For example, Fillmore says *win* only allows definite IOs, but in the natural-sounding Web example below, it licenses sluicing, which only indefinite IOs are supposed to do:

- (23) I just **won**, but I **don't know what**. Neither the winning email, nor the claim form mention which prize I won. [W]

Fillmore also argues that "indefinite" IOs should behave like regular indefinite NPs. Indefinites can antecede pronouns, as in an example based on (Heim, 1982).

- (24) **A sailor** walked in. **She** sat down. (pronoun anteceded by indefinite)

Thus, indefinite IOs should license pronouns as well. But this prediction does not seem to hold:

- (25) Mark **lifted**. ??They were heavy.

This awkward example shows that “indefinite” IOs do not behave like regular indefinites with respect to pronouns.

Moreover, regular indefinites introduce novel discourse referents, and thus must not be understood as co-referring with a previously mentioned discourse referent. In (26), *a book* in the second sentence cannot refer to the book introduced in the first sentence, because the indefinite article conveys novelty (Heim, 1982).

- (26) I bought **a book** on Friday, and then read **a book** all weekend.

If “indefinite” IOs patterned like indefinites, then they should also introduce novel discourse referents. Thus, Fillmore predicts that an “indefinite” IO should never be understood to co-refer with a previously mentioned discourse referent. However, in the web examples below, the IOs of *read* and *eat* are most saliently understood as co-referring with the book and pizza mentioned earlier in those sentences. This inference may be cancelable; (27b) could continue with *but we fed the pizza to the children and ate sandwiches instead*; but the point is that there is a salient reading in which the IO co-refers with a prior discourse referent. Such a reading is not predicted if “indefinite” IOs are truly indefinite.

- (27) a. Your book arrived last week and I’ve been **reading** all weekend and LOVING it [W]
b. We bought pizza and **ate** in the park. [W]

Thus, “indefinite” IOs do not behave as indefinites should, questioning Fillmore’s classification. Moreover, AnderBois (2012) shows that while “indefinite” IOs are supposed to both license sluicing and introduce novel discourse referents, these two diagnostics do not always line up as they should. Thus, Fillmore’s classification does not divide the data in a clear way.

The classification system of C&G does not divide the data cleanly either. According to C&G’s diagnostic, when there is an “existential” IO as in (28a), we infer that John didn’t eat *anything*, so it is bizarre to add that he ate grapes. In contrast, when there is an “anaphoric” IO as in (28b), we infer that John didn’t apply to the particular job that the IO is anaphoric to; but it’s possible that he applied to a different one.

- (28) a. There was bread on the table but John **didn’t eat**. ??He ate **grapes** instead.
b. There was a good job available but John **didn’t apply**. He applied to **a different job**.

However, it seems that sometimes an “existential” IO (such as the IO of *eat*) can mean that the speaker didn’t eat a specific, previously mentioned thing, but instead ate something else:

- (29) I’d been feeling sluggish all day. I got home and we made **dinner** but **didn’t eat** instead I ate **a piece of bread with some peanut butter** about 30 minutes before running. [W]
(30) I absolutely *loathe* **buffet food** to start with, so that’s why I **didn’t eat**. I only remember eating the favor, which was **a huge sugar cookie**. [W]

Thus it seems that neither the indefinite/definite nor existential/anaphoric distinctions are particularly clear. That’s why I have considered IOs that might fall into any of these categories. At least in general, however, some IOs do seem to require an antecedent (*I noticed*) whereas others do not (*I ate*). Thus, while the indefinite/definite and existential/anaphoric distinctions may be hazy, there is still a real phenomenon to explain. C&G note that no one has ever successfully derived this distinction from a verb’s lexical semantics; I wonder if the distinction could be better explained by looking at how strongly each verb predicts its object in the sense of Resnik (1993), because it seems that *eat* tends to involve food whereas *notice* could have any number of things as its complement: a problem, a person, a haircut, or even a CP such as *that it was raining*. Thus it makes sense that *notice* would only allow an IO when this could be recovered from the immediate context, whereas *eat* could allow them more generally. Relating this idea to speakers’ goals, it is unlikely to be significant in a conversation that someone *noticed* without knowing whether they noticed *an error* or *that it was raining*, whereas it may be significant that someone *ate* without knowing whether they ate a burger or quinoa. I hope that future work will shed more light on these issues.

5 Conclusion

Exploring two aspects of the “recoverability” criterion for English IOs, I have argued that speakers’ real-world goals influence the degree to which an object must be recoverable, and I have presented corpus results suggesting that IOs are more frequent in communities where that object is more recoverable because it forms part of a routine. In both of these ways, I have argued that whether a verb allows an IO can be negotiated in context as speakers try to maximize the relevant information conveyed while minimizing unnecessary complexity (Sperber and Wilson, 1986). When the IO is not relevant, or easy to recover in that particular speech community, it is likely to be omitted. Thus, this grammatical phenomenon is shaped by speakers’ goals and common ground. Zooming out, I see IOs as just one example of how social factors influence syntax and semantics (e.g., Clark, 1997).

References

- AnderBois, Scott. 2012. Indefiniteness and the typology of implicit arguments. In *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL 30)*, 232–240. Santa Cruz, CA.
- Clark, Herbert H. 1997. Dogmas of understanding. *Discourse Processes* 23:567–59.
- Condoravdi, Cleo, and Jean M. Gawron. 1996. The context-dependency of implicit arguments. In *Quantifiers, Deduction, and Context*. CSLI Publications.
- Davies, Mark. 2008. The Corpus of Contemporary American English. 450 million words, 1990-present. Available online at <http://corpus.byu.edu/coca/>. Brigham Young University.
- Eckert, Penelope, and Sally McConell-Ginet. 1992. Think practically and look locally: Language and gender as community-based practice. *Annual Review of Anthropology* 21:461–490.
- Fara, Delia G. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 45–81.
- Fillmore, Charles J. 1986. Pragmatically controlled zero anaphora. In *Proceedings of Berkeley Linguistics Society* 12, 95–107. Berkeley, Calif.: Berkeley Linguistics Society.
- Gillon, Brendan. 2006. English relational words, context sensitivity and implicit arguments. Ms., McGill University.
- Goldberg, Adele. 2001. Patient arguments of causative verbs can be omitted: The role of information structure in argument distribution. *Language Sciences* 503–524.
- Heim, Irene. 1982. The Semantics of Definite and Indefinite Noun Phrases. Doctoral dissertation, University of Massachusetts, Amherst.
- Levin, Beth, and Malka Rappaport Hovav. 2013. Manner and result: The view from *Clean*. Manuscript, Stanford University / The Hebrew University of Jerusalem.
- Rappaport Hovav, Malka. 2008. Lexicalized meaning and the internal structure of events. In *Theoretical and cross linguistic approaches to the semantics of aspect*, ed. S. Rothstein, 13–42. John Benjamins.
- Rappaport Hovav, Malka, and Beth Levin. 1998. Building verb meanings. In *The Projection of Arguments: Lexical and Compositional Factors*, ed. M. Butt and W. Geuder, 97–134. CSLI Publications, Stanford.
- Recanati, François. 2007. It is raining (somewhere). *Linguistics and Philosophy* 30:299–345.
- Resnik, Philip S. 1993. Selection and Information: A Class-Based Approach to Lexical Relationships. Doctoral dissertation, University of Pennsylvania.
- Roberts, Craige. 1996. Information structure: Towards an integrated formal theory of pragmatics. In *OSUPL Volume 49: Papers in Semantics*, ed. J. H. Yoon and A. Karthol. The Ohio State University Department of Linguistics.
- Scott, Kate. 2006. When less is more: Implicit arguments and relevance theory. In *UCL Working Papers in Linguistics*, 1–25. University College London.
- Sperber, Dan, and Deirdre Wilson. 1986. *Relevance: Communication and cognition*. Harvard University Press.

Department of Linguistics
 Margaret Jacks Hall
 Stanford University
 Stanford, CA 94305
 lelia@stanford.edu