Communicating with Cost-based Implicature: a Game-Theoretic Approach to Ambiguity

Hannah Rohde
University of Edinburgh
3 Charles Street
Edinburgh, UK
hannah.rohde@ed.ac.uk

Scott Seyfarth
UC San Diego
9500 Gilman Drive
La Jolla, CA, USA
sseyfarth@ucsd.edu

Brady Clark
Northwestern University
2016 Sheridan Road
Evanston, IL, USA
bzack@northwestern.edu

Gerhard Jaeger
University of Tübingen
Wilhelmstraße 19
Tübingen, Germany
gerhard.jaeger@uni-tuebingen.de

Stefan Kaufmann
Northwestern University
2016 Sheridan Road
Evanston, IL, USA
kaufmann@northwestern.edu

Abstract

A game-theoretic approach to linguistic communication predicts that speakers can meaningfully use ambiguous forms in a discourse context in which only one of several available referents has a costly unambiguous form and in which rational interlocutors share knowledge of production costs. If a speaker produces a low-cost ambiguous form to avoid using the high-cost unambiguous form, a rational listener will infer that the high-cost entity was the intended entity, or else the speaker would not have risked ambiguity. We report data from two studies in which pairs of speakers show alignment of their use of ambiguous forms based on this kind of shared knowledge. These results extend the analysis of cost-based pragmatic inferencing beyond that previously associated only with fixed lexical hosts.

1 Introduction

A growing body of work demonstrates that joint communication tasks yield alignment of referring expressions, highlighting the role of interlocutors' experience of shared common ground in establishing conventions (Brennan & Clark, 1996; Garrod & Pickering, 2004). Less well-established, however, are predictions regarding which form∼meaning mappings interlocutors will converge on. To address this, we evaluate alignment in contexts where interlocutors' common ground includes the costs of producing particular forms.

Consider the shapes in Figure 1. In a context that contains only the first item, a speaker can efficiently draw attention to it by saying "Look at the circle." In a context with all three shapes, a more specific referring expression—such as "blue circle"—is required to unambiguously indicate that same item. However, if it is necessary to draw attention to the third item, the speaker may need to accept either inefficiency or ambiguity. Since there is no efficient label (e.g., "circle") for the third item's unique shape, it is costly to unambiguously refer to it in the context of Figure 1—a longer or more obscure expression is necessary (e.g., "the triangle-and-square thing" or "the blue shape that's not a circle"). On the other hand, the speaker can avoid producing a costly expression by instead using an ambiguous expression such as "the blue thing" or "the blue shape." The question is whether a listener can be expected to infer that the intended referent is the difficult-to-describe shape, even though "the blue shape" could in principle also refer to the blue circle.

An accurate inference about the intended referent of "the blue shape" requires the following chain of reasoning: The listener would have to realize that had the speaker intended to refer to the blue circle, a relatively short unambiguous expression would have sufficed ("the blue circle"); since the speaker used a low-cost ambiguous expression "the blue shape" instead of the available low-cost unambiguous name,

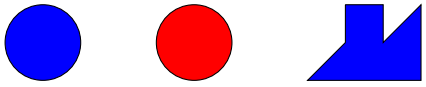

Figure 1: Variation in cost of referring expressions

Proceedings of the 16th Workshop on the Semantics and Pragmatics of Dialogue, September 19–21, 2012, Paris, France.

the implicature is that the circle is not the intended referent, leaving the difficult-to-describe item as the preferred target.

Using ambiguous forms to convey meaning depends in part on the listener's ability to diagnose the source of the ambiguity: Does the ambiguity likely arise from the speaker's own production decisions or from other factors that make the expression noisy or unclear (see Schlangen & Fernandez, 2007)? We focus here on productions whose ambiguity listeners can be confident originates with the speaker herself.

This paper presents two studies testing speaker alignment in dialog games with superimposed costs and rewards that are predicted to guide production and comprehension of ambiguous forms. Our results show that speakers' use of ambiguous expressions reflects the relative costs of available forms. Rather than avoiding ambiguity, speakers show behavior that is in keeping with theories of communicative efficiency that posit that speakers make rational decisions about redundancy and reduction.

2 Game Theory and Implicature

Game theory is an area of applied mathematics, which aims to provide a framework for analyzing the behavior of individuals (*players*) in strategic situations (*games*) (von Neumann and Morgenstern, 1944; see Clark, 2011, Chapter 4 for an introduction). It is used to describe games in which players have choices regarding their behavior and preferences over possible outcomes. The outcomes depend on both players' choices: While some games are zero-sum, meaning that success requires that only one player can win, in other types of games, both players can succeed if they coordinate their actions. Linguistic communication is typically argued to be an example of this second type (Lewis, 1969).

Games are characterized by shared knowledge, meaning that players know the overall structure of the game. They know what moves are possible in which situation by each player, what consequences are associated with each move, and what preferences each player has. Crucially, all players know that the other players know these facts. A game-theoretic approach makes an assumption of player rationality.

In a recent computational model that treats language use as a cooperative game, Golland, Liang, and Klein (2010) show that a rational speaker's decisions about whether or not to use ambiguous referring expressions can be significantly improved by embedding a model of the listener that represents the shared game knowledge described above. In their communication game with a human listener, unambiguous expressions were preferred over ambiguous-yet-accurate expressions when the speaker selected an expression that optimized utility for the listener.

However, in that type of listener-oriented model, a speaker's choice of expression is based solely on maximizing the probability that the listener understands the speaker's intent. In this paper, we consider how a game-theoretic approach offers further predictions about the players' behavior when the various (ambiguous and unambiguous) referring expressions also have different costs, when the players share knowledge of costs, and when the players know that they share knowledge. In particular, this approach suggests cases in which literal ambiguity might actually be preferred in pursuit of efficient communication.

The prediction from such an approach is that an ambiguous form can be used to refer unambiguously if an unambiguous form is costly and other meanings can be conveyed at low cost (Clark, 2011; Jäger, 2008). In other words, a listener who knows the relative costs of unambiguously referring to X (high-cost) or Y (low-cost) may reason that a speaker using an ambiguous word "X-or-Y" (low-cost) intends to convey X, or else she would have said "Y".

This type of reasoning has been used to explain the conventional use of "some" (Jäger, 2008). Having heard a speaker use the word "some", the listener is faced with two possible interpretations: AT-LEAST-ONE-AND-POSSIBLY-ALL or else AT-LEAST-ONE-AND-NOT-ALL. A rational listener is said to reason that, had the speaker intended to convey the meaning ALL, she would have used the low-cost (short and easy to produce) form "all". The speaker, having instead used a low-cost but ambiguous form "some", can be taken to implicate that the intended interpretation is not ALL, but is instead a meaning that would have been costly to produce unambiguously: AT-LEAST-ONE-BUT-NOT-ALL.

This account of "some" formulates in game-theoretic terms the well-known Gricean account, which focuses on the amount of information conveyed. In the Gricean version, the literal meaning of "some", AT-LEAST-ONE-AND-POSSIBLY-ALL, conveys less information than "all". Its meaning is strengthened to a more informative meaning of AT-

LEAST-ONE-BUT-NOT-ALL via implicature: A cooperative speaker who obeys the maxim of Quantity and intends to convey the more informative meaning ALL would have said "all", but since she did not, the meaning AT-LEAST-ONE-BUT-NOT-ALL is favored.[1] Grice's recognition of the importance of speaker intention echoes a game-theoretic approach to signaling and the calculation of what must be true in order for a rational speaker to have produced a particular utterance (Stalnaker, 2005).

The AT-LEAST-ONE-BUT-NOT-ALL implicature associated with the use of "some" is what Grice called a generalized conversational implicature: the implicature AT-LEAST-ONE-BUT-NOT-ALL is typically associated with the proposition expressed. What remains an open question is whether this type of cost-based inferencing applies beyond a fixed lexical host like "some". The next sections describe two studies aimed at measuring alignment in a communication game with explicit superimposed costs and rewards for production and comprehension.

3 Study 1: Communicating about Objects with Divergent Costs

We created a networked interactive two-player chat environment (see Figure 2) in which pairs of players could communicate about a set of objects. Costs and rewards were made explicit via points, and players shared both knowledge of the cost/reward structure as well as a shared goal of working together to communicate successfully. In contrast to Study 2 in the next section, the costs imposed in Study 1 served to highlight a single high-cost entity in each category, creating a bias to conventionalize the meaning of an ambiguous form to refer to that entity. In production, the prediction is that players will use a low-cost *ambiguous* word to refer to an object with a high-cost *unambiguous* name, as long as other objects can be unambiguously referred to with relatively low-cost names. In comprehension, the prediction is that players will more often interpret ambiguous words to refer to objects with a costly unambiguous name than to objects whose unambiguous name is associated with a mid or low cost (henceforth 'high-cost objects', 'mid-cost objects', and 'low-cost objects').

3.1 Participants

10 pairs of English speakers from Northwestern University received $10 to participate in the study.

3.2 Methods

Materials The game involved a set of objects in two categories—three flowers and three trees. Players could communicate using a set of eight referring expressions: six unambiguous names and two ambiguous words. The costs varied among the unambiguous names, but the ambiguous words were both low-cost. Table 1 shows the point costs associated with the eight forms. The point values themselves are less important than the relative values: In this study, the cost of the most expensive name in each category ("Tulip"/"Pine Tree") was more than four times the cost of the least expensive name and more than twice the cost of the mid-cost name.[2]

Name	*Cost*	*Name*	*Cost*
"Rose"	-60	"Apple Tree"	-60
"Daisy"	-120	"Palm Tree"	-120
"Tulip"	-280	"Pine Tree"	-250
"Flower	-80	"Tree"	-80

Table 1: Referring expressions and their costs (Study 1)

Task Players were seated in separate rooms in front of computers showing the interactive game interface depicted in Figure 2. They were told that they would take turns as Sender and Receiver in a game that involved communicating about a set of objects. As the Sender, a player would see a gnome highlight an object with a spotlight, and the Sender's task was to send a message to the other player so that the other player (the Receiver) could guess what the highlighted object must have been. Sending a word consisted of pressing a button on the screen which

[1] This logic is spelled out in the Stanford Encyclopedia of Philosophy (Davis, 2010) in terms of the interaction of cost (maxim of Manner) and information (maxim of Quantity):

> Assuming that the accepted purpose of the conversation requires the speaker to say whether or not all athletes smoke, a speaker who said "Some athletes smoke" would be infringing the Quantity maxim if she meant only what she said. So she must have meant more. If she believed that all athletes smoke, she would have said so. Since she did not, she must have meant that some but not all athletes smoke. As a bonus, she achieved brevity, in conformity to the maxim of Manner.

[2] Alternate materials were constructed with abstract shapes and nonsense names, but a pilot study found that participants had difficulty learning the names. Variants of Study 1 were conducted with first names (e.g., "Ann") for unambiguous names for objects in plant and vessel categories and family names or nonsense names (e.g., "Puliniki") for ambiguous words; the results matched those presented here with flower/tree materials.

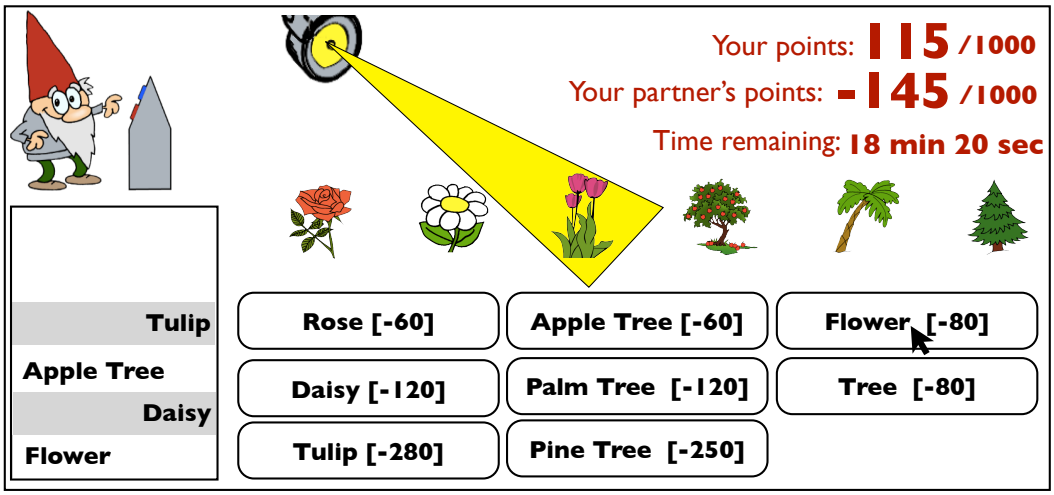

Figure 2: Player's view of the interactive chat environment for both Study 1 and 2 (point values from Study 1)

displayed the name and cost; pressing the button resulted in an immediate point decrement for the Sender. If the Receiver successfully identified the intended object, then both players got an immediate reward (Sender and Receiver scores increased). Otherwise, there was no reward and the Sender had to retry until communication succeeded—the penalty in that case being the continued point decrements for sending multiple words. The reward for successful communication was +85 points for each player. If either player reached the target score of 1000 points, the game ended and both players were free to leave. Otherwise, games continued for 20 minutes or until the gnome had highlighted a total of 60 objects. Scores could become negative if the point decrements of production outstripped the point increments for successful communication.

For each pair, we calculated a unique sequence of objects to be highlighted so that it would be impossible to reach 1000 points in less than 60 turns without successfully coordinating the use of ambiguous words. The first 10 turns were intended as practice turns, involving only low-cost and mid-cost objects.

Shared Knowledge Both players were able to see the available referring expressions and their costs and were told that the other player could as well. Both players were informed of the reward structure (+85 points each for successful communication, 1000 points to end early) and told that the other player had likewise been informed. The game interface showed the current scores of the two players, the time remaining, and a scrolling chat window listing previous words sent and received.

3.3 Results and Discussion

Of the 10 pairs, 5 consistently coordinated their referring expressions, allowing an early exit from the game. Two pairs played for the full 20 minutes, struggling to coordinate their efforts, and their limited attempts at coordination failed to yield an early exit to the game. Of the remaining 3 pairs, 1 pair did not make a serious attempt at using the ambiguous words, and the other 2 pairs used them repeatedly but had difficulty finding a strategy while doing so, although all 3 did eventually manage to exit the game before reaching the time or turn limit.

Table 2 shows a transcript from one pair of players, listing the first 26 moves of their game. The transcript demonstrates how the players developed a coordinated strategy for using the ambiguous words: the use of an ambiguous word by Player 1 when the intended referent was not a high-cost object (which led to Player 2's initial guess that the high-cost object was the target), the use of an ambiguous word by Player 2 (which Player 1 failed to interpret as a reference to the high-cost object), and eventually their convergence. After the success shown in the last line of the table, the players continued to reliably use "Flower" and "Tree" to refer to the tulip and pine tree, and the game ended after 44 moves when Player 2 reached 1000 points.

For the analysis, we measured the effect of one within-players factor (the target object's unambiguous cost) on two binary outcomes: whether the Sender sent an ambiguous word (production outcome) and whether an ambiguous word resulted in successful communication (comprehension out-

Sender	*Highlighted Target*	*Form Used*	*Receiver's Guess*
1	daisy (mid)	'Flower'	tulip (high)
1	daisy (mid)	'Daisy'	daisy (mid)
2	palm (mid)	'Palm'	palm (mid)
1	palm (mid)	'Tree'	pine (high)
1	palm (mid)	'Palm'	palm (mid)
2	apple (low)	'Apple'	apple (low)
1	palm (mid)	'Tree'	pine (high)
1	palm (mid)	'Palm'	palm (mid)
2	daisy (mid)	'Daisy'	daisy (mid)
1	tulip (high)	'Tulip'	tulip (high)
2	apple (low)	'Apple'	apple (low)
1	pine (high)	'Pine'	pine (high)
2	tulip (high)	'Flower'	daisy (mid)
2	tulip (high)	'Flower'	daisy (mid)
2	tulip (high)	'Flower'	tulip (high)
1	pine (high)	'Pine'	pine (high)
2	pine (high)	'Tree'	pine (high)
1	rose (low)	'Rose'	rose (low)
2	rose (low)	'Rose'	rose (low)
1	palm (mid)	'Palm'	palm (mid)
2	tulip (high)	'Flower'	rose (low)
2	tulip (high)	'Flower'	daisy (mid)
2	tulip (high)	'Flower'	tulip (high)
1	pine (high)	'Tree'	apple (low)
1	pine (high)	'Tree'	pine (high)
2	pine (high)	'Tree'	pine (high)
⋮	⋮	⋮	⋮

Table 2: Excerpt of a game transcript from two successful players in Study 1 ('Apple Tree', 'Palm Tree', and 'Pine Tree' are abbreviated without the word 'Tree'). Consecutive rows with the same Sender show retries.

come). We also measured the effect of trial number on the interpretation of ambiguous words in order to evaluate the time course of the Receivers' cost-based inferencing. For that, the three-way outcome of Receiver guess was treated as three binary variables (high-cost-or-not, mid-cost-or-not, low-cost-or-not). For all analyses, means, and figures, we consider only non-retry moves. We report the logistic-regression coefficient estimate and p-value (based on the Wald Z statistic; Agresti, 2002) for the factors *cost* and *trial number* (both treated as numeric factors) with random participant-specific intercepts and slopes.

Additionally, Figure 3 shows the overall rates of

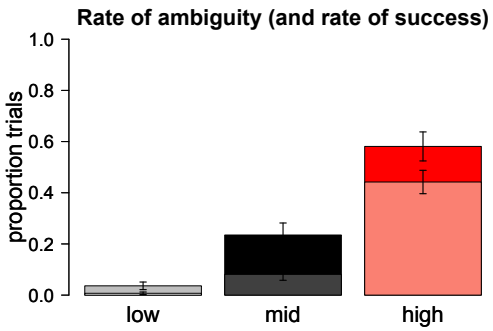

Figure 3: The percentage of trials with a low-, mid-, or high-cost target in which a Sender produced an ambiguous word. The lower part of each bar represents the percentage of trials in which the use of an ambiguous word was successful.

use and success for the ambiguous words, measured over proportions of trials. The height of each bar in the graph shows the proportion of trials for each cost—low, mid, high—where an ambiguous word was used. The shaded (lower) portion of the bar shows the proportion of those ambiguous words that resulted in successful communication.

As predicted, Senders produced an ambiguous word more often if the gnome had highlighted a target object whose unambiguous name was high cost (β_{cost}=1.94, p<0.001): 58.9% of high-cost-target trials yielded an ambiguous word, whereas only 24.8% of mid-cost-target trials and 4.6% of low-cost-target trials yielded ambiguity.

Receivers likewise paid attention to cost, correctly guessing the target more often when an ambiguous word was used for a high-cost target than for a mid- or low-cost target (β_{cost}=1.68, p<0.005): Trials with an ambiguous word yielded successful communication 74.1% of the time if the intended target was high cost, compared to 37.8% and 37.5% success when an ambiguous word was used for mid-cost and low-cost targets, respectively. In other words, of the 58.9% of trials in which an ambiguous word was produced for a high-cost target, 74.1% of those trials yielded successful communication (compared to less than half the time for trials in which an ambiguous word was used for a low- or mid-cost target), as depicted in the 'high' bar of Figure 3.

We also ask whether the interpretation of an ambiguous word changes over successive trials. Restricting the analysis to trials in which an ambiguous word was used, we find that the interpreta-

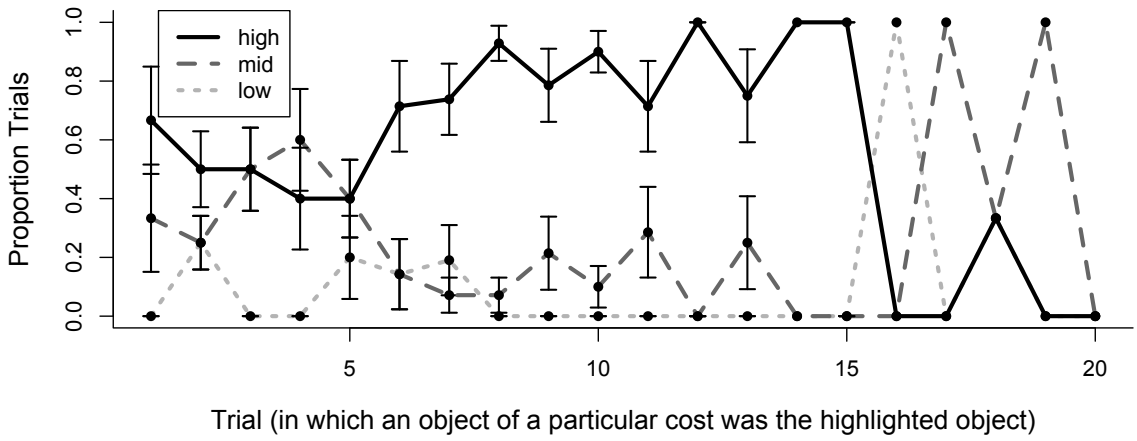

Figure 4: Rate at which Receivers in Study 1 interpreted ambiguous words as referring to objects with high-, mid-, and low-cost unambiguous names. Errors bars show standard error of the mean.

tion of ambiguous words favors high-cost objects over time (β_{trial}=0.29, p<0.01) and disfavors, albeit not significantly, mid-cost and low-cost objects (mid: β_{trial}=-0.22, p=0.09; low: β_{trial}=-0.23, p=0.11). Figure 4 shows the probability of a Receiver interpreting an ambiguous word as referring to a high, mid, or low-cost object. Trial number in Figure 4 (and in the regression) represents the number of trials for which an object of that cost has been highlighted—e.g., the datapoints at Trial=10 are the 10th trials, within each game, in which the gnome highlighted a high-cost object (either a tulip or a pine tree) and the Sender sent an ambiguous word ('Flower' or 'Tree') and the Receiver guessed a high-cost object (solid line), a mid-cost object (dashed line), or a low-cost object (dotted line).

What is apparent in Figure 4 is that the data after Trial=15 becomes noisier (more fluctuation) and more sparse (limited/no error bars). This drop-off corresponds to the point in the game when most successful players had reached 1000 points and left, so the data for the later trials represents the behavior of pairs of players who had failed to coordinate their referring expressions. These players overall used fewer ambiguous words, and because of this, many were watching their scores become more and more negative. Data from all players up through Trial=15 was analyzed in the time course logistic regression.

These results show that players can and do coordinate their use of referring expressions, conventionalizing the meaning of an ambiguous form to associate it with the object whose unambiguous name is the most costly to produce.

4 Study 2: Communicating about Objects with Similar Costs

As a further test of the predictions of a game-theoretic model, a second study was conducted that shifted the cost structure within each category. Compared with the costs in Study 1, the high-cost names in Study 2 were less costly than before, and the point difference between the low-cost, mid-cost, and high-cost names was smaller. The revised costs were predicted to reduce the likelihood that players would coordinate their use of referring expressions. The target score and the reward for successful communication stayed the same, but the stakes were lower (i.e., the cost structure imposed lower costs overall) so that it was possible for players to reach the target score in less time without making recourse to the ambiguous words. Rational players could choose to waste less effort attempting to align their use of referring expressions because the benefit of alignment was potentially outweighed by the points lost in rounds in which successful communication required the Sender to send more than one word.

4.1 Participants

10 pairs of English speakers from Northwestern University received $10 to participate. None had participated in Study 1.

4.2 Methods

The game environment contained the same set of six objects. The game rules and shared knowledge of those rules matched Study 1. The only difference was the costs (Table 3): The most expensive name

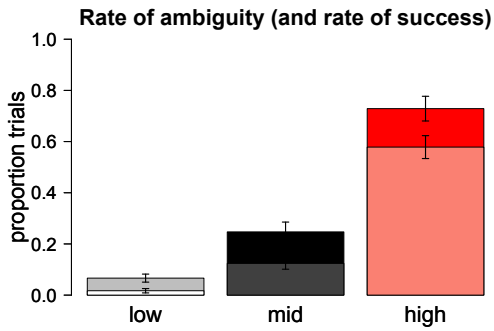

Figure 5: The percentage of trials with a low-, mid-, or high-cost target in which a Sender produced an ambiguous word. The lower part of each bar represents the percentage of trials in which the use of an ambiguous word was successful.

in each category cost slightly more than two times the least expensive name and not more than one and a half times the mid-cost name. Successful communication was still rewarded with +85 points to both players and the game ended when either player reached 1000 points. Since the absolute costs for the low- and mid-cost objects were similar to Study 1 while the absolute cost for the high-cost object was reduced, the average point cost in Study 2 was reduced and therefore it was possible to end the game after 48 turns rather than 60 without coordination.

Name	*Cost*	*Name*	*Cost*
"Rose"	-80	"Apple Tree"	-80
"Daisy"	-140	"Palm Tree"	-135
"Tulip"	-165	"Pine Tree"	-170
"Flower	-80	"Tree"	-80

Table 3: Referring expressions and their costs (Study 2)

4.3 Results and Discussion

Of the 10 pairs, 8 coordinated their referring expressions, allowing an early exit from the game. Contrary to prediction, the imposition of lower costs did not reduce players' motivation to conventionalize. This can be seen in reliable effects of cost on production and comprehension, as in Study 1. Figure 5 shows the overall rates of use and success for ambiguous words in Study 2.

Senders produced an ambiguous word most often if the highlighted object was high cost (β_{cost}=2.56, p<0.001): 72.6% of high-cost-target trials yielded an ambiguous word, whereas only 25.6% of mid-cost-target trials and 6.4% of low-cost-target trials yielded ambiguity. Receivers likewise paid attention to cost, correctly guessing the target more often when an ambiguous word was used for a high-cost target than for a mid- or low-cost target (β_{cost}=1.18, p<0.001): Trials with an ambiguous word yielded successful communication 82.5% of the time if the intended target was high cost, compared to 51.0% and 21.4% success when an ambiguous word was used for mid-cost and low-cost targets, respectively.

The time course differs slightly from Study 1, however. Receivers in Study 2 did not show a reliable rise in their preference to interpret ambiguous words as referring to high-cost targets. Again restricting the analysis to trials in which an ambiguous word was used (see Figure 6), the only reliable effect is that ambiguous words become less likely to be interpreted as referring to low-cost objects over time (β_{trial}=-0.65, p=0.05). The effect of trial number on the use of ambiguity for mid-cost targets is again not significant, though the coefficient is positive here (it was negative in Study 1), meaning that ambiguity tended to favor the mid-cost target slightly over time (β_{trial}=0.13, p=0.32). The slight increase here in the rate at which ambiguous words were interpreted to refer to high-cost objects is not significant (β_{trial}=0.09, p=0.22), unlike in Study 1. Given the different cost structure, the point in the game when most successful players had reached 1000 points and left comes at Trial=11. Figure 6 shows the subsequent sparseness after that point, and the time course regression includes data only up to Trial=11.

Across the two studies, Sender/Receiver pairs who coordinated their use of ambiguous forms were better able to communicate. Two pairs in Study 2 converged on a pattern of usage whereby an ambiguous word was used to refer to the object with the mid-cost unambiguous name. This could explain the time course result whereby the slope for mid-cost guesses for ambiguous words was positive (though not significantly so) in this study but not in Study 1. Another difference between the two studies is that in Study 2, convergence in the use of ambiguity in one category did not guarantee convergence in the other: two pairs used 'Tree' reliably but not 'Flower'.

In terms of our predictions, players did show sensitivity to the differences in the cost structure, but what characterized the behavior of players in Study 2 was a more immediate and sustained use of ambiguous words as referring to high-cost objects for most pairs and an openness to assign the ambigu-

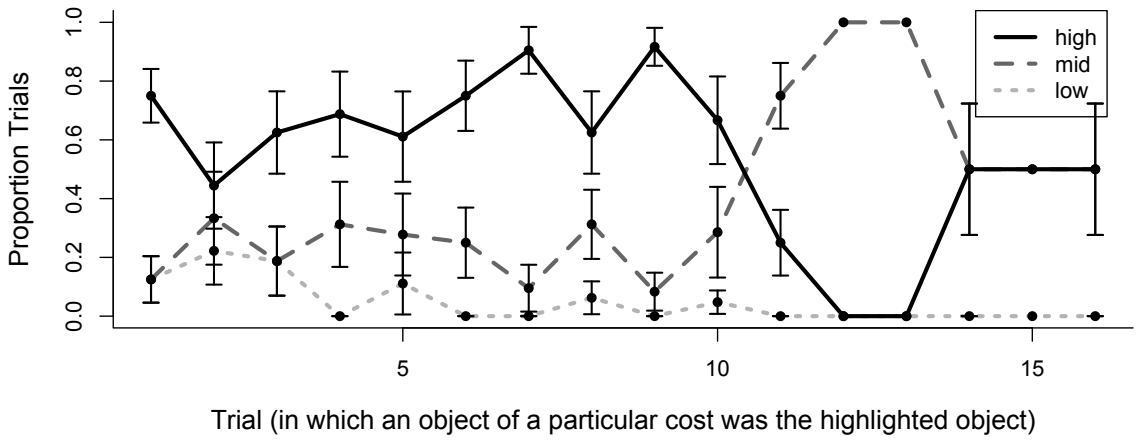

Figure 6: Rate at which Receivers in Study 2 interpreted ambiguous words as referring to objects with high-, mid-, and low-cost unambiguous names. Interpretation of axes and error bars as in Figure 3.

ous word to a mid-cost object for a small subset of pairs. The players' behavior seems to point to a greater willingness to experiment with the ambiguous words in a context like Study 2 where, despite their experimentation, they could see their scores increasing more quickly to the target value. Alternatively, rather than casting Study 2 as the context with increased experimentation, one can ask why players did not experiment more in Study 1. Perhaps the higher production costs in Study 1 made players avoid risking ambiguity and possibly having to retry.

Lastly, one can ask if players simply used a trial-and-error strategy for finding an efficient alignment without recourse to the pragmatic inference required for the cost-based implicature. To rule this out, we considered the trials in which a Sender *first* sent an ambiguous word. We pooled the data from the two studies since each participant could only contribute one datapoint. In keeping with the cost-based implicature account, Receivers inferred, more often than chance, that the high-cost object was the intended object ($\chi^2(2)=11.54$, $p<0.005$).

5 General Discussion

In keeping with the game-theoretic prediction, we saw that ambiguous words can be used meaningfully to refer to entities with costly unambiguous names, crucially if other referents can be identified with low-cost unambiguous names. This extends the claim that listeners draw cost-based implicature beyond the case of a fixed lexical host like "some".

We also saw sensitivity to relative costs: In comparing the two studies, the trajectory for how ambiguous words were interpreted over time in Study 1 (where the unambiguous names had more divergent costs) differs from the trajectory in Study 2 (where the unambiguous names had more similar costs). Only in Study 2 did players ever assign the ambiguous word to a mid-cost item, and only in Study 2 did a pair of successful players coordinate their use of one ambiguous word but not both. Contrary to predictions, however, the lower stakes in Study 2 did not yield a reduction in the players' overall ability to coordinate referring expressions.

This line of research raises an important question about how production cost should be measured. For the studies here, we imposed our costs arbitrarily—i.e., we could just as well have assigned the cost of 'Rose' to 'Tulip' or to 'Daisy'. One could imagine instead a cost metric that reflects properties such as length (as in Figure 1) or the presence of non-native phonemes or other articulation-based complexity. Alternatively, there is evidence that frequency contributes to production difficulty: Speakers are slower to name objects with low-frequency names than high-frequency names (Oldenfield & Wingfield, 1965). Speakers also experience difficulty, as measured by their disfluency, when describing objects which have not yet been mentioned (Arnold & Tanenhaus, 2007), are unfamiliar or lack a name (Arnold, Kam, & Tanenhaus, 2007). To avoid the inference step of assessing what phonological/lexical/pragmatic properties speakers perceive as costly, the proof-of-concept studies presented here used externally imposed costs to test the use of cost-based implicatures. The next stage of this re-

search will extend the experimental framework described here to the kinds of costs that are imposed naturally in regular conversation.

If cost does influence choice of referring expression, one must still ask whether its role is automatic or strategic (Horton, 2008). By presenting this work from the standpoint of calculable implicatures, we have framed the questions in strategic terms. The factors that guide speakers' strategic selection of referring expressions may depend not only on the costs associated with production (as we have shown here) but also on speakers' estimates of the costs and benefits of successful communication and of the degree of coordination between speaker and hearer (van Deemter, 2009).

It is also possible that our participants had a more automatic reaction to the salience of high-cost objects — maybe they just associated an ambiguous form with the most salient object of that category, where cost indicated salience. This scenario is compatible with a game-theoretic account — the reasoning being that it would be unnecessarily costly to refer to a prominent entity with a full name when a reduced or ambiguous form could be used.[3]

Lastly, these results fit with existing work on the role of reduction in communication — namely, work showing that speakers make rational decisions about redundancy and reduction and do not necessarily avoid ambiguity (Levy & Jaeger, 2007; Jaeger, 2010; Piantadosi, Tily, & Gibson, 2011). Like this previous work, we argue that ambiguity arises from a rational communication process. In our case, ambiguity arises in contexts in which the explicit costs of production are part of speakers' common ground.

Acknowledgments

This research was supported by an Andrew W. Mellon postdoctoral fellowship to the first author. The results from this paper have been presented at the 25th Annual CUNY Conference on Human Sentence Processing in New York and the 36th Penn Linguistics Colloquium in Philadelphia, both in March 2012. We thank Judith Degen, Roger Levy, and Kenny Smith for helpful discussion. We also thank research assistants Elizabeth Mazzocco, Anant Shah, Alex Djalali, and John Lee.

[3] Clark (2011) characterizes the role of salience in reduction (here, pronominalization) in a game-theoretic framework:

> It is cheaper to refer to a prominent element with a pronoun. It is correspondingly more marked (hence, costlier) to refer to a more prominent element with a description or name when a pronoun could be used. Here the speaker and hearer are presumably using some notion of salience to guide their choice. (p. 252)

References

Agresti, A. (2002). *Categorical data analysis*. Wiley.

Arnold, J. E., Kam, C. L. H., & Tanenhaus, M. K. (2007). If you say thee uh you are describing something hard: The on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33(5)*, 914–930.

Arnold, J. E., & Tanenhaus, M. K. (2007). Disfluency effects in comprehension: How new information can become accessible. In E. Gibson & N. Perlmutter (Eds.), *The processing and acquisition of reference.* MIT Press.

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1482–1493.

Clark, R. (2011). *Meaningful games: Exploring language with game theory*. MIT Press.

Davis, W. (2010). Implicature. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2010 ed.).

Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, *8*, 8-11.

Golland, D., Liang, P., & Klein, D. (2010). A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (p. 410-419).

Horton, W. S. (2008). A memory-based approach to common ground and audience design. In I. Kecskes (Ed.), *Intention, common ground, and the egocentric speaker-hearer* (pp. 189–222). Mouton de Gruyter.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage information density. *Cognitive Psychology*, *61*(1), 23–62.

Jäger, G. (2008). Applications of game theory in linguistics. *Language and Linguistics Compass*, *2/3*, 408–421.

Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. In B. Schlökopf, J. Platt, & T. Hoffman

(Eds.), *Advances in neural information processing systems.* Cambridge: MIT Press.

Lewis, D. (1969). *Convention*. Harvard University Press.

Oldenfield, R., & Wingfield, A. (1965). Response latencies in naming objects. *The Quarterly Journal of Experimental Psychology*, *17(4)*, 273–381.

Piantadosi, S., Tily, H., & Gibson, E. (2011). The communicative function of ambiguity in language. *Cognition*, *122*, 280–291.

Schlangen, D., & Fernandez, R. (2007). Speaking through a noisy channel: Experiments on inducing clarification behaviour in human-human dialogue. In *Proceedings of Interspeech 2007 (ICSLP2007).*

Stalnaker, R. (2005). Saying and meaning, cheap talk and credibility. In A. Benz, G. Jäger, & R. van Rooij (Eds.), *Game theory and pragmatics* (p. 83-100). Palgrave Macmillan.

van Deemter, K. (2009). Utility and language generation: The case of vagueness. *Journal of Philosophical Logic*, *38(6)*, 607-632.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.