NORTHWESTERN UNIVERSITY

A Game-Theoretic Model of Grounding for Referential Communication Tasks

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Linguistics

By

William Thompson

EVANSTON, ILLINOIS

December 2009

ABSTRACT

A Game-Theoretic Model of Grounding for Referential Communication Tasks

William Thompson

Conversational grounding theory proposes that language use is a form of rational joint action, by which dialog participants systematically and collaboratively add to their common ground of shared knowledge and beliefs. Following recent work applying *game theory* to pragmatics, this thesis develops a game-theoretic model of grounding that formalizes the core claims of grounding theory. This game-theoretic model is based on the concept of *signaling games*, originally proposed as a model of linguistic convention. In order to account for grounding, this thesis proposes to extend signaling games with an observation model, which allows for the possibility that the actions a participant takes may only be partially observable to others. This game-theoretic model is applied to the domain of *referential communication tasks*, a type of task commonly used in psycholinguistic experiments. The model generates predictions about how dialog participants in such tasks package referential expressions into installments, by calculating an optimal trade-off of cost and uncertainty. These predictions are experimentally evaluated with a novel variant of an online referential communication task.

Acknowledgements

When I started graduate school at Northwestern University in 1995, I had no inkling that it would be well over a decade before I would finish. After the passage of all these years, it is with great pleasure that I perform this final task of thanking all those who have made it possible to successfully conclude this chapter of my life.

First of all I would like to thank the members of my committee. To Stefan Kaufmann goes the credit for bringing me back to academia after a long absence as a research engineer at Motorola. I remember particularly well one of Stefan's seminars on semantics, in which I first learned about how game theory could be applied to problems in linguistics. Without the interactions I've had with Stefan over the past few years, this dissertation would not be what it is. To Darren Gergle I owe the inspiration to do experimental work in the area of conversational grounding. Darren's series of puzzle task studies served as a model for me to follow. His energy, enthusiasm, and feedback played an important role in the development of the ideas expressed in this dissertation. To Sid Horton goes my greatest debt of gratitude. I can honestly say that there is not a single aspect of this dissertation that hasn't been impacted by his intelligent and insightful input. Regular meetings with Sid provided a steady drumbeat of progress that kept me moving forward, even during those times when my energies flagged. Every graduate student should be lucky enough to have someone like Sid on their committee.

I have had many other influential teachers and professors over the course of my academic career. Going way back, Tom Otte and George Anderson first encouraged my interest in writing long papers in the 5th and 6th grades. More recently, my initiation into linguistics occurred at the University of Chicago, where I fondly remember the courses of Amy Dahlstrom, Karen Landahl, and James McCawley. My knowledge of linguistics has subsequently been considerably deepened at the University of British Columbia and Northwestern University, under the tutelage of Douglas Pulleyblank, Rose-Marie Déchaine, Janet Pierrehumbert, Beth Levin, and Chris Kennedy. I also owe thanks to David Traum, Matthew Stone, and Justine Cassell for teaching a fascinating course on the formal pragmatics of face-to-face conversation during the 2005 edition of ESSLLI at Edinburgh. This course is where I first entertained the idea of exploring installment noun phrases as a dissertation topic.

From among my colleagues in academia, I would like to single out James Sneed German, David Huffaker, Francisco Iacobelli, Paul Tepper, and William Turkel. These colleagues have provided many hours of stimulating conversation on a wide range of topics. I give special credit here to William Turkel for teaching me how to be a graduate student, and to Paul Tepper for getting me up to speed in current work on conversational agents. From among my colleagues at Motorola, I would like to single out Janet Cahn, Corey Miller, and Paul Davis. Paul Davis in particular has been a close friend and collaborator, who has provided me with strong encouragement over the past couple of years to complete this dissertation. Also from Motorola, I thank my past and present managers Harry Bliss, Steven Nowlan, Mark Randolph, and Bhavan Gandhi. They encouraged me to continue in my studies and provided me with the time to do so.

My friends and family are the ones that have kept me grounded during the process of writing this dissertation. Sincere thanks to my good friends Angie Alexander, Kathleen Conley, Alvin Gayles, Mike Mahany, RoseMary Mahany, and Alex Meininger. From my family, I thank in particular Ann Grandau, Berta Sigle, and Wesley Thompson. My mother-in-law Ann Grandau provided me with motivation whenever I needed it by offering me a "kick in the pants" to stop slacking. In similar fashion, my beloved Tante Bertl was always ready to offer me the sage advice "Erst die Arbeit, dann das Vergnügen". My brother Wesley Thompson provided me with a concrete and relevant example of someone who had successfully completed his Ph.D., despite having started his graduate studies well after I had.

Finally, I dedicate this dissertation to three special people in my life. First, to my mother Clara Thompson. From an early age, she instilled in me a deep love of learning. She will understand me when I say that this dissertation has been completed as part of the Doctorate Obtaining Division of the Adventure Club. Second, to my wife Laura Grandau. Laura is not just my wife, but also my best friend. Her constant love and support have meant everything to me. It is now my turn to return the favor, and to support her while she completes her dissertation work. Third, to my beautiful daughter Julie Berta Thompson Grandau. Watching her as she acquires her own linguistic knowledge has taught me new awe for the power and expressiveness of language.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Introduction

A convergence of ideas from philosophy (Austin, 1962; Grice, 1989; Lewis, 1969; Searle, 1969; Stalnaker, 1978), psycholinguistics (Clark, 1992, 1996), and pragmatics (Levinson, 1983; Horn and Ward, 2004) has sought to explicate a wide range of linguistic phenomena by claiming that they follow from general principles governing rational and cooperative behavior. From psycholinguistics in particular has come the widely influential *conversational grounding theory* of Herbert Clark and colleagues (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1987, 1989; Clark and Brennan, 1991; Clark, 1996). This theory proposes that language use in dialog is a form of rational joint action, executed by dialog participants in order to systematically and collaboratively add to their common ground of shared knowledge and beliefs. In this thesis, I address the issue of constructing a formal model of conversational grounding, with the ultimate goal of making the core ideas and intuitions of grounding theory precise enough to be used in a computational implementation of a conversational agent. To this end, I propose that *game-theoretic* models (Fudenberg and Tirole, 1991; Myerson, 1991; Osborne and Rubinstein, 1994) are particularly well-suited to the task. Game theory, to-date the dominant formal approach to analyzing multiagent interaction (Shoham and Leyton-Brown, 2009), provides well-suited analytical tools for formalizing the core ideas and intuitions of grounding theory.

The direction taken by this thesis is inspired by recent work applying game theory to linguistic pragmatics (Parikh, 2001; van Rooij, 2004; Benz et al., 2005; Jaeger, 2008).

These authors have sought to reconstruct the Gricean notion of the *cooperative principle* (Grice, 1989) using game-theoretic tools and analyses. Stalnaker (2005, pp. 87-88) describes the motivation behind this effort:

> As many people have noticed, Gricean ideas naturally suggest a game-theoretic treatment. The patterns of iterated knowledge and belief that are characteristic of game-theoretic reasoning are prominent in Grice's discussion of speaker meaning, and the patterns of strategic reasoning that Grice discussed in the derivation of conversational implicatures are patterns that game theory is designed to clarify.

The central claim of this thesis is that conversational grounding theory is similarly amenable to such reconstruction in game-theoretic terms (see also de Jaegher (2005, 2008)). Game theory provides a mathematically precise framework for describing multi-agent interactions that can clarify the core ideas of grounding theory, and their precise relationships to one another. Formalizing grounding theory in this way has two advantages. First, it sharpens its predictive strength with respect to empirically observed dialog behaviors. Second, it makes it possible to use grounding theory to created principled implementations of computational conversational agents, with the hope that such agents will prove superior to those implemented with ad hoc dialog strategies.

The focus of this thesis is on a class of task-oriented dialogs frequently used in psycholinguistic experiments, known as *referential communication tasks* (Krauss and Weinheimer, 1966; Clark and Wilkes-Gibbs, 1986; Horton and Keysar, 1996; Gergle et al., 2004a). Successful completion of such tasks requires participants to engage in information exchange that crucially relies on resolving referential descriptions about task objects.

The game-theoretic model I propose for such tasks is based on a class of games known as *signaling games* (Lewis, 1969; Spence, 1973; Crawford and Sobel, 1982; Stalnaker, 2005). A signaling game is a Bayesian game of asymmetric information, where an informed player must communicate with an uninformed player in order for the two to coordinate on an action. The informed player knows the true state of the world, and can send a signal to the uninformed player. The uninformed player must then select an action that determines the payoffs for both players. A standard signaling game provides a reasonable first approximation of typical referential communication tasks.

However, standard signaling games cannot directly account for grounding behaviors, because they don't allow for the possibility of a player having imperfect information about the communicative act of another player. In standard signaling games, the actions of the players are always completely observable to each other. Therefore, in order to model this type of uncertainty in communication, I propose a new type of signaling game, which I call *signaling games with partially observable actions*. This type of game extends signaling games with an observation model, which allows for the possibility that messages from the informed player are only partially observable to the uninformed player. The uninformed player receives observations of messages, which provide probabilistic information about the messages that generate them. A solution to a signaling game with partially observable actions is a strategy profile where the two players maximize their expected utility, taking into account the payoffs of outcomes, the costs of sending particular messages, and exogenously determined uncertainty regarding the content of messages.

This thesis experimentally evaluates some predictions of the game-theoretic model using an online version of a referential communication task. In this task, two participants worked together to identify and manipulate groups of two-dimensional shapes with features that vary along several dimensions, such as the color and configuration of their parts. The experiment was designed to elicit complex referring expressions from the participants, in the form of *installment noun phrases* (Clark and Brennan, 1991). During the experiment, both the cost of turn taking and the cost of task error were manipulated. The game-theoretic model predicts that these cost manipulations should affect how dialog participants package their contributions into installments, with installment frequency and length determined by calculating an optimal trade-off of cost and uncertainty. Results are reported in Chapter 5.

1.1. Thesis Overview

The rest of the thesis is structured as follows:

Chapter 2 provides a more detailed look at conversational grounding theory. It presents four core ideas of the theory: (1) language use is *joint action*, (2) joint actions are coordinated via the *common ground*, (3) the minimum amount of effort that dialog participants expend to add something to the common ground is determined by the *grounding criterion*, and (4) the maximum amount of effort that dialog participants expend to add something to the common ground is determined by the *principle of least collaborative effort*. The second part of the chapter is an overview of existing experimental work with referential communication tasks. It focuses on the contextual factors that influence how participants package referential descriptions into installments.

The game-theoretic model of grounding for referential communication tasks is presented in two parts. First, Chapter 3 provides necessary background material on game theory, leading up to a description of *signaling games*, which were first introduced by Lewis (1969) as a formal theory of linguistic convention. Signaling games provide a good first approximation of a model of communication in referential communication tasks, but they are not adequate to model grounding behaviors, because they do not allow for imperfect information with respect to the content of messages.

Chapter 4 builds on Chapter 3 by defining *signaling games with partially observable actions*, a model of communication that allows for imperfect information regarding the content of a message. Explicit connections are made between this formal model of communication and conversational grounding theory as presented in Chapter 2. The chapter continues by applying the model to a case study of factors that influence the size of referential description installments. The model predicts that varying the values of certain parameters, such as turn cost or task error cost, will cause dialog participants to vary the installment size of referential descriptions according to an expected value calculation that takes these costs into account. The chapter concludes with a comparison of the game-theoretic model to other formal approaches to conversational grounding.

Chapter 5 describes an experiment designed to test the predictions of the game-theoretic model. In the referential communication task used in the experiment, pairs of participants worked together online to identify and manipulate groups of two-dimensional shapes, with features that varied along several dimensions. The design of these stimuli was intended to elicit installment type referential descriptions from participants. In the experiment, both the cost of turn taking and the cost of task error were manipulated. As turn costs go up, the model predicts that dialog participants will generate longer, less incremental installments. Conversely, as task error cost goes up, the model predicts shorter, more incremental installments. These predictions are evaluated in the experiment.

Chapter 6 concludes the thesis by reviewing its theoretical and empirical contributions, and by presenting ideas for future research.

CHAPTER 2

Conversational Grounding Theory

Research in psycholinguistics and conversational analysis has revealed that a significant proportion of conversational behaviors are dedicated to actively keeping dialog participants coordinated: indicating understanding, making repairs, and soliciting feedback (Schegloff, 1968; Sacks et al., 1974; Clark and Wilkes-Gibbs, 1986). This observation has been the inspiration for conversational grounding theory, which has been developed by H. Clark and colleagues in a series of publications (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1987, 1989; Clark and Brennan, 1991; Clark, 1996). This theory has become widely influential in psycholinguistics, and has had an impact in computational linguistics as well (Traum and Allen, 1992; Traum, 1994; Allen et al., 2001).[1] This chapter provides an overview of the core ideas of grounding theory, and summarizes a series of psycholinguistic experiments that provide support for some of its claims.

2.1. The Core Ideas

Grounding theory takes the Gricean idea that language use is a type of rational action, and applies it to the mechanisms and processes by which dialog participants coordinate sequences of dialog actions. According to grounding theory, successful language use in dialog requires participants to coordinate their knowledge, beliefs, and behaviors on a variety of levels, from their attentional focus up through the ultimate goal of the activity

[1]Although not uncontroversial in the scope of its claims (see Traum, 1999; Keysar and Barr, 2005; DeVault and Stone, 2006).

to which the dialog is directed. Grounding theory further claims that participants act in such a way that maximizes the likelihood of achieving their goals, while simultaneously minimizing the amount of effort that they must expend in order to do so. Various formulations of the core ideas of grounding theory have been presented in the works of Clark and colleagues, but the core of the theory is succinctly captured by the following four claims:

(1) *Communicative acts are joint actions*: Conversation is a joint activity composed of a sequence of communicative acts, each of which is a joint action. Joint actions are composed of participatory individual actions that share a joint goal, and are executed in coordination with each other.

(2) *Coordination of joint actions is achieved through the common ground*: In order to succeed, communicative acts require participants to coordinate on both content and process. Coordination is achieved by relying on the common ground of knowledge and beliefs shared by the participants. Furthermore, a successfully executed communicative joint action advances the goals of the agents by incrementally adding to this common ground.

(3) *The minimum amount of grounding effort required to add something to the common ground is determined by the grounding criterion*: The participants in a joint action try to establish the mutual belief that the contributor has succeeded in adding to the common ground to a criterion sufficient for current purposes. Grounding is the collective process by which this mutual belief is achieved.

(4) *The maximum amount of grounding effort that will be expended to add something to the common ground is determined by the principle of least collaborative effort*: speakers and addressees try to minimize collaborative effort, the work both speakers and addressees do from the initiation of each contribution to its mutual acceptance.

Each of these claims is addressed in the following sections.

2.1.1. Joint Actions

The first claim is that language use in dialog is a specific type of joint action, executed by agents in order to advance the goals of a joint activity. A joint action is defined to be an action that is "carried out by an ensemble of people acting in coordination with each other" (Clark, 1996, p. 3). A joint action has constitutive parts – the "participatory" actions of the individuals involved. But these constituent parts must be coordinated with each other in order to be successful. They must be directed towards a joint goal, and they must be synchronized properly to achieve their desired effect. Joint actions may also be hierarchical, composed from smaller joint actions that achieve sub-goals of the higher-level joint action (Clark and Schaefer, 1987, 1989), and ultimately the goals of the joint activity in which the agents are taking part.

Clark and Wilkes-Gibbs (1986) provide evidence for this view of language use by looking at patterns of linguistic behaviors in a referential communication task. Example 2.1 is from an experiment in which participants work together to arrange a set of tangram figures. One of the participants (called "A") knows how the figures are to be arranged, while the other (called "B") must do the arranging. Notice how in this example a single referential description is split across multiple "installments", interleaved with an acknowledgement from B that provides feedback to A about the incremental description.

(2.1) **A.** And the next one is the one with the triangle to the right...

B. Okay.

A. With the square connected to it.

Clark and Wilkes-Gibbs (1986) found many other examples where the participants worked together in this fashion to identify a figure, with a significant percentage of the utterances serving as acknowledgements of understanding or requests for repairs. Many are marked with a question intonation, "trial marking" the utterance in order to elicit feedback from the addressee. Clark and Wilkes-Gibbs contrast these sorts of behaviors with the predictions of what they term the *literary model* of language use, which depicts an idealized speaker as generating a unitary speech act (Austin, 1962; Searle, 1969) uniquely identifying a referent in context, and an idealized hearer who simply extracts the relevant content from the referring expression in order to determine the identity of the referent. With examples like theses, they show that the literary model of reference does not do justice to the interactional nature of language use as shown in example 2.1, and misses the joint nature of the activity.

2.1.2. Common Ground

If communicative acts are joint actions, then how do participants coordinate on their execution? The answer, according to Clark, lies in their *common ground*, which is "the sum of their mutual, common, or joint knowledge, beliefs, and suppositions" (Clark, 1996, p. 93). This notion of common ground is derived from the work of philosophers Lewis (1969), Schiffer (1972), and Stalnaker (1978). A proposition P is said to be part of the common ground among a group of agents if it is *common knowledge* or *common belief* among the agents that P is true (Lewis, 1969). One way of characterizing common knowledge of a proposition P among a group of agents is as an infinite hierarchy of statements: everyone knows P, everyone knows that everyone knows P, and so on. Common belief is defined in a similar fashion, with *belief* replacing *knowledge* as the epistemic operator.[2] Since Lewis (1969) first provided an analysis of common knowledge, the notion has come to play an important role in fields that formally deal with multiagent reasoning, such as game theory (Aumann, 1976; Aumann and Brandenburger, 1995) and epistemic logic (Fagin et al., 1995; van Ditmarsch et al., 2007).

Common knowledge is an important concept because there are situations where having common knowledge is a prerequisite for achieving perfect coordination (Clark and Marshall, 1981; Rubinstein, 1989; Halpern and Moses, 1990). For example, Clark and Marshall (1981) argue that definite reference can succeed only if the intended referent of a referential description is common knowledge between speaker and hearer. To see that this is so, consider a referential expression e used by a speaker to identify a referent r to a hearer. To be successful, the speaker must believe that e will indicate r to the hearer. The hearer, on the other hand, will recover r from e only if he believes that the speaker believes that this is the referent that she intends for him to identify. Clark and Marshall (1981) demonstrate various scenarios where these kinds of *higher-order beliefs*

[2] Terminological confusion is rife in this field (see Lee (2001) for an attempt to resolve this confusion). In this thesis, I will use the terms *mutual knowledge* and *common knowledge* (and their belief variants) in the way that is standard in the game-theoretic literature (Geanakoplos, 1994). Mutual knowledge of a proposition P among a group of agents G holds if every member of G knows P. Common knowledge of P among G holds if everyone knows P, everyone knows that everyone knows P, etc.

are required up to an arbitrary level. At the limit, this implies that common knowledge of the intended identity of a referent is required to guarantee successful use of a definite referential description.

Clark and Marshall (1981) claim that this state of affairs gives rise to a paradox: common knowledge of the intended referent is required for successful use of a definite description, but agents with bounded resources cannot be expected to compute this infinite conjunction in finite time and with limited memory. Clark and Marshall consider two types of solutions to this "paradox": (1) arbitrarily limiting the level at which higher-order beliefs are calculated, and (2) using heuristics to infer common knowledge from shared bases, such as visual co-presence with the intended referent. They argue that the second type of solution is more plausible than the first, since there is no principled way of setting the limit on how many higher-order beliefs are computed. Even more fundamentally, it is quite implausible that agents perform any degree of higher-order thinking beyond two or three levels on any kind of regular basis. Statements like "A believes that B believes that A believes that..." rapidly become incomprehensible. Such statements are even more difficult to process when more than two agents are involved.

However, this paradox is not quite what it seems. Although the iterated higher-order beliefs formulation of common knowledge has been frequently taken as evidence that common knowledge is impossible to obtain for agents with limited resources (see Clark (1996) for discussion), there have been alternative definitions from early on that negate this problem. Lewis (1969) already recognized the potential problem of the iterated higher-order belief model, and suggested that iterated higher-order beliefs form a chain of implicit implications, and do not form a basis for a realistic model of how agents actually represent common knowledge. Soon thereafter, Aumann (1976) provided the first formal definition of common knowledge, and he used a set-theoretic representation that does not require an infinite number of deductive steps to decide whether or not a proposition P is common knowledge among a group of agents (see Section 4.2.2). Aumann's definition implies the iterated higher-order model, but allows for common knowledge to be computed in a finite number of steps for finite domains.

Even so, the paradox arises in another form. From the field of distributed systems in computer science, there is a well-known problem of common knowledge called the "coordinated attack problem" (Halpern and Moses, 1990; Fagin et al., 1995; Morris and Shin, 1997), related to the issue of establishing reliable communication protocols over a computer network. One version of the story goes as follows. Two divisions of an army, each commanded by a general, must decide on whether or not to attack the enemy. The attack fails if either the enemy is prepared, or if one of the divisions fails to make a move. One of the generals discovers that the enemy is unprepared, and sends a message to the second general to this effect. However, there is some chance that the messenger will be captured, and hence the second general will not receive this message. Even if the second general does manage to receive the message, he knows that the first general will remain unsure that the second general knows that the enemy is unprepared, and hence he sends back a messenger to the first general with a confirmation that he did in fact receive the message. Again, this confirmation may fail to arrive, and hence the first general must send a confirmation of the confirmation back to the second general. However, no matter how many confirmations of confirmations pass back and forth, the generals will never obtain full common knowledge that the enemy is unprepared, and without this common knowledge a coordinated attack may fail.[3]

It is a general result that any possibility for communicative error makes common knowledge of the communicated content impossible (Halpern and Moses, 1990). This result is particularly relevant for theories of communication in dialog, because human language is inherently ambiguous, susceptible to noisy channel problems, and errors in both production and comprehension. In these circumstances, where there is a certain likelihood (however small it might be in a given situation) that a communicative act can fail in some way, it is in principle impossible to achieve common knowledge. Fortunately however, it turns out that full common knowledge is not an absolute requirement to achieve coordinated actions among agents *on average*. Approximate common knowledge, in the form of probabilistic common beliefs, is sufficient in many cases of practical interest (Monderer and Samet, 1989; Halpern and Tuttle, 1993; Morris and Shin, 1997). This result opens the door to pursuing the idea that communication is joint action without requiring the impossibly high criterion of full common knowledge for the coordination of joint actions to succeed (see Section 4.2.2).

Common ground as probabilistic common belief is also fully consistent with Clark's vision of grounding theory. As mentioned above, Clark and Marshall (1981) argue that speakers and hearers solve the knowledge paradox by using "copresence heuristics" as evidence for common ground. Using these heuristics, evidence for concluding that a piece of information is in the common ground is derived from shared community membership,

[3]Rubinstein (1989) re-creates essentially the same problem in a game-theoretic setting, and similarly arrives at the conclusion that coordination can fail when common knowledge does not obtain, no matter how many levels of higher-order beliefs have been achieved

physical copresence, and "linguistic copresence". These types of evidence are not uniformly reliable, and they can be graded with respect to each other (for example, physical copresence is typically better evidence than linguistic copresence (Clark and Marshall, 1981)). As stated in Clark (1996, p. 98): "People tacitly evaluate shared bases for quality, recognizing that pieces of common ground range in likelihood from 0 to 1." If this is so, then probabilistic common belief is the best we can hope for in general. Furthermore, the very notion of a "grounding criterion" (Section 2.1.3) fails to make sense if we require full common knowledge for communication for successful coordination.

2.1.3. Grounding Criterion

Now that we know what the common ground is, and its role in the coordination of joint activities, what is *grounding*? Grounding is the process by which new information is added to the common ground. This involves both the initial presentation of this information, and any extra work that follows in order to make sure the initial presentation was mutually understood. In general, the effects of communicative actions are on the mental states of addressees, to which a speaker does not have direct access. Therefore, there is no sure way for a speaker to know with absolute certainty that an utterance has been understood as was intended. There is always a chance that an error in understanding has occurred, and this possibility is what motivates the existence of grounding behaviors, actions that are used by dialog participants to decrease this possibility (Traum, 1994).

The minimum amount of effort required in order to reach a sufficient degree of confidence in the success of a communicative act is determined by the *grounding criterion*. In dialog, participants expend as much effort as required so that "the contributor and his or her partners mutually believe that the partners have understood what the contributor meant to a criterion sufficient for current purposes" (Clark and Brennan, 1991, p. 129). For Clark, grounding is the collective process by which dialog participants try to reach this mutual belief.[4]

The perceived importance of the current purposes of a dialog determines a floor on how much grounding must be performed, i.e., the minimum degree of certainty the participants require that mutual understanding has been achieved. A casual dialog about the weather is not as important as a 9-1-1 call to emergency services, and standards of grounding will be correspondingly different. We expect to see more extensive grounding of information in the latter case than in the former. Within a single dialog, we expect grounding behaviors to adapt dynamically to the importance of a particular exchange. In collaborative task-oriented dialog, the importance of a particular dialog contribution is often determined by characteristics of the task to which the talk is directed.

Grounding behavior is shaped not only by the relative importance of the information being grounded, but also by the communication media available to dialog participants (Clark and Brennan, 1991). Communication media, such as audible spoken language or visually observable gestures and actions, possess different properties that impact the cost and reliability of the information communicated. For example, speech is ephemeral, and generally only one dialog participant can speak at a given time. Visual feedback, however, while still potentially ephemeral, can be provided in parallel with speech. The visual medium can also be cheaper to use and more reliable in its communicative effects

[4]Clark is not always clear about what he means by the term "mutual belief", but in general, it corresponds to what in this thesis I call "common belief", following standard usage in game theory and in epistemic logic. See Section 2.1.2 for discussion.

(though this depends on the particulars of the situation), and these properties often make additional explicit linguistic feedback redundant. In short, a communication medium can affect grounding behavior by how much it "costs" to use, how much noise or error it introduces, by the granularity of action it permits, and whether or not its use by one participant locks up resources that forces synchronization (preventing overlap) of communicative actions.

A rational dialog participant must make a decision about the best type of contribution to make based on a calculation of the costs incurred, and the requirements of the current purpose of the dialog. How does a dialog participant deal with these constraints? According to (Clark and Brennan, 1991), they must continuously make *trade-offs* on the different kinds of costs they incur, choose the available medium that bests serves the goals at hand, and make sure that the contribution that is made is cost effective with respect to the purposes of the dialog. I show in Chapter 4 that these kinds of calculation make perfect sense from the viewpoint of a *Bayesian decision maker* who is interested in maximizing their utility. Such a decision maker will examine the possible outcomes of their actions with their associated payoffs, look at the costs involved in achieving these outcomes, and select the action that achieves the maximum expected value of payoff minus costs. Doing this calculation in a multiagent setting is precisely the domain of game theory, as we will see in Chapter 3. Thus, the presentation of grounding theory as given by Clark and colleagues leads us directly to game theory as an appropriate formalization of its core ideas and intuitions.

2.1.4. Least Collaborative Effort

The grounding criterion establishes a floor on the effort required to ground a communicative act. What prevents us from putting in extra, potentially unnecessary effort? Why not ground a contribution indefinitely in order to achieve arbitrary closeness to full common knowledge? This type of unnecessary effort is prevented by the *principle of least collaborative effort*: "the participants try to minimize their collaborative effort – the work both do from the initiation of each contribution to its mutual acceptance" (Clark and Brennan, 1991, p. 135).

This principle is not as trivial as it might first seem. The key here is the word "collaborative". A consequence of the fact that each participant is responsible for minimizing collaborative effort is that dialog participants share responsibility for making sure that a dialog contribution is as efficient as possible. This means, for example, that speakers will design their utterances with the beliefs and capabilities of the addressee in mind. If the speaker perceives that the addressee is distracted (such as driving a car), or that communication is over a noisy channel, they should adjust their contribution accordingly, such as timing the utterance for when a listener can dedicate full attention to the speaker. Addressees should also provide the most efficient grounding feedback possible, given the demands of the grounding criterion. In some situations this might be a simple continuer (such as "okay", or a head nod), in other circumstances it might be a more costly verbatim repetition of what they heard (Clark, 1996). The more costly types of feedback will only be used in case it is the only course of action that satisfies the grounding criterion.

In this thesis, I will focus on how this principle plays out with respect to what Clark (1996, p. 235) calls "packaging":

> Packaging is always an issue in contributing to discourse: How large a contribution should the two participants try to complete if they are to minimize their joint effort? If there were a presentation and acceptance phase for each word separately, conversation would double in length. On the other hand, if each contribution were a paragraph long, a minor misunderstanding at the beginning might snowball into a major misunderstanding by the end. With limited working memory for what the speaker said, the two people would have great trouble repairing it. The optimal size of a contribution ought to be somewhere in between.

Example 2.1 already showed us one instance of dialog where a speaker decided to package up a referential description into multiple installments. Example 2.2 is similar, this time taken from (Cohen, 1984). In this dialog, two participants are engaged in the task of collaboratively building a toy water pump. One subject (called "S") knows how the pump should be put together, while the other subject (called "J") must do the actual assembly.

(2.2) **S:** Okay now, the small blue cap we talked about before?

J: Yeah.

S: Put that over the hole on the side of the tube—

J: Yeah.

S: —that is nearest to the top, or nearest to the red handle.

In the first contribution, S refers back to an object that has been discussed before. S's utterance here is marked with question intonation, prompting J to respond with an acknowledgment of understanding, which J provides. Notice here that the goal of this utterance is pure referent identification – there is no predicate per se (see the discussion of these types of communicative acts in Cohen (1984)). The predicate that applies to this referent is supplied in the next contribution from S, which is split over two dialog turns. S pauses after the first installment, again prompting J to respond with an acknowledgement of understanding of the contribution-so-far. Only after this acknowledgment is received does J proceed with the next installment of the referring expression. In both of these cases, S is "projecting" to J what kind of evidence S wants in order to incrementally ground the contributions (Clark, 1996).

In this example, S could have made the dialog contribution differently. The speaker could have eliminated the pause and generated the entire referring expression as a single *utterance unit* (Traum and Heeman, 1997). However, by splitting the referring expression into two installments, the speaker was able to gain confidence that the hearer understood the speaker's intent before completing the contribution. It appears that in this particular dialog situation, the perceived benefit from incrementally grounding the referring expression out-weighed the cost incurred by forcing an extra dialog turn.

Packaging up a contribution into multiple installments is something that people routinely do in dialog, but one that has received little attention in computational approaches to dialog systems. For example, an influential account of referring expression generation takes the approach that referring expressions are packaged as single utterances that uniquely identify a referent in context (Dale and Reiter, 1995). Examples 2.1 and 2.2 show that this result obtains for stretches of dialog containing multiple installments, but not necessarily for individual utterances taken in isolation. To the extent that the issue of packaging is dealt with at all in computational work on dialog, it is done in an ad hoc way with hand-crafted rules. It would clearly be beneficial for the performance of a computational conversational agent if it could do this sort of packaging automatically, based on characteristics of the dialog and task context.

Example 2.3 shows yet another case of packaging, this time taken from Gergle (2006). In this experiment, pairs of participants collaborated to solve an online puzzle task. Here, visual information on the state of the puzzle was available to both the helper (called "H") and the worker (called "W"):

(2.3) **H:** OK, and the orange

W: [moved correct piece]

H: Um, touching the right corner, right top corner of the dark blue.

Once again, we see a referring expression broken into two installments. However, this time the hearer did not respond with an explicit grounding speech act acknowledging understanding, but let the action of moving the piece speak for itself. In this case, visual feedback from the worker's actions makes linguistic feedback redundant – it is common knowledge to both participants that the helper can see what the worker is doing, and he therefore lets his actions do the talking. Extra verbal feedback is unnecessary, and is therefore eschewed. The principle of least collaborative effort predicts precisely this sort of behavior.

In the remainder of this thesis, I focus on the issue of packaging in order to demonstrate the utility of a game-theoretic model of grounding. According to grounding theory, installment size is determined by a speaker who is (1) jointly acting with the addressee, (2) in order to add new information to their common ground, (3) while respecting the grounding criterion, and (4) attempting to minimize collaborative effort in the process.

This process, as described above, should be sensitive to the means of grounding at the disposal of the participants, and the demands of the task they are trying to perform. Chapters 3 and 4 pursue this line of thought further.

If dialog participants were to adhere rigorously to both the grounding criterion and the principle of least collaborative effort, we should expect to see optimal dialog behaviors – optimal sequences of initial contributions and subsequent grounding actions. Therefore, conversational grounding theory as presented is a *normative* model of dialog, much as game theory is a normative model of strategic multiagent interaction. Of course this is an idealization. What we expect to see in reality might be an approximation of such optimal behavior, or a tendency to act in a way that accords with these general principles. The degree to which we in fact do approximate the ideals of grounding theory is an open question, partially addressed by the empirical studies reviewed in Section 2.2, and the experiment described in Chapter 5.

2.2. Referential Communication Tasks

Many psycholinguistic studies have taken the form of referential communication tasks, originally developed by Krauss and Weinheimer (1964, 1966, 1967). These tasks involve pairs of people that must collaboratively identify a set of objects – typically to arrange them into a particular sequence or configuration. In the basic paradigm, one of the pair takes the role of the *helper*. She is informed about the target sequence or configuration of objects which serves as the goal of the task. However, the helper is unable to manipulate the objects herself, and must communicate this goal state to the other subject, the *worker*. The worker has a complementary role in the task – he is initially uninformed about the goal state, but he has the ability to manipulate the task objects. In order for him to manipulate them correctly, the helper and the worker must engage in dialog to identify the target objects and the configuration into which they must be arranged. Successful completion of these tasks therefore requires the helper and the worker to engage in information exchange that crucially relies on generating and resolving referential descriptions about task objects.

The remainder of this section is a review of selected studies that have used referential communication tasks to investigate properties of language use. Many of these studies have been performed explicitly to test qualitative predictions of conversational grounding theory. This review focuses on this aspect of these studies, and pays special attention to the factors that have been found to influence how participants package referential descriptions into installments. The experiment described in Chapter 5 is inspired by and builds upon this body of work on referential communication tasks.

2.2.1. Building a Water Pump

Cohen (1984) is one of the first studies to focus on a fine-grained analysis of the linguistic behaviors dialog participants use when they are given different communicative media to use. This study describes a referential communication task where pairs of participants collaborate in order to build a toy water pump. The helper was informed about how the assembly was to occur, while the worker who had to perform the task was instructed by the helper on how to do so. The study employed several different modalities for communication between the helper and the worker, but the focus of the analysis is on audio communication (over a telephone) versus keyboard communication (chat-style). The resulting dialogs were coded with a set of speech act types as shown in Table 2.1.

Table 2.1. Speech act typology used in Cohen (1984)

Communicative Act	**Example**
`Request(Assembly Action)`	"put that on the hole"
`Request(Orientation Action)`	"the other way around"
`Request(Pick-up)`	"take the blue base"
`Request(Identify-Referent)`	"there is a little piece of rubber"
`Request(Inform([relation]))`	"and you've got the base on it?"
`Request(Informif([Identified-referent]))`	"got it?"
`Request(Achieved([relation]))`	"and the purpose of that is to cover up that hole"
`Label`	"that's a plunger"

Cohen (1984) found that speech-based dialogs differed markedly from keyboard-based dialogs in the distribution of the types of speech acts found. Speech-based dialogs were more "granular" and incremental, containing large numbers of `Request(Identify-Referent)` acts. We already saw an instance of this behavior in Example 2.2. This example turns out to be typical of speech-based dialog in this study. Consistent with this finding, Cohen (1984) also found that speech-based dialogs had on average over twice as many independent request acts (of the types shown in Table 2.1) than chat-based dialogs. On the other hand, chat-based dialog produced larger chunks of text, which contained more or less complete instructional units. This is demonstrated in the following stretch of chat-based dialog (Cohen, 1984):

(2.4) **B:** put the pink valve on the two pegs in that blue cap...

N: ok

B: now, put the little blue cap over the hole in the large tube near the plunger handle...

J: ready

S: forgot one thing…use the red thing that looks like a nail to plug the plunger so it will work…

The contributions of the helper in example 2.4 (called "B") are less granular than is typical for the speech-based dialogs. Referring expressions are embedded in the larger instructional units, thereby backgrounding them in the conversation. Cohen (1984) provided no explanation for why the modality of communication impacted packaging in this way, but this result is not unexpected given the predictions of grounding theory. Turn taking is costlier in chat-based dialog than it is for speech dialog, reducing the attractiveness of turn taking behaviors. Conversely, text is persistent in ways that speech is not, so incremental grounding is less important. These two facts suggest why it is that speech-based dialog involved more fine-grained behaviors in this study.

2.2.2. Arranging Tangrams

In a classic referential communication task, Clark and Wilkes-Gibbs (1986) performed an experiment in which pairs of participants worked together to arrange abstract tangram figures into a target configuration. For each trial, the helper viewed twelve tangram cards arranged in a sequential order. The worker had replicas of the same figures, but his were randomly organized into a matrix at the start of the trial. The task of the helper and the worker was to collaborate in order to get the worker's cards into the same order as the helper's cards. In order to accomplish this, they communicated with each other using speech. Visual contact between them was prevented by the presence of an opaque screen.

Table 2.2. Noun phrase types from Clark and Wilkes-Gibbs (1986)

Type of noun phrase	**Example**
Elementary	"Number 4's the guy leaning against the tree."
Episodic	"Number 7's the goofy guy that's falling over, with his leg kicked up"
Installment	see example 2.1
Provisional	"And the next one is also the one that doesn't look like anything. It's kind of like that tree?"
Proxy	A: "And number 12 is, uh, . . . " B: "Chair"

The resulting transcripts from the trials were transcribed for various communicative behaviors, including turn taking, back-channel responses, repairs, false starts, and intonational features. The main types of noun phrases used by helpers are shown in Table 2.2. Of these noun phrase types, the elementary type was the most common, and it became even more common in later trials than in earlier trials. This type of noun phrase requires fewer words on average than the other types, and its predominance in later trials is a major reason for the finding that participants used fewer words to complete the task for later trials. The other types of noun phrases – episodic, installment, provisional, and proxy – were all used to some degree as well, though they decreased in frequency for later trials.

Clark and Wilkes-Gibbs (1986) claim that each of these noun phrase types *projects* what kind of response the addressee should make, setting different standards for the kind of feedback desired. Elementary and episodic noun phrases project implicated acceptance, installment and proxy noun phrases project explicit acceptance, while provisional noun phrases project self-expansion. Adding a "try marker" (question intonation) to these noun phrase types projects a request for an explicit verdict from the addressee. In general, the speaker will project the amount of evidence he desires from the addressee, in order to satisfy the demands of his grounding criterion (Clark, 1996). But projecting more costly feedback from an addressee is generally only done for a reason, because the extra time taken for the explicit feedback takes time and processing resources. This is in accordance with the principle of least collaborative effort.

Focusing on just the episodic and installment noun phrase types, Clark and Wilkes-Gibbs (1986) distinguish between them on the basis of the evidence they project. Episodic noun phrases do not project implicated acceptance, while installment noun phrases project explicit acceptance. Other than this, the two types are the same in that they result from packaging a single referential description into multiple chunks, which are uttered by the speaker in sequence. For the rest of this thesis, I do not distinguish between these two types of noun phrase, because the difference does not figure into the analysis. I adopt the term "installment" to refer to a portion of a referential description that is uttered by a speaker as an independent dialog turn, regardless of whether it projects implicit or explicit acceptance by the hearer.

2.2.3. Driving Directions

Brennan (1990, 2005) describes a collaborative matching task in which pairs of participants worked together to maneuver a car icon on a map to a target location. The helper had private access to the target location on her view of the map. The worker used a mouse in order to maneuver the car to the target location, following the directions of the helper.

In one condition, the helper could see the location of the worker's car icon in real time, as the worker maneuvered the car with the mouse. In another condition, there was no view of the worker's car icon. For each trial, it was common knowledge between helper and worker which condition they were in. An action-language log was created from each trial, recording utterances, mouse clicks, and x-y coordinates of the car icons. From these logs, action transcripts were created that graphed the distance between the target and worker icons as a function of time. This was used as an indicator of worker comprehension of helper instructions. For a subset of trials, the language of the participants was transcribed and aligned with the action transcripts.

Task performance, measured in time from start to completion, was more than twice as long in the verbal-only condition than when visual evidence was also available. Fewer than half as many words were spoken in the visual condition as in the verbal-only condition. The worker spent almost four times as much time in the verbal-only condition verifying the correctness of the worker's icon position at the end of the task. In general, workers used many fewer backchannels and acknowledgments in the visual condition than in the verbal-only condition. Responsibility for verifying the correctness of the worker icon position shifted from the worker to the helper when it was common knowledge that the helper could see the worker's icon. At times, the helper stopped mid-utterance to generate a timely deictic cue, based on real-time visual feedback from the worker's icon. Many exchanges in the visual condition involved no speech from the worker at all, who tended to let his actions do the talking.

The grounding criterion and the principle of least collaborative effort can be invoked to explain these findings. Visual feedback is cheap and reliable. If the more expensive verbal behaviors aren't needed to satisfy the grounding criterion, then least collaborative effort tells us they won't be used. Without visual information however, the grounding criterion requires additional verbal grounding behaviors in order for the dialog participants to achieve a sufficient degree of confidence that common ground has been achieved. This study therefore provides support for a conversational grounding theory explanation of dialog coordination behaviors.

2.2.4. Solving Puzzles

Gergle, Kraut and Fussell (Kraut et al., 2002; Gergle et al., 2004a,b) performed a series of experiments that involved pairs of participants working together online to solve a virtual jigsaw puzzle, such as the one shown in Figure 2.1. For each trial the helper had private access to a solved version of the puzzle, and her task was to provide verbal instructions to the worker on how to build the puzzle. The worker's task was to follow these instructions by selecting from a set of graphically represented puzzle pieces and moving them into the workspace in the proper configuration. In a subset of trials the helper shared a view of the worker's workspace. In other trials, the helper had no view of the worker's workspace, and their interaction was restricted to audio communication only. From the worker's point of view, the only difference across the shared and not-shared conditions was whether or not he was aware that the visual workspace was shared with the helper. As an orthogonal condition, the experimenters also manipulated the visual complexity and temporal dynamics of the puzzle, on the hypothesis that this type of complexity would enhance the value of visual information.

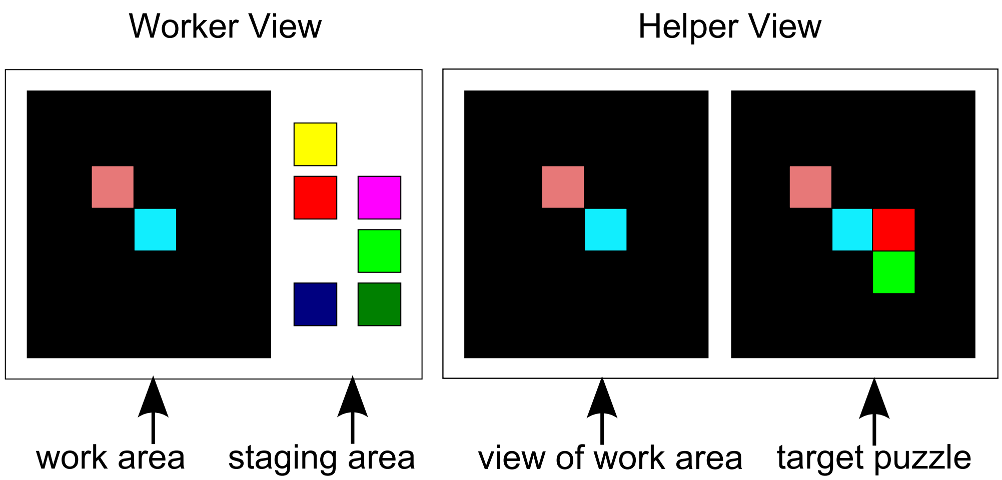

Figure 2.1. Puzzle task helper and worker views

Both communication efficiency and communicative processes showed differences across the shared vs. not-shared condition. The pairs were about a third quicker at solving puzzles, with fewer words per unit of time, when a high fidelity view of the workspace was shared, and this difference was greater when task complexity was higher. The worker's efficiency in this respect was affected to a greater degree than the helper's. Figure 2.2 (from (Gergle et al., 2004c)) illustrates how language changes when a shared visual workspace is available. When the shared visual workspace is available, responsibility for checking the state of the puzzle shifts from the worker to the helper, there are fewer verbal grounding acts, fewer words, and fewer turns.

Gergle et al. (2004c) performed a related study using a chat window instead of an audio channel. The independent variables of the study were dialog history persistence (one turn versus six turns), linguistic complexity (simple colors versus more complex plaids), and workspace visible (shared vs. not-shared). As predicted, lexical complexity,

Shared Visual Workspace

H: okay all blue with 2 vertical darker stripes
H: mostly grey with 1 bottom red
W: [moves piece in]
H: S of blue
W: [positions piece]
H: grey with 1 horizontal white stripe
W: [moves piece in]
H: NW of blue
W: [positions piece]

No Shared Visual Workspace

H: the first block we need has 1 white stripe at the very top...with a thinner yellow stripe about a cm below it
W: [moves piece in]
W: ok, got it
H: the other one we need has a thin yellow stripe, then a thick white one 1 cm below, then another thick white one ~3 cm below, then another thin yellow stripe 1 cm below
W: ok, does that have a faint blue stripe in the center?
H: it should be divided down the center...with a plaid light blue diagonal stripes on the left half, and dark blue solid on the right half
W: [moves the piece in]
H: do you need more clarification?
W: no, I've got the second piece

Figure 2.2. Sample puzzle task transcripts from (Gergle et al., 2004c)

shared visual workspace, and chat persistence all had an impact on task performance. Pairs were substantially faster in the condition where solids were used instead of plaids, they were over twice as fast when there was a shared visual workspace, and there was a small but reliable increase in performance when there was a larger chat dialogue history. There was an interaction between these conditions – having a shared visual workspace and a larger dialogue history mattered more when the lexical complexity was higher. On the other hand, having a larger dialogue history mattered only when there was no shared visual workspace, suggesting that the improvements in grounding provided by a shared visual workspace overwhelmed any further advantages from the larger dialogue history.

Having a shared visual workspace had a large impact on the structure of dialogs. Without a shared visual workspace, pairs used almost 2.5 times as many words to complete a puzzle, and nearly twice as many utterances. This is due to the fact that helpers could infer workers' comprehension by directly observing the effects of physical actions on the puzzle pieces, allowing the helpers to create shorter, more incremental descriptions that could be cut short when the physical evidence made them superfluous. Without a shared visual workspace helpers became far more explicit with their directions, and workers became more explicit in declaring their state of understanding. This latter effect is particularly significant from the viewpoint of conversational grounding theory – the only thing that changes for the worker in the no-shared workspace condition is that he knows (and knows that the helper knows) that his actions are not being observed (cf. Brennan, 2005).

Having a persistent dialog history also affected dialogue structure, though not to the same degree as having a shared visual workspace. With a dialog history, pairs generated more utterances, with fewer words per utterance, although the total number of words was approximately the same as the no dialog history condition. Example 2.5 shows a stretch of dialog in the dialog history available condition.

(2.5) **H:** fourth has a light blue cross

H: with green on the top

H: and black on the bottom

W: with dark right next to the light blue

H: yes

W: ok, where

In this example, the helper generated small, incremental referential installments, and waited for feedback from the worker before stopping. This was made possible by the fact that the helper knew that her incremental descriptions would persist in the worker's view, available for review. Because the helper was uncertain in advance how much descriptive content would be necessary to actually identify the piece, sending out incremental descriptions allowed her to minimize collaborative effort by permitting the worker to cut her short when sufficient information to identify the referent had been transmitted.

2.2.5. Bicycle Repair

Kraut et al. (2003) conducted a series of studies, examining the impact of a shared visual workspace on a collaborative physical task. In these studies, sets of two participants engaged in bicycle repair tasks, such as attaching a bicycle seat to a frame. The helper in each trial was an expert on bicycle repairs, and her task was to assist a novice worker by providing instructions on how to perform the repair. The experiment had three conditions: (1) a side-by-side condition, where the participants were co-located, (2) an audio-visual condition, where the helper and the worker were physically remote from one another, and the worker was outfitted with a head-mounted camera and a full duplex audio link with the helper, and (3) an audio-only condition, where the helper and the worker were physically remote and shared only an audio link.

Dependent measures included task performance (completion rate, completion time, and repair quality), and conversational efficiency (number of utterances, types of speech acts generated). The experimenters found that participants in the side-by-side condition scored highest on task performance measures and had the most efficient dialogs, in terms of number of utterances. They found no significant difference in these measures between the audio-visual and audio-only conditions. [5] However, an analysis of the conversations showed that when visual information was available, workers spent less effort in describing the state of the task, and less effort in describing their internal state of understanding. With video information, helpers were more likely to proactively give assistance without explicit requests for help, and also more likely to proactively provide clarifications of prior instructions, probably because they had visual evidence when workers were having comprehension problems. Once again, visual information had a pervasive effect on the nature of the linguistic contributions to the dialog.

2.2.6. Building Lego Models

Clark and Krych (2004) also investigated the impact of a shared visual workspace on grounding behaviors. In each trial of the experiment, the helper had private access to a target Lego model, and the worker's task was to assemble a set of Lego pieces into a configuration that matched this target model. In one condition, both participants could see the worker's workspace, while in another condition the helper could not see the workspace. The purpose of this manipulation was to determine how verbal and visual feedback from workers would affect helpers' dialog contributions, and conversely how the knowledge that worker actions were (or were not) visible to the helper would affect the workers' contributions.

Several outcome variables were analyzed, including completion times and error rates, numbers of words and turns, and number and types of gestures, and their timing with

[5] This result contrasts with those in the other studies described in this section. The experimenters hypothesized that this unexpected result was due to limitations in the visual link technology they employed.

respect to utterances. Unsurprisingly, task performance improved when a shared visual workspace was available. Tasks took significantly less time to complete, and both helpers and workers used fewer words and turns overall. Interestingly however, helpers used more words per "turn" when the workspace was visible, and workers used fewer words per turn. Many of the exchanges looked like the following:

(2.6) **Doris** Take a short blue

Betty [Retrieves a short blue block]

Doris [Looks at Betty's block.] Put it at the end of the yellow close to the green.

Betty [Places the blue block on the yellow block.]

Doris [Looks at result.] Take a . . .

With a visible workspace, workers relied upon the ability of helpers to observe their actions as evidence for understanding, in many cases eliminating the need for the worker to verbally confirm or request clarification of the helper's instructions. This meant that a helper could complete large chunks of the task with no intervening dialog turns from the worker. There were many cases of a fine-grained coordination of behaviors between helper and worker, with verbal directions from the helper interleaved with physical actions from the worker.

A shared visual workspace provides the ultimate in incremental grounding – the visual updates provide continuous, uninterrupted feedback on the state of the task, what the worker is doing, and by proxy how well the worker understands the helper's directions. The visual channel operates in parallel with the audio channel, hence information across the modalities does not compete for resources, and allows for asynchronous (simultaneous) behaviors to occur. One way of looking at Example 2.6 is that the helper is receiving the most incremental form of feedback possible – an almost word by word grounding of the addressee's understanding of the helper's utterance. Given the powerful nature of visual evidence, the sharp reduction in linguistic feedback from the worker is predicted by the principle of least collaborative effort, since it is in many cases superfluous.

2.3. Summary

The studies reviewed in this section have focused primarily on how performance in referential communication tasks is affected by the choice of communication modality. These modalities included text-based chat, speech, and visual workspace. Different patterns of grounding behaviors resulted from adopting different modalities, providing qualitative empirical support for aspects of grounding theory. Text-based chat communication typically results in the longest installments of the modalities studied, presumably because turn taking is relatively expensive, and because the persistent nature of text provided the recipient of a message plenty of time to review the longer messages and extract the relevant content. Speech-based communication led to shorter increments than for text, with more incremental grounding. This is predicted by grounding theory if turn taking is cheaper for speech than it is for chat, and the ephemeral nature of speech placed greater demands on the recipient's memory. Finally, when there is a shared visual workspace incremental grounding is taken to its extreme, and linguistic feedback from addressees is greatly diminished.

In addition to manipulating the communication modality, the puzzle task studies also manipulated properties of the task itself, by varying stimulus complexity and task difficulty. These manipulations affected the structure of dialog in ways predicted by grounding theory. Stimuli and tasks that are more complex to encode lexically led participants to rely more heavily on visual feedback. However, making visual feedback less useful (for example, by introducing a delay in visual feedback) led participants to return to the speech modality for providing feedback. The puzzle studies therefore showed (as grounding theory predicts) that properties of the task and of the communication modality interact in non-trivial ways to affect the participants' communicative behaviors.

These studies provide support for general predictions of grounding theory. However, many questions remain to be answered. For example, can we disaggregate turn cost from wholesale changes in the choice of communication modality? According to grounding theory, turn cost is the primitive notion that is used to explain differences among modalities in terms of the grounding behaviors that result from using them (Clark and Brennan, 1991). If costs are the primitives that explain these differences, then it shouldn't be necessary to make wholesale changes to communication modalities (such as using chat vs. speech) in order to obtain them. It should be sufficient to manipulate costs *within* a single communication modality in order to obtain different grounding behaviors.

Task "importance" is another primitive of grounding theory. Is it therefore possible to change a dialog participant's grounding criterion (and therefore grounding behaviors) by directly manipulating the reward or cost of the task actions he executes? For example, can we affect dialog behavior by making particular task actions relatively expensive, or by introducing "bad" consequences from making task errors? According to grounding theory, these types of manipulation should also affect grounding behaviors, such as the incrementality of referential expression installments. However, the experiments described in this section do not test this prediction.

Finally, how do we formalize notions like "cost", the "grounding criterion" and "least collaborative effort" in such a way that they could be implemented computationally? If grounding theory is really a useful way to describe and predict human behavior in dialog, it should also prove useful for computational conversational agents. Conversational agents now interact with many thousands of people on a daily basis, in the form of automated speech help systems and multimodal user interfaces to applications (Paek and Pieraccini, 2008). All of these systems rely on dialog management software to make decisions about what kinds of grounding actions to execute. If more of these decisions could be made in a principled fashion, based on properties of the modalities being used and properties of the task to which the dialog is directed, this could potentially result in conversational agents that perform better (and are more flexible) than those implemented with ad hoc dialog strategies.

CHAPTER 3

Games and Communication

This chapter presents necessary background material on game theory and how it can be used to model communication among agents, eventually leading to a formal model of conversational grounding. Section 3.1 introduces *extensive games with perfect information*, a type of game that provides an explicit model of players that interact by taking sequential actions, such as two people taking turns in a dialog. However, an extensive game with perfect information does not model players that have private information, and without the possibility for private information, meaningful communication cannot occur. Therefore, in Section 3.2, I turn to a well-known game-theoretic model of communication, called *signaling games*. Signaling games are based on an extension to extensive games with perfect information that allows players to have private information. A signaling game is a two player game where one player knows the true state of the world, and she has the ability to generate a signal that potentially communicates information about this state to the other player, who then takes an action that determines the payoff for both players. This model is a good first approximation of a referential communication task, although (as I show) it is inadequate as a model of grounding for such tasks.

3.1. Sequential Games

Game theory is the mathematical study of strategic, multiagent interaction.[1] A game consists of any situation in which an agent must decide among a set of alternative actions to take, the agent has preferences over the outcomes of these actions, and the outcomes that occur depend at least in part on the choices that other agents make. The most basic definition of a game is called a *strategic* game, and it consists of the following components (Osborne and Rubinstein, 1994):

Definition 3.1.1. A (finite, n-person) *strategic game* is a tuple $\langle N, A, (\succsim_i)\rangle$, where:

- N is a finite set of *players*, indexed by i;
- $A = \times_{j\in N} A_j$, where A_j is a finite set of *actions* available to player j. Each vector $a = \langle a_1, \ldots, a_n\rangle \in A$ is called an *action profile*;
- For each player $i \in N$, there is a *preference relation* $\succsim_i$ over $A = \times_{j\in N} A_j$.

A strategic game consists of a set of agents, a set of action profiles containing an action for each agent, and a set of preference relations over action profiles, with one preference relation for each agent. Given a small set of standard and natural assumptions about these preference relations (such as transitivity), it is common to represent them using numerical *utility* (or *payoff*) functions, which map from action profiles to the set of real numbers (Myerson, 1991). For each preference relation $\succsim_i$ we can create a function $u_i : A \mapsto \mathbb{R}$ such that whenever $a \succsim_i a'$ we have $u_i(a) \geq u_i(a')$. [2]

[1] In this thesis, I have relied on the overviews of game theory provided by Myerson (1991), Osborne and Rubinstein (1994), and Shoham and Leyton-Brown (2008). For the relationship of game theory to linguistics, I have been guided by the chapters in Benz et al. (2005). The formal definitions in this chapter are based primarily on those from Osborne and Rubinstein (1994).

[2] There are an infinite number of ways to assign utility values to action profiles that result in equivalent games, where "equivalent" means that they result in the same set of solutions to the game. Therefore,

Crucially, each player's preference relation (or utility function) is defined over the set of action profiles, and therefore the payoff to an individual player depends upon the actions taken by all players, not just their own action. Defining preferences over action profiles instead of individual actions is what differentiates game theory from single agent models of decision making. To a first approximation, we can identify action profiles $a \in A$ as analogs to Clark's notion of *joint actions* (see Section 2.1). The individual "participatory actions" of a joint action (Clark and Schaefer, 1987, 1989) can be mapped to the individual components of each action profile, selected from each player's set of individual actions A_i.

Since each player's payoff depends not only on their own action, but also on the actions of the other players, the best choice of action for player $i \in N$ depends upon what i believes the players $j \in N\backslash\{i\}$ will do. The actions of each of these other players in turn depends on *their* beliefs about what all the others will do, including player i. In this way, it already follows from Definition 3.1.1 that patterns of iterated knowledge and belief are relevant to how a perfectly rational player will choose her actions. Once again, we can make a direct comparison with grounding theory, where it is claimed that joint actions are coordinated via dialog participants' common ground of information and beliefs.

Definition 3.1.1 doesn't tell us directly what actions rational players should take; all it does is give us the means to represent the structure of a game. We require three additional concepts in order to say what actions a player should take. First, we need to define a set of possible *player strategies*, which assign an action for each possible game situation the player might find herself in. Second, we need to define the *outcomes* of a game, where an outcome determines the payoff to each agent based on the strategy that each agent

it is important to not assign too much significance to particular numerical values of utility functions, as long as they preserve the essence of the underlying preference relation over outcomes.

pursued. Third, we need to define some type of *solution concept*, which tells us which strategies are good (in that they lead to desired outcomes), and which ones are not.

For a strategic game, a strategy is trivial to define: a strategy for player i is simply a one-shot choice of action, $a_i \in A_i$. Given a strategy for each player, an *outcome* is also trivial to define: it is an action profile $a \in \times_{j \in N} A_j$. What's left is to define is a solution concept that tells the players which strategies they should adopt. The typical solution concept for a strategic game is an *equilibrium*, and the most commonly used equilibrium concept is the *Nash equilibrium*:[3]

Definition 3.1.2. A *Nash equilibrium for a strategic game* $\langle N, A, (\succsim_i) \rangle$is a profile $a^* \in A$ of actions with the property that for every player $i \in N$,

$$(a^*_{-i}, a^*_i) \succsim (a^*_{-i}, a_i) \text{ for all } a_i \in A_i$$

For a strategic game, a Nash equilibrium is an action profile where no player has incentive to unilaterally deviate and choose another action. Each player's action is a *best response* to the actions of the other players. Therefore, for each player i, if she knows that players $j \in N \backslash \{i\}$ will play their actions in the equilibrium profile, i has no incentive to deviate. In general, it is possible for there to be multiple Nash equilibria for a given game, making it necessary for the players to coordinate on one of them in particular. Various refinements of the Nash equilibrium can help narrow down this set to a smaller number.

A strategic game models a one shot decision where all actions take place simultaneously. Each player selects an action from her set of available actions, and afterwards

[3]This definition uses standard notation where a_{-i} indicates the vector of actions in an action profile excluding the action of player i, while (a_{-i}, a_i) and (a_{-i}, a'_i) indicate two action profiles that differ only in the value of player i's action.

each player receives a payoff that corresponds to the resulting action profile. A one-shot process like this is obviously not well-suited for modeling communication, because a communicative act must be able to transmit relevant information *before* a subsequent, payoff relevant action is taken. Therefore, the basis for the game-theoretic model developed in the rest of this chapter will be an alternative representation known as an *extensive game*. Extensive games explicitly model the sequential structure of interaction; one player can take action before another does, and the action a player chooses can therefore influence subsequent decisions of other players. Informally, an extensive game is a tree (a *game tree*), with nonterminal nodes labelled by players, the branches labelled by actions, and terminal nodes labelled by payoffs. In this section, I focus on the simplest type of extensive game, where each player's move is fully observable to all other players. This type of game is called an *extensive game with perfect information.*[4]

Figure 3.1 depicts an extensive game with perfect information that has three players, and three stages where actions can be selected. Each interior node of the tree is labelled with one of the three players: 1, 2, or c. The branches of the tree that follow these nodes are labelled with actions. These represent the choices that the players have at that particular point in the game. For player 1, the action choices for the two nodes she is associated with are a and a'. For player 2, the action choices are b and b'. Player c is a little different: it can select either action l or r, but it does so according to a probability distribution. Player c is called "chance" (or "nature"), and its choices at any point in the tree are associated with numerical values that represent probabilities. The terminal

[4]Mathematically, strategic games and extensive games are arguably equivalent (see Myerson (1991) and Osborne and Rubinstein (1994)). However, showing this equivalence requires taking a more subtle approach to the nature of an action in a strategic game. See the cited references for more discussion of this issue.

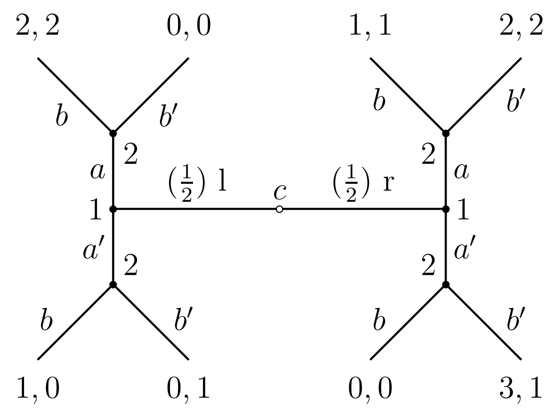

Figure 3.1. Extensive game with perfect information

nodes of the tree are labelled with tuples of numbers. Each tuple represents the payoffs that the players receive if the path of play leads to a given terminal node.[5]

Figure 3.2 shows two potential paths of play through the game tree in Figure 3.1. For the blue path, chance first chooses l. Since this is a game of perfect information, every other player observes this choice, and in particular player 1 knows that l was chosen. Player 1 at this point selects action a. Player 2 now has the opportunity to choose, after having observed both chance and player 1 take their moves. In this case, he decides on action b. This takes us to a terminal node, where the game concludes with player 1 and player 2 both receiving a payoff of 2. The green path represents an alternative play of the game, corresponding to the terminal history (r, a', b'), resulting in a payoff of 3 for player 1 and a payoff of 1 for player 2.

Formally, an extensive game with perfect information can be defined with the following components (Osborne and Rubinstein, 1994):

[5]Chance does not receive a payoff, as it is disinterested in the outcome of the game.

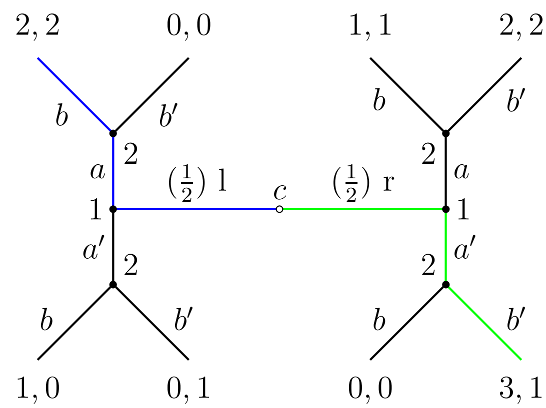

Figure 3.2. Two paths of play in an extensive game

Definition 3.1.3. An *extensive game with perfect information* is a tuple $\langle N, H, P, f_c, (\succsim_i)\rangle$, where:

- N is a finite set of *players*, indexed by i;
- H is a set of action sequences, such that:
  - The empty sequence $\emptyset$ is a member of H, and
  - If $(a^k)_{k=1,\ldots,K} \in H$, then so is $(a^k)_{k=1,\ldots,L}$, when $L < K$.

  Each $h \in H$ is a *history*. A history $(a^k)_{k=1,\ldots,K} \in H$ is *terminal* if there is no a^{K+1} such that $(a^k)_{k=1,\ldots,K+1} \in H$. The set of terminal histories is denoted Z. The set of actions after the nonterminal history $h \in H \backslash Z$ is denoted $A(h) = \{a : (h, a) \in H\}$;
- P is a function that assigns to each nonterminal history a member of $N \cup \{c\}$. P is the *player function*, $P(h)$ being the player who takes an action after the history h. If $P(h) = c$, then *chance* determines the action after history h;

- f_c is a function that associates with each $h \in H$ where $P(h) = c$ a probability measure $f_c(\cdot|h)$ on $A(h)$;
- $\succsim_i$ is a *preference relation* for player i on probability distributions over the set of terminal histories Z.

Unlike the one-shot model of games given by Definition 3.1.1, Definition 3.1.3 allows players to make multiple action choices over the course of a game. At any given point in a game, a player will have observed all actions of all players that have taken place earlier in its history, hence the player makes her choice with full knowledge of what has happened so far. Whereas a strategy in a strategic game is trivially defined as a one-shot action choice, the more complex structure of extensive games leads us to a define a strategy as function. This function assigns an action from the set $A(h)$ (where A is defined as in 3.1.3) to each $h \in H$ where $P(h) = i$. For extensive games, therefore, a strategy is a plan-like construct with conditional logic: for player i, *if* i reaches a particular point in the game tree, *then* player i will take a particular action. In the game shown in Figure 3.1, a strategy for player 1 selects an action from the set $A(h)$ where h is one of (l) or (r), while a strategy for player 2 selects an action for each of (l, a), (l, a'), (r, a), and (r, a'). The following strategies for players 1 and 2 are consistent with the two paths of play shown in Figure 3.2: $s_1((l)) = a$, $s_1((r)) = a'$, $s_2((l, a)) = s_2((r, a)) = b$, and $s_2((l, a')) = s_2((r, a')) = b'$.

Given a strategy profile such as this one, and given a set of probability distributions $f_c(\cdot|h)$ for each nonterminal history where $P(h) = c$, we can define the outcome $O(s)$ of a strategy profile $s = (s_i)_{i \in N}$ in an extensive game with perfect information as a probability distribution over terminal histories, where $O(s)(h)$ denotes the probability of terminal history $h \in Z$ given strategy profile s. The probability assigned to each terminal history $h = (a^k)_{k=1,\ldots,K}$ is equal to $\prod_{k=1}^{K} \hat{p}_s(a^k|(a^1,\ldots,a^{k-1}))$, where:

- $\hat{p}_s(a^k|(a^1,\ldots,a^{k-1})) = 1$ if $P((a^1,\ldots,a^{k-1})) = i$, $i \neq c$, and $s_i((a^1,\ldots,a^{k-1})) = a^k$, and
- $\hat{p}_s(a^k|(a^1,\ldots,a^{k-1})) = f_c(a^k|(a^1,\ldots,a^{k-1}))$ if $P((a^1,\ldots,a^{k-1})) = c$.

Intuitively, the probability of each terminal history is determined by whether or not each component action of the history is selected by the strategy s_i of the player i whose turn it is to select an action, and by the probability distribution determined by $f_c(\cdot|h)$ if it is chance's turn. Given definitions for strategies and outcomes, we now define a Nash equilibrium for an extensive game with perfect information (Osborne and Rubinstein, 1994):

Definition 3.1.4. A Nash equilibrium for an extensive game with perfect information is a strategy profile s^* such that for every player $i \in N$ we have

$$O(s^*_{-i}, s^*_i) \succsim_i O(s^*_{-i}, s_i) \text{ for every strategy } s_i \text{ of player } i$$

Definition 3.1.4 is similar to Definition 3.1.2, with the exception that action profiles are replaced with outcomes as they are defined for extensive games. The definition is stated in terms of preference relations over these outcomes. In the remainder of this thesis however, it will be more convenient to deal with utility functions rather than directly with preference relations. For preference relations that are modeled as probability distributions over the set of terminal histories, the standard way to do this is by performing an expected value calculation. That is, for every player $i \in N$ there is a utility function $u_i : Z \mapsto \mathbb{R}$, and $O(s) \succsim_i O(s')$ if and only if $\sum_{h \in Z} O(s)(h) u_i(h) \geq \sum_{h \in Z} O(s')(h) u_i(h)$.

Turning back to Figure 3.2, we can see that the strategy profile $s = (s_1, s_2)$ defined earlier for the two players is a Nash equilibrium. Given this strategy profile, we get $O(s)((l, a, b)) = O(s)((r, a', b')) = \frac{1}{2}$. Neither player 1 nor player 2 could choose a strategy that returns a higher expected value given the strategy of the other. Therefore, (s_1, s_2) is a Nash equilibrium. However, the game has more than one Nash equilibrium. Another equilibrium consists of $s' = (s'_1, s'_2)$, where s'_1 is defined as $s'_1((l)) = s'_1((r)) = a$, and s'_2 is defined as $s'_2((l, a)) = s'_2((r, a')) = b, s'_2((l, a')) = s'_2((r, a)) = b'$. This equilibrium is not as intuitive as the previous one: player 1's choice of $s'_1((r)) = a$ only makes sense because player 2 is threatening to play the non-optimal move $s'_2((r, a')) = b$, were player 1 to play a' after chance played r. But this type of "threat" from player 2 is not credible – were player 2 to actually arrive at this point in the game, rational self-interest dictates that he select b' because it gives him a higher payoff.

This type of non-intuitive equilibrium is eliminated by a standard refinement of the Nash equilibrium, known as the *subgame perfect equilibrium*. A *subgame* $\Gamma(h)$ of an extensive game $\Gamma = \langle N, H, P, f_c, (\succsim_i) \rangle$ is the game induced by taking just the sub-tree rooted at the non-terminal node following the history h (Osborne and Rubinstein, 1994):

Definition 3.1.5. A *subgame perfect equilibrium* of an extensive game with perfect information $\Gamma = \langle N, H, P, f_c, (\succsim_i) \rangle$ is a strategy profile s^* in Γ such that for any history h the strategy profile $s^*|_h$ is a Nash equilibrium of the subgame $\Gamma(h)$. ($s^*|_h$ is the strategy profile for $\Gamma(h)$ induced by the strategy profile s^* of Γ in which $\Gamma(h)$ is embedded).

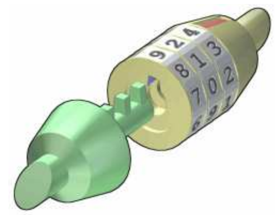

Figure 3.3. Combination lock

This refinement of the Nash equilibrium eliminates non-intuitive equilibria like s' for the game in Figure 3.1. For a subgame perfect equilibrium of this game, player 2 can no longer make the non-credible threat of playing b after the history (r, a'), because this is not a Nash equilibrium of the subgame rooted at this history. In Chapter 4, the subgame perfect equilibrium will form the basis of a related equilibrium concept that will be applied to the game-theoretic model of grounding proposed in that chapter.

Definition 3.1.3 provides us the means to model the sequential structure of a game, and the means to model non-deterministic play (via chance moves). However, we still lack an adequate model of communication. To see this, let's consider how we would use this definition to model a simple referential communication task, which I call the *lock task*. In this task, two agents collaborate to open a combination lock, such as the one shown in Figure 3.3. One agent, the *helper*, knows the combination, while the other agent, the *worker*, does not. On the other hand, the worker has the physical ability to rotate the disks on the lock, while the helper does not. Imagine that the worker is standing next to the combination lock, while the helper is in a remote location communicating with the worker by phone. The helper has a sheet of paper with the correct combination written on it, and her job is to communicate this information to the worker. The worker's job is to use this information to open the lock.

This simple task presents us with two asymmetries between the participants: (1) an *informational* asymmetry, which requires communication in order to solve the task in a reasonable amount of time, and (2) an asymmetry in *ability to affect physical state*, because the worker can take actions that change the lock settings while the helper cannot. These two types of asymmetry are typical of the referential communication task experiments described in Section 2.2. For the basic version of the lock task, I assume that the goals of the helper and the worker are perfectly aligned. Imagine that opening the lock provides access to some treasure, the value of which is to be equally divided between them. The helper and the worker therefore have every incentive to cooperate on opening the lock in order to gain the reward. This perfect alignment of interests is also typical of the experiments described in Section 2.2. The basic version of the lock task therefore serves as an abstract model of a referential communication task.

Figure 3.4 shows an attempt to model the lock task using Definition 3.1.3. It depicts a game with three players: h, w, and c. These represent the helper, the worker, and chance, respectively. The first move belongs to chance, selecting the solution to the combination lock. In this scenario, the lock has only one rotating disk, and the disk has only two possible values, either d_1 or d_2. After chance selects the solution to the lock (with probability $\frac{1}{2}$ for either possible value), the helper is informed of chance's selection, and then has the option to choose either action m_1 or m_2. These represent (respectively) the messages "the solution is 1" and "the solution is 2". After choosing a message to send, the worker observes this selection and can then choose either action s_1 or s_2. These

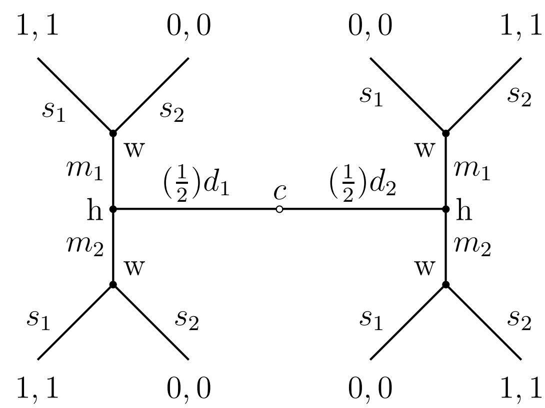

Figure 3.4. Lock task as an extensive game with perfect information

represent (respectively) “setting the lock disk to 1” or “setting the lock disk to 2”. After setting the lock disk, the game is over, and each player receives the payoffs indicated at the terminal nodes of the game tree, depending on their choices and the selection of chance.

The obvious problem with this attempt to model communication is that the message that the helper sends is irrelevant to the play of the game. It is a game of perfect information, so the worker knows everything that the helper knows, and in particular, the worker knows which lock solution that chance selected. In order to model communication, therefore, we need to augment Definition 3.1.3 in order to allow for the possibility that players can have private information. Referring to the game in Figure 3.4, the initial selection of chance should be revealed to the helper, but not to the worker. The message that the helper selects would then have significance to the worker, because it has the potential to update the beliefs of the worker in a way that is relevant to the payoff of the game. Extending the extensive model with the ability to model private information, and therefore communication, is the topic of the next section.

3.2. Signaling Games

In order to extend an extensive game with perfect information to model players with private information, we need to add a mechanism for representing uncertainty. Specifically, we need a model of the uncertainty of a player about payoff relevant information that another player might have. This model of uncertainty is provided by an *extensive game with imperfect information*, which is a straightforward generalization of extensive games with perfect information. The formal definition includes the following components (Osborne and Rubinstein, 1994):

Definition 3.2.1. An *extensive game with imperfect information* is a tuple $\langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i)\rangle$, where:

- N, H, P, f_c and $(\succsim_i)_{i \in N}$ are the same as in Definition 3.1.3;
- For each player $i \in N$, there is a partition $\mathcal{I}_i$ of $\{h \in H : P(h) = i\}$ with the property that $A(h) = A(h')$ whenever h and h' are in the same member of the partition. $\mathcal{I}_i$ is the *information partition* of player i, and a set $I_i \in \mathcal{I}_i$ is an *information set* of player i.

Definition 3.2.1 introduces a set of *information partitions* $\mathcal{I}_i$, one for each $i \in N$. An information set $I_i \in \mathcal{I}_i$ can be thought of as the set of histories that are compatible with what player i knows to be true at that point of the game. One of these histories must be the true one, but player i doesn't know which one, unless the set contains only a single history.

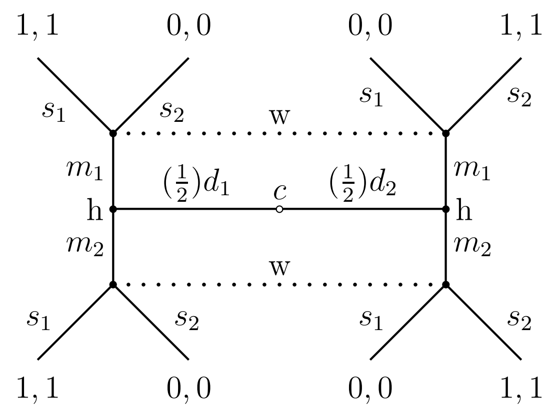

Figure 3.5. Lock task as a signaling game

Figure 3.5 presents the two-digit version of the lock task as a type of extensive game with imperfect information. The only visual difference between this figure and Figure 3.4 is the addition of two dashed lines between two sets of nodes: the pair of nodes after each branch where the helper takes action m_1, and the pair of nodes after the helper takes action m_2. Each of these dashed lines indicate a pair of nodes in the game tree that the worker cannot distinguish between. These are the worker's information sets, $I_{m_1} = \{(d_1, m_1), (d_2, m_1)\}$ and $I_{m_2} = \{(d_1, m_2), (d_2, m_2)\}$. They are to be interpreted as follows: the worker can observe the actions of the helper, and therefore knows whether she played m_1 or m_2, but the worker does not know the original action selected by "chance", either d_1 or d_2. Hence, the worker's two information sets each include two histories: after observing message m_1, the worker cannot distinguish between (d_1, m_1) and (d_2, m_1), and after observing message m_2, the worker cannot distinguish between (d_1, m_2) and (d_2, m_2).

On the other hand, the helper's information sets are singletons: $I_{d_1} = \{(d_1)\}$ and $I_{d_2} = \{(d_2)\}$. The helper therefore knows with certainty the action of chance, and hence the correct solution to the lock. This information is her private, payoff relevant information. The helper's choice of action – the message that she sends – can now play an important role in the game, by giving the worker evidence about the true state of the world. The helper knows the solution to the lock, while the worker does not. The worker can observe messages from the helper, and use this information to decide which action to take.

The game in Figure 3.5 represents an approach to modeling communication that was first proposed in Lewis (1969). Lewis, in an effort to model linguistic convention, defined a special type of two player game called a *signaling game*. Signaling games have subsequently been used in formal approaches to linguistic pragmatics (Benz et al., 2005; Stalnaker, 2005), theoretical biology (Maynard Smith, 1982) and economics (Spence, 1973). Formally, signaling games are members of a proper subset of extensive game with imperfect information, called *Bayesian extensive games with observable actions*. In this type of game, every player observes with full certainty the actions of every other player, and the only uncertainty in the game is about an initial move of chance that distributes private payoff relevant information among the players (Osborne and Rubinstein, 1994). In the literature on game theory, this private information is called the player's *type* (Harsanyi, 1967). Each player knows their own type, but only has access to probability distributions over the other players' types.

A signaling game is defined as a two-player Bayesian extensive game with observable actions in which chance selects a game to be played according to a commonly known distribution, player 1 is informed of that choice and chooses an action, and player 2 then chooses an action without knowing chance's choice, but knowing player 1's choice (Shoham and Leyton-Brown, 2009). Because a signaling game has such a restricted form,

	w_1	w_2	w_3	w_4
h_1	$1, 1$	$0, 0$	$1/2, 1/2$	$1/2, 1/2$
h_2	$0, 0$	$1, 1$	$1/2, 1/2$	$1/2, 1/2$
h_3	$1/2, 1/2$	$1/2, 1/2$	$1/2, 1/2$	$1/2, 1/2$
h_4	$1/2, 1/2$	$1/2, 1/2$	$1/2, 1/2$	$1/2, 1/2$

Figure 3.6. Nash equilibria for a signaling game

it is possible to define a strategy for a signaling game in a very specific way. Namely, we can view the helper's strategy as a function from the helper's types to her messages, and the worker's strategy as a function from messages to actions (see Benz et al. (2005)). For the game in Figure 3.5, there are four possible strategies of this nature for the helper, and four for the worker:

$$h_1 : \{\langle d_1, m_1\rangle, \langle d_2, m_2\rangle\} \qquad w_1 : \{\langle m_1, s_1\rangle, \langle m_2, s_2\rangle\} \tag{3.1}$$

$$h_2 : \{\langle d_1, m_2\rangle, \langle d_2, m_1\rangle\} \qquad w_2 : \{\langle m_1, s_2\rangle, \langle m_2, s_1\rangle\}$$

$$h_3 : \{\langle d_1, m_1\rangle, \langle d_2, m_1\rangle\} \qquad w_3 : \{\langle m_1, s_1\rangle, \langle m_2, s_1\rangle\}$$

$$h_4 : \{\langle d_1, m_2\rangle, \langle d_2, m_2\rangle\} \qquad w_4 : \{\langle m_1, s_2\rangle, \langle m_2, s_2\rangle\}$$

Given the set of individual player strategies defined above, there are 16 possible strategy profiles. Figure 3.6 represents these profiles in matrix form, with the helper's strategies on the left, and the worker's strategies on the top. The entries in the matrix cells represent the expected payoffs to the players, given that they play the indicated strategies. The six cells that are highlighted represent the Nash equilibria of the game.

Not all of these equilibria are created equal. The two equilibria (h_1, w_1) and (h_2, w_2) each have a payoff of 1, while the other four have payoffs of $\frac{1}{2}$. The first two equilibria establish a unique mapping from the helper's type to a message, enabling the worker to use the message received in order to inform his decision on which action to take. These are called *separating equilibria*: the messages of the informed player provide the uninformed player with enough information to adequately distinguish between payoff relevant courses of action. For the other four equilibria, successful communication does not take place. The messages bear no useful information, and therefore the best the worker can do is guess at which action to take, leading to an expected payoff of only $\frac{1}{2}$. These are called *babbling equilibria*.

Much of the literature on signaling games is concerned with when meaningful communication occurs, i.e., what conditions ensure that there are separating equilibria, and when players have the incentive to speak the truth (Crawford and Sobel, 1982; Farrell and Rabin, 1996; Stalnaker, 2005). However, this topic is not the focus of this thesis. In a referential communication task, I assume that players will tell the truth, that their messages have semantic content, and that the players have common interests. The goal is to model grounding behavior, and to simplify matters I will assume that cooperative communication is a given. Therefore, I will only take into account games where truthful communication occurs. This can be achieved by restricting the player's action functions in such a way that babbling equilibria are eliminated from consideration.

Signaling games provide a good first approximation of a model for communication during a referential communication task. However, they are still insufficient as a model of grounding for such tasks. This is because they are defined as Bayesian games with observable actions. Therefore, each player's moves are fully observable to the other players. Given this situation, establishing common ground is simple – the helper just has to decide which message to send, and the worker will receive the intended message. There may be uncertainty associated with the message itself (e.g., ambiguity in its meaning), but there will be no uncertainty with respect to which message was sent. Grounding behaviors, like the ones described in Chapter 2, are simply unnecessary in this case. In order to model grounding, therefore, we must relax the restriction that uncertainty exists only with respect to a player's type. We must also allow for the possibility of uncertainty about which action a player has executed. Extending signaling games in this manner is the topic of the next chapter.

CHAPTER 4

Communication with Uncertainty

While signaling games successfully model communication, they do not model the possibility that communicative acts can fail to have their intended effect on listeners. Without this possibility, we have no way to model the types of grounding behaviors described in Chapter 2 (see the discussion in Section 2.1.) Therefore, we need to extend signaling games to allow for the possibility of communicative error.[1] In this chapter, I accomplish this by generalizing signaling games to *signaling games with partially observable actions.* This type of game can model the uncertainty of players about moves that have occurred in the past, such as the uncertainty of a dialog participant about which message another participant actually sent. When this kind of uncertainty is included in the model, conversational grounding behaviors are justified as actions that reduce uncertainty in the conversational common ground.

4.1. Partially Observable Actions

Definition 3.2.1 already provides the ingredients for modeling uncertainty of the type that allows us to model grounding behaviors. In a general extensive game[2], a player's actions can be "unobserved" by other players. Additionally, chance can play a role beyond

[1]By using the term "error", I don't just necessarily mean speech recognition problems due to a noisy channel. I'm using the term to refer to any potential problem that the recipient of a message might have recovering its intended meaning. Problems might arise due to any combination of a noisy channel, lexical ambiguities, listener distraction, etc.

[2]From now on, the term "extensive game" refers to an extensive game with imperfect information, as in Definition 3.2.1.

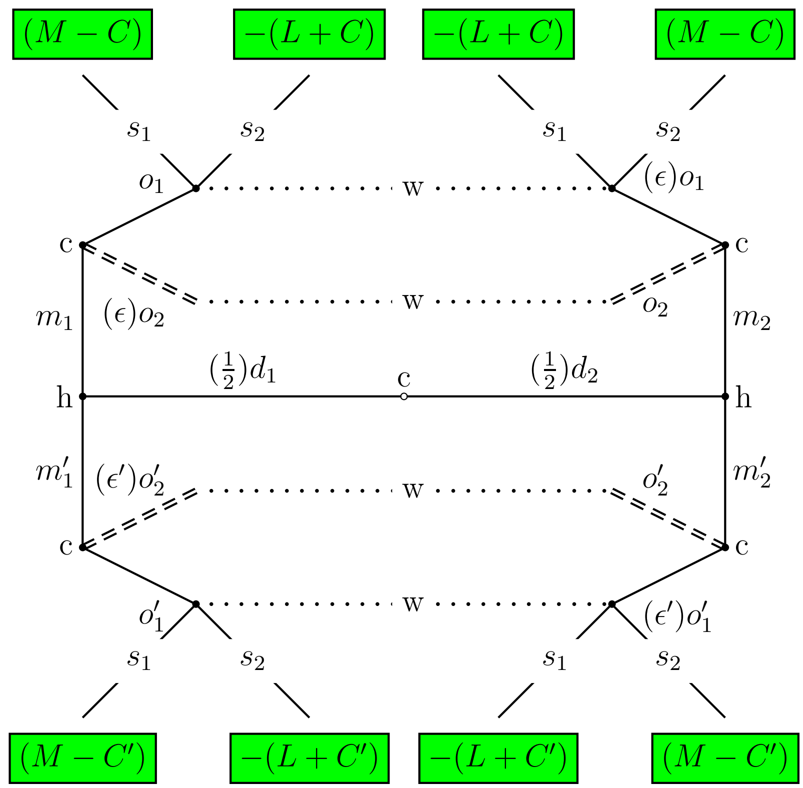

Figure 4.1. Lock task as a signaling game with partially observable actions.

the initial distribution of payoff relevant private information to players. Figure 4.1, which extends the signaling game in Figure 3.5 demonstrates both of these features.

The game in Figure 4.1 is to be interpreted as follows. In similar fashion to the signaling game depicted in Figure 3.5, the initial move is made by chance, who selects lock solution d_1 or d_2 with probability $\frac{1}{2}$. The next move belongs to the helper. Her information sets at this point distinguish between the two histories (d_1) and (d_2), hence she knows the solution to the lock. If history (d_1) occurs, then the helper has two action choices, m_1 or m_1'. If history (d_2) occurs, then the helper has two other action choices, m_2 or m_2'. Both m_1 and m_1' are to be interpreted as alternative ways of communicating the same message, namely "the lock solution is 1". Likewise, m_2 and m_2' are to be interpreted as alternative ways to communicate "the lock solution is 2". This game tree therefore provides an abstract model of alternative ways to communicate the same information.[3]

Looking forward, we can see that choosing m_1 or m_2 leads to payoffs $(M - C)$ or $-(L+C)$, while choosing m_1' or m_2' leads to payoffs $(M - C')$ or $-(L+C')$. The value M represents the payoff for both players when the worker chooses the "correct" action (i.e, setting the lock to 1 when the solution is 1, and to 2 when the solution is 2). Conversely, the value L represents the penalty for both players when the worker chooses the incorrect action, such as setting the lock to 2 when the solution is 1. The values C and C' are the *costs* associated with particular messages, where C is the cost of m_1 and m_2, and C' is the cost of m_1' and m_2'. These costs will play an important role later in this chapter in tying the game-theoretic model back to conversational grounding theory's grounding criterion and principle of least collaborative effort.

After the helper makes her choice, the next move is again made by chance. Chance selects an *observation* of the message action selected by the helper.[4] Chance selects either observation o_1 or o_2 after m_1 and m_2 respectively, and observation o_1' or o_2' after m_1' and m_2'. If the helper sends message m_1, chance selects o_1 with probability $(1 - \epsilon)$, and o_2 with probability ϵ. If the helper sends message m_2, chance selects o_2 with probability $(1-\epsilon)$, and o_1 with probability ϵ. After messages m_1' and m_2', the observation probabilities are $(1 - \epsilon')$ for o_1' and o_2', respectively. Chance's probability distributions

[3]The structure of this game eliminates the possibility of babbling equilibria. However, if desired it would be easy to re-introduce this possibility by expanding the helper's set of alternatives to m_1, m_1', m_2, and m_2' at each of her choice points.

[4]The double dashed lines at this point indicate that the remainder of the game tree has been elided, in order to reduce the visual complexity of the diagram. The worker's choices after each of these elisions is either s_1 or s_2, and the payoffs are identical to their sister (unelided) nodes in the tree.

over these observations provides the worker with exogenously determined probabilistic information about which message the worker sent. Since by assumption these probability distributions are commonly known between the helper and the worker, they represent both the uncertainty of the worker *and* the helper that the helper's message was properly understood by the worker.

Both ϵ and ϵ' are to be interpreted as the likelihood of communicative error, given a helper message. That is, it represents the possibility that the message that the helper sends is not understood properly for some reason, whether due to a noisy channel, listener distraction, or lexical ambiguity. It is the level of uncertainty that both helper and worker have about the success of a communicative act. For this game (and for reasons that will become clear later) we require that $0 \leq \epsilon \leq \epsilon' \leq \frac{1}{2}$. The boundary cases for these values are $\epsilon = \epsilon' = 0$, and $\epsilon = \epsilon' = \frac{1}{2}$. In the first case, where both $\epsilon = \epsilon' = 0$, the game in Figure 4.1 reduces to a garden variety signaling game, such as the one shown in Figure 3.5. With no possibility for error, the helper's messages become completely observable to the worker, and the only remaining uncertainty is the worker's uncertainty about the move of chance that begins the game. In the second case, where both $\epsilon = \epsilon' = \frac{1}{2}$, the helper's messages bear no useful information for the worker, since they are completely ambiguous. Useful communication can therefore occur only between these two extremes.

After receiving a message observation from chance, the worker then takes his action, selecting either s_1 or s_2. The payoffs to the players are then determined by which terminal history was followed in the game tree. As mentioned above, and unlike the signaling game shown in Figure 3.5, the payoffs in this game take into account the cost of a message. This cost is deducted from the payoff received at the end of the game. Furthermore, since in this game both helper and worker receive the same payoff (i.e, this is a game of pure coordination), the payoffs of the helper and worker are both deducted by the same amount. The specific amount depends on the message that is sent: m_1 and m_2 each cost a positive amount C, while m_1' and m_2' each cost a positive amount C'. For this game, I stipulate that $C \leq C'$. With the above restriction that $\epsilon \leq \epsilon'$, this leads to the following interpretation of the game in Figure 4.1: m_1' (respectively m_2') is at least as "precise" as m_1 (respectively m_2), but this extra precision comes at a potential price, since C' is at least as big as C.

I refer to games of the type shown in Figure 4.1 as *signaling games with partially observable actions*. They are straightforward generalizations of standard signaling games, and can be defined as follows:

Definition 4.1.1. A *signaling game with partially observable actions* is a two-player extensive game $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i), (\succsim_i)\rangle$ where an initial move of chance selects a game to be played according to a commonly known distribution. Player 1 is informed of that choice and chooses an action, at which point chance chooses an *observation* for player 2 according to a commonly known distribution. Player 2 then chooses an action knowing only the observation he received, and the probability distributions of the chance moves.

If chance always selects a single observation with probability 1, then a signaling game with partially observable actions reduces to a standard signaling game. Therefore Definition 4.1.1 is a simple generalization of standard signaling games. It could easily be extended in several ways, such as permitting more than two players, and by allowing the players to communicate with a sequence of messages before the final payoff relevant action is taken. I have refrained from doing so here in order to highlight the close relationship of signaling games with partially observable actions to standard signaling games. I will show later in this chapter that Definition 4.1.1 is sufficient to model the core properties of grounding theory described in Chapter 2.

Now that I have defined signaling games with partially observable actions, three additional auxiliary definitions are required in order to come up with solutions for them: (1) a *player strategy*, (2) an *outcome* to a game given a strategy profile, and (3) an *equilibrium* concept to select good strategies from among the set of possible strategies. For signaling games with partially observable actions, I will simply apply standard tools developed for general extensive games with imperfect information. In what follows, I will closely follow Osborne and Rubinstein (1994) in particular.

A strategy for an extensive game is defined as a function from information sets to probability distributions over actions, and is called a *behavioral strategy.* Let $A(I_i)$ be the set of actions available to player i at information set I_i. Then a behavioral strategy of player $i \in N$ is a collection $\beta(I_i)_{I_i \in \mathcal{I}_i}$ of probability measures, where $\beta(I_i)$ is a probability measure over $A(I_i)$. If $\beta(I_i)$ assigns probability 1 to an action in $A(I_i)$ for each $\beta(I_i)_{I_i \in \mathcal{I}_i}$, then player i has a *pure strategy.* A pure strategy can alternatively be denoted as a function from information sets to actions, rather than to probability distributions over actions.

For extensive games with imperfect information, a player must decide how to play without always knowing what history has actually occurred. However, a player can use his knowledge of the structure of the game, and his beliefs about the strategies of the other players, in order to come up with a probability distribution over the set of histories in his current information set. This fact is important, because a player's beliefs about the probabilities of histories should influence which decision he will make. For example, if the worker in Figure 4.1 receives observation o_1 from chance, then he can calculate the likelihood that the helper sent message m_1 versus message m_2, and select his response accordingly. Figuring out which message is more likely to have occurred should influence which action the worker will take.

As this discussion makes clear, optimal play in a game with imperfect information demands that players calculate the probabilities of histories that might have occurred. Therefore, a solution concept has to include these probabilities as a component. This leads us to the concept of an *assessment* (Osborne and Rubinstein, 1994): a pair (β, μ), where β is a profile of behavioral strategies, and μ is a function that assigns to every information set a probability measure on the set of histories in the information set. The function μ is called a *belief system*, and $\mu(I)(h)$ is the probability that player $P(I)$ assigns to the history $h \in I$, conditioned on I being reached. Given an assessment (β, μ), an *outcome* is defined as follows (Osborne and Rubinstein, 1994):

Definition 4.1.2. The *outcome* $O(\beta, \mu|I)$ of (β, μ) conditional on I is the probability distribution over terminal histories determined by β and μ conditioned on I being reached. If $h^* = (a^1, \ldots, a^K)$ is a terminal history, then:

- $O(\beta, \mu|I)(h^*) = 0$ if there is no sub-history of h^* in I
- $O(\beta, \mu|I)(h^*) = \mu(I)(h) \times \prod_{k=L}^{K-1} \beta_{P((a^1,\ldots,a^k))}(a^1, \ldots, a^k)(a^{k+1})$ if the subhistory $h = (a^1, \ldots, a^L)$ of h^* is in I, where $L < K$.

Once player i reaches an information set I_i, she can use μ to calculate the probability of any history $h \in I_i$, conditioned on I_i having been reached, and then use her knowledge of the players' strategies β to figure out the outcome moving forward, i.e., a probability distribution over terminal histories. Given an outcome, it is straightforward for player i to calculate the expected utility for player i of (β, μ) given I_i. $EU_i(\beta, \mu|I_i)$ is simply the sum of utilities for each terminal history, weighted by its outcome probability.

Finally, we now have the ingredients to give a definition for an equilibrium concept that applies to extensive games with imperfect information. This equilibrium concept is a generalization of Definition 3.1.5, which applied to games of perfect information (Osborne and Rubinstein, 1994):

Definition 4.1.3. A *sequential equilibrium* of an extensive game with imperfect information is an assessment (β, μ) that is both *sequentially rational* and *consistent*, where:

- (β, μ) is *sequentially rational* if for every player $i \in N$ and every information set $I_i \in \mathcal{I}_i$ we have

$$O(\beta, \mu|I_i) \succsim_i O((\beta_{-i}, \beta_i'), \mu|I_i) \text{ for every strategy } \beta_i' \text{ of player } i$$

- (β, μ) is *consistent* if there is a sequence $((\beta^n, \mu^n))_{n=1}^{\infty}$ of assessments that converges to (β, μ) and has the property that each belief system μ^n is derived from β^n using Bayes' rule.

Definition 4.1.3 imposes two conditions on an equilibrium profile. The first condition, *sequential rationality*, is a straightforward extension of the notion of a subgame perfect equilibrium (Definition 3.1.5). It basically says that each player will choose an action that maximizes the expected utility of an outcome, at every point in the game where the player chooses an action. The second condition, *consistency* says that the probability of histories in such an information set must be derived from behavioral strategies using Bayesian updating. The consistency condition is stated in terms of limits because we need a way to calculate the probabilities of histories in information sets that have zero probability of being reached, given a particular strategy profile.

4.1.1. Solving the Lock Task

With these definitions in order, I now apply them to the game in Figure 4.1 in order to identify an equilibrium strategy profile. For the sake of concreteness, I assume that $(M+L) > 0$, that $0 < \epsilon' < \epsilon < \frac{1}{2}$, and that $C < C'$. I first define the information partitions $\mathcal{I}_h$ and $\mathcal{I}_w$ of the helper and worker, and the range of possible behavioral strategies β_h and β_w induced by these information partitions. The helper's information partition $\mathcal{I}_h$ contains two information sets, while the information partition $\mathcal{I}_w$ of the worker contains four information sets:

$$
\begin{aligned}
I_{h_1} &= \{(d_1)\} & \qquad I_{w_1} &= \{(d_1, m_1, o_1), (d_2, m_2, o_1)\} \\
I_{h_2} &= \{(d_2)\} & \qquad I_{w_2} &= \{(d_1, m_1, o_2), (d_2, m_2, o_2)\} \\
& & \qquad I_{w_3} &= \{(d_1, m_1', o_1'), (d_2, m_2', o_1')\} \\
& & \qquad I_{w_4} &= \{(d_1, m_1', o_2'), (d_2, m_2', o_2')\}
\end{aligned}
\tag{4.1}
$$

Given her information sets, the helper's behavioral strategy is some function from I_{h_1} into a probability distribution over over m_1 and m_1' and from I_{h_2} into a probability distribution over over m_2 and m_2'. Since there are only two action choices for each information set, we can represent the probability distribution over the action choices as follows, where $\delta_1, \delta_2 \in [0, 1]$:

$$\beta_1(I_{h_1})(m_1) = \delta_1 \qquad \beta_1(I_{h_1})(m_1') = 1 - \delta_1 \tag{4.2}$$
$$\beta_1(I_{h_2})(m_2) = \delta_2 \qquad \beta_1(I_{h_2})(m_2') = 1 - \delta_2$$

The worker's behavioral strategy is a function from his four information sets into probability distributions over the actions s_1 and s_2. Again, since there are only two action choices for each information set, we can represent the probability distributions simply as follows, where $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in [0, 1]$:

$$\beta_2(I_{w_1})(s_1) = \alpha_1 \qquad \beta_2(I_{w_1})(s_2) = 1 - \alpha_1 \tag{4.3}$$
$$\beta_2(I_{w_2})(s_2) = \alpha_2 \qquad \beta_2(I_{w_2})(s_1) = 1 - \alpha_2$$
$$\beta_2(I_{w_3})(s_1) = \alpha_3 \qquad \beta_2(I_{w_3})(s_2) = 1 - \alpha_3$$
$$\beta_2(I_{w_4})(s_2) = \alpha_4 \qquad \beta_2(I_{w_4})(s_1) = 1 - \alpha_4$$

The issue now is to identify values for δ_1, δ_2, α_1, α_2, α_3, α_4, and a belief system μ, such that the constraints of a valid sequential equilibrium are satisfied. In what follows, I make the further assumption that $\delta_1 = \delta_2$, and refer to each value as simply δ.[5] Temporarily,

[5]This is justified in that both cases are symmetrical and shouldn't be different in principle. It simplifies matters by eliminating some uninteresting equilibria that are accidental properties of the particular game in Figure 4.1.

I also make the assumption that $0 < \delta < 1$, which implies that each of the worker's information sets has a positive probability of being reached. Given these assumptions, the belief system μ can be derived as follows (taking as example information set I_{w_1}):

$$\begin{aligned}
\mu(I_{w_1})((d_1, m_1, o_1)) &= \frac{\frac{1}{2}\delta(1-\epsilon)}{\frac{1}{2}\delta(1-\epsilon) + \frac{1}{2}\delta\epsilon} \\
&= \frac{1-\epsilon}{1-\epsilon+\epsilon} \\
&= 1-\epsilon \\
\mu(I_{w_1})((d_2, m_2, o_1)) &= \frac{\frac{1}{2}\delta(\epsilon)}{\frac{1}{2}\delta(1-\epsilon) + \frac{1}{2}\delta\epsilon} \\
&= \frac{\epsilon}{1-\epsilon+\epsilon} \\
&= \epsilon
\end{aligned}$$

Similar calculations generate the rest of the belief system μ:

$$\begin{aligned}
&\mu(I_{w_2})((d_2, m_2, o_2)) = 1-\epsilon && \mu(I_{w_2})((d_1, m_1, o_2)) = \epsilon \\
&\mu(I_{w_3})((d_1, m_1', o_1')) = 1-\epsilon' && \mu(I_{w_3})((d_2, m_2', o_1')) = \epsilon' \\
&\mu(I_{w_4})((d_2, m_2', o_2')) = 1-\epsilon' && \mu(I_{w_4})((d_1, m_1', o_2')) = \epsilon'
\end{aligned} \tag{4.4}$$

Given this belief system, we can now calculate the expected utility of β and μ (again taking as example information set I_{w_1}):

$$\begin{aligned}
EU(\beta, \mu | I_{w_1}) &= \sum_{z \in Z} O(\beta, \mu | I_{w_1})(z) \times u(z) \\
&= \sum_{h \in I_{w_1}} \mu(I_{w_1})(h) \times \sum_{s \in \{s_1, s_2\}} \beta_2(I_{w_1})(s) \times u((h, s)) \\
&= (1-\epsilon)(\alpha_1)(M-C) - (1-\epsilon)(1-\alpha_1)(L+C) \\
&\quad +(\epsilon)(1-\alpha_1)(M-C) - (\epsilon)(\alpha_1)(L+C) \\
&= \alpha_1 (M - 2\epsilon M + L - 2\epsilon L) + \epsilon (M+L) - (L+C)
\end{aligned}$$

Note that $\alpha_1 (M - 2\epsilon M + L - 2\epsilon L) > 0$ for $M, L > 0$ and $\epsilon < \frac{1}{2}$. Therefore, the worker's utility maximizing strategy is to set $\alpha_1 = 1$. By similar arguments, we can show that the worker's best strategy β_2 is to set $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 1$. Thus the worker's best strategy is a *pure strategy*, and given this pure strategy the expected utility values for the worker reduce to:

$$EU(\beta, \mu | I_{w_1}) = (1-\epsilon)M - \epsilon L - C$$

$$EU(\beta, \mu | I_{w_2}) = (1-\epsilon)M - \epsilon L - C$$

$$EU(\beta, \mu | I_{w_3}) = (1-\epsilon')M - \epsilon' L - C'$$

$$EU(\beta, \mu | I_{w_4}) = (1-\epsilon')M - \epsilon' L - C'$$

Now we return to the helper's strategy β_1. Since we now know the strategy of the worker, we can use this information to determine what the helper should do, given information set I_{h_1} or I_{h_2}. For information set I_{h_1} there are three possibilities to consider: (1) $EU(\beta, \mu | I_{w_1}) = EU(\beta, \mu | I_{w_3})$, (2) $EU(\beta, \mu | I_{w_1}) < EU(\beta, \mu | I_{w_3})$, and (3) $EU(\beta, \mu | I_{w_1}) > EU(\beta, \mu | I_{w_3})$. The situation for information set I_{h_2} is the same, except the comparisons are between $EU(\beta, \mu | I_{w_2})$ and $EU(\beta, \mu | I_{w_4})$. For case (1), the particular value of δ is irrelevant, since the helper doesn't prefer one outcome to the other. For case (2), it is clear that the helper should choose $\delta = 0$, and for case (3) the helper should choose $\delta = 1$.

Given these arguments, we now see that the helper also has a pure strategy for either case (2) or case (3), and the value of δ (either 0 or 1) depends on the particular values of C, C', ϵ, and ϵ'. In fact, we can state a precise relationship among these values. For example, let's assume that case (2) from the previous paragraph holds, so that $EU(\beta, \mu | I_{w_1}) < EU(\beta, \mu | I_{w_3})$. This implies the following:

$$\begin{aligned}
EU(\beta, \mu | I_{w_1}) &< EU(\beta, \mu | I_{w_3}) \\
(1-\epsilon)M - \epsilon L - C &< (1-\epsilon')M - \epsilon' L - C' \\
C' - C &< (1-\epsilon')M - \epsilon' L - (1-\epsilon)M + \epsilon L \\
C' - C &< \epsilon(M+L) - \epsilon'(M+L) \\
C' - C &< (\epsilon - \epsilon')(M+L) \\
\frac{C' - C}{M+L} &< (\epsilon - \epsilon')
\end{aligned}$$

If we denote $C' - C$ as $\Delta(C)$ and $\epsilon - \epsilon'$ as $\Delta(\epsilon)$, we get the following constraint on the relationship among C, C', ϵ, and ϵ':

$$\frac{\Delta(C)}{M + L} < \Delta(\epsilon) \tag{4.5}$$

Inequality 4.5 shows us a simple way to evaluate which of two messages the helper should use, where one message is more "precise" than another, and yet this precision comes at a greater cost. We simply take the difference in the costs of the two messages (normalized with respect to $M + L$), and compare it to the gain in precision that results. If the normalized difference in the costs is less than the gain in precision, the helper should choose the higher precision message. If, on the other hand, we replace $<$ with $>$ in this inequality, the helper should choose the cheaper (and less precise) message. This relationship between cost and error is the core result of this section, because it gives us a simple way to characterize the trade-off of cost and precision that is the essence of both the grounding criterion and the principle of least collaborative effort (as shown in Section 4.2).

To conclude this section, here is a summary of the sequential equilibrium solution to the game in Figure 4.1. The solution relies on assumptions that $M + L > 0$, that $C < C'$, and that $0 \leq \epsilon' < \epsilon < \frac{1}{2}$. For specificity, I will deal with the case where Inequality 4.5 holds, so that the higher precision (and higher cost) messages m_1' and m_2' are preferred to their lower precision (and lower cost) counterparts m_1 and m_2. Given these assumptions, the following belief system and behavioral strategies form a sequential equilibrium (Definition 4.1.3):

$$\text{(4.6)} \qquad \begin{aligned} \mu(I_{w_1})((d_1, m_1, o_1)) &= 1 - \epsilon & \qquad \mu(I_{w_1})((d_2, m_2, o_1)) &= \epsilon \\ \mu(I_{w_2})((d_2, m_2, o_2)) &= 1 - \epsilon & \mu(I_{w_2})((d_1, m_1, o_2)) &= \epsilon \\ \mu(I_{w_3})((d_1, m_1', o_1')) &= 1 - \epsilon' & \mu(I_{w_3})((d_2, m_2', o_1')) &= \epsilon' \\ \mu(I_{w_4})((d_2, m_2', o_2')) &= 1 - \epsilon' & \mu(I_{w_4})((d_1, m_1', o_2')) &= \epsilon' \end{aligned}$$

$$\text{(4.7)} \qquad \begin{aligned} \beta_1(I_{h_1})(m_1') &= 1 & \qquad \beta_2(I_{w_1})(s_1) &= 1 \\ \beta_1(I_{h_2})(m_2') &= 1 & \beta_2(I_{w_2})(s_2) &= 1 \\ & & \beta_2(I_{w_3})(s_1) &= 1 \\ & & \beta_2(I_{w_4})(s_2) &= 1 \end{aligned}$$

There is one technical detail to take care of. Given the strategy profile (β_1, β_2) shown in 4.7, there is actually a probability of 0 that the worker will ever find himself in information sets I_{w_1} or I_{w_3}, since the helper will never employ actions m_1 or m_2. I finessed this problem at the beginning of this section by assuming that the helper's value for δ was strictly greater than 0 and strictly less than 1. However, as demonstrated above, for the equilibrium play the value of δ *is* either equal to 0 or to 1. Therefore, in order to satisfy the consistency requirement on the belief system, we need to identify a sequence $((\beta^n, \mu^n))_{n=1}^{\infty}$ of assessments that converges to (β, μ) and has the property that each belief system μ^n is derived from β^n using Bayes' rule. Consider the following strategy profile

$\beta^\gamma = (\beta_1^\gamma, \beta_2^\gamma)$:

$$\beta_1^\gamma(I_{h_1})(m_1') = 1 - \gamma \qquad \beta_1^\gamma(I_{h_1})(m_1) = \gamma$$

$$\beta_1^\gamma(I_{h_2})(m_2') = 1 - \gamma \qquad \beta_1^\gamma(I_{h_2})(m_2) = \gamma$$

$$\beta_2^\gamma = \beta_2 \text{ for all } I_w \in \mathcal{I}_w$$

As $\gamma \to 0$, we see that $\beta^\gamma \to \beta$. For each member of this sequence, the belief system shown in 4.6 remains unchanged, and is derived from the strategy profile via Bayes' rule. Therefore, this sequence of assessments satisfies the consistency requirement of Definition 4.1.3.

4.2. Connections to Grounding Theory

The intention of this chapter has been to develop a game-theoretic model that formalizes core aspects of conversational grounding theory. In this section I now take stock of how the model deals with the four main claims of grounding theory: (1) language use is joint action, (2) joint actions are coordinated via the common ground, (3) the minimum amount of effort that dialog participants expend to add something to the common ground is determined by the grounding criterion, and (4) the maximum amount of effort that dialog participants expend to add something to the common ground is determined by the principle of least collaborative effort.

4.2.1. Joint Actions

According to grounding theory, a joint action is one that is carried out by an ensemble of people that are acting in coordination with one another (see Section 2.1.1). The abstract game-theoretic model developed here captures at least a part of this definition. Game theory, as we have seen, is a formal theory of multiagent decision making. Players receive payoffs based on what every player does, not just on what they do in isolation. In a game of pure coordination, such as the one in Figure 4.1, the players share the same payoffs in every outcome, therefore their interests are perfectly aligned. In such a game, players must coordinate with each other in order to achieve the best possible outcome. The equilibrium solution to the game in Figure 4.1 maximizes expected utility for both players. The helper sends the best possible message, one which produces an optimal degree of confidence for a given level of cost, and the worker subsequently takes the best action given the information gained from this message. The actions of both players are directed towards maximizing their joint expected reward.

However, the game-theoretic model does not capture all of what Clark (1996) includes in his concept of a joint action. According to Clark, a joint action coordinates individual participatory actions at the level of both *content* and *process*. Coordinating on process means that dialog participants make fine-grained judgements about *timing* their contributions, which can be either verbal or non-verbal in nature. For example, in a face-to-face conversation a dialog participant can acknowledge understanding through head nods, or through backchannel responses uttered in parallel with an on-going contribution. If a workspace is visible in a task-oriented dialog, a worker can perform task actions in parallel with a helper's instructions, thereby grounding the helper's instructions without any verbal responses (see Section 2.2 for examples). In cases like these, the observations by the helper of the worker's actions form a more-or-less continuous stream. This would be very difficult to capture with a game tree.

The game-theoretic model developed in Section 4.1 is too coarse-grained to capture joint actions at this level of detail. Though it would be possible in principle to make it more fine-grained in order to capture more of the coordination phenomena described in Clark (1996), the extensive game notation is too cumbersome to make this a rewarding exercise. In Chapter 6, I suggest an approach using *multiagent decision processes* in order to partially overcome this problem. This formalism provides a compact means for implicitly defining large and complex game trees, more elaborate than the one shown in Figure 4.1. It can therefore be used to create a more fine-grained model of grounding behaviors in dialog.

4.2.2. Common Ground

According to grounding theory, joint actions are coordinated via agents' common ground of shared knowledge and beliefs. Section 2.1.2 described how full *common knowledge* is not always obtainable, and agents must sometimes settle for a lesser form of common ground, such as (probabilistic) common belief. Fortunately, these lesser forms of common ground still potentially allow for agents to achieve a high degree of coordinated behavior. This fact is particularly fortunate for the case of signaling games with partially observable actions, because in general this type of game does not allow for full common knowledge of communicated content. On the other hand, this type of game *does* allow for probabilistic common belief.

This can be shown formally using standard definitions of (common) knowledge and belief. Here I follow the set-theoretic approach established by Aumann (1976), which dovetails nicely with theory of extensive games with imperfect information.[6] We've already seen the basic ingredients of this approach in Definition 3.2.1. First, we have a set of N *players*. Second, we have a set Ω of possible *states*, which I will take to be a subset of histories H.[7] Third, for each $i \in N$, we have an *information function* denoted by Π_i, which represents the information that player i has. At a given state (history) ω, the set $\Pi_i(\omega)$ represents the set of states (histories) that player i believes are possible, given her information. This information function is determined by the structure of the game, and is closely related to the *information partitions* $\mathcal{I}_i$ of each player.[8]

In the set-theoretic approach, an event E is represented as a subset of Ω: $E \subseteq \Omega$. An agent i "knows" E at state ω if $\Pi_i(\omega) \subseteq E$, i.e., every state that i believes is possible at ω is member of the event E. The general event of "i knows E" is defined as $K_i(E) = \{\omega : \Pi_i(\omega) \subseteq E\}$. This event is the union of those elements of player i's information partition that support knowledge of E.

Other types of knowledge are built on top of these definitions. An event E is said to be *mutual knowledge* at ω if $\Pi_i(\omega) \subseteq E$ for every $i \in N$. Mutual knowledge of event E means that every agent knows that E is the case, but mutual knowledge implies nothing about what knowledge each agent has about the knowledge of others. An event E is said to be *evident knowledge* if for each $i \in N$ we have $E \subseteq K_i(E)$. The prototypical

[6] For a relatively recent overview of common knowledge from a game-theoretic perspective, see Geanakoplos (1994). The exposition and notation of this section borrows heavily from Monderer and Samet (1989).

[7] In particular, in this section I take Ω to be the subset of all possible histories after the helper has sent her message, and before the worker has taken his action. This is because I focus here on the epistemic status of the players after communication has occurred, and before the final action is taken. For a more complete picture of the relationship of epistemic logic to extensive games, see van Ditmarsch et al. (2007); Gerbrandy (2007).

[8] For example, when $P(\omega) = i$, then $\Pi_i(\omega) = \{I_i \in \mathcal{I}_i : \omega \in I_i\}$.

evident knowledge event is a public announcement that is simultaneously audible and understandable to a group of agents. An occurrence of an evident knowledge event (such as a public announcement) entails by its very nature knowledge that the event has occurred. This leads us to the following definition of common knowledge:

Definition 4.2.1. An event C is *common knowledge* at ω iff there exists an evident knowledge event E such that $\omega \in E$ and for all $i \in N$,

$$E \subseteq K_i(C)$$

In words, an event C is common knowledge among a group of agents N if and only if an event E occurs that is evident knowledge to all of them, and E entails C.[9] To use an example from (Monderer and Samet, 1989), imagine that an auctioneer publicly announces to a group of agents that “the price of the picture is $1000”. Assuming that there is no possibility that one of the agents may have misheard or misunderstood the auctioneer, the public announcement event will be evident knowledge to all of them. Since the public announcement event entails the event that the price of the picture is indeed $1000, this latter fact will become common knowledge among the group of agents.

With this formal apparatus in place, it is now easy to show that the helper and the worker will always fail to achieve common knowledge of communicated content in the game shown in Figure 4.1. Let the set of states Ω consist of the set of eight possible

[9]This definition of common knowledge entails iterative definitions of common knowledge, but requires only a finite number of steps to compute for finite Ω (Monderer and Samet, 1989; Geanakoplos, 1994).

histories at the stage before the worker takes his action:

$$\Omega = \{(d_1, m_1, o_1), (d_1, m_1, o_2), (d_2, m_2, o_1), (d_2, m_2, o_2),$$
$$(d_1, m_1', o_1'), (d_1, m_1', o_2'), (d_2, m_2', o_1'), (d_2, m_2', o_2')\}$$

Given this Ω, and the structure of the game in Figure 4.1, the information partition induced by the information function Π_h of the helper is represented as follows:

$$\Pi_h = \{\{(d_1, m_1, o_1), (d_1, m_1, o_2)\}, \{(d_1, m_1', o_1'), (d_1, m_1', o_2')\},$$
$$\{(d_2, m_2, o_1), (d_2, m_2, o_2)\}, \{(d_2, m_2', o_1'), (d_2, m_2', o_2')\}\}$$

This partition represents the fact that the helper knows the action of chance at the beginning of the game, and the helper knows what message she sent to the worker. However, the helper does not know which observation the worker received. Conversely, the information partition induced by the information function Π_w of the worker should represent the fact that he does know what observation he received, but is uncertain about the initial action of chance and the message that the helper sent:

$$\Pi_w = \{\{(d_1, m_1, o_1), (d_2, m_2, o_1)\}, \{(d_1, m_1', o_1'), (d_2, m_2', o_1')\},$$
$$\{(d_1, m_1, o_2), (d_2, m_2, o_2)\}, \{(d_1, m_1', o_2'), (d_2, m_2', o_2')\}\}$$

Now as an example, consider the case where $\omega = (d_1, m_1, o_1)$ is the actual state of the world. That is, the solution to the lock is d_1, the helper has sent message m_1, and the worker has received observation o_1. At this state, we have $\Pi_h(\omega) = \{(d_1, m_1, o_1), (d_1, m_1, o_2)\}$ and $\Pi_w(\omega) = \{(d_1, m_1, o_1), (d_2, m_2, o_1)\}$. Therefore, it fails to be the case that either player has even *individual knowledge* that the event $E = \{\omega\}$ has occurred, since neither $\Pi_h(\omega)$ nor $\Pi_w(\omega)$ is a subset of E. So common knowledge of E is clearly impossible in this situation. This conclusion holds no matter which state in Ω is the true state.

Since common knowledge is impossible here, we need a weaker notion of common ground to explain how communication and coordination occurs during the game shown in Figure 4.1. This weaker notion is provided by Monderer and Samet (1989), who extend the set-theoretic definitions of (common) knowledge to the case of probabilistic (common) belief. An agent i is said to believe event E with probability at least p at state ω ("i p-believes E at state ω") if he assigns probability p or higher to the set of states contained in E, given that ω is the true state of the world. Formally, we have $B_i^p(E) = \{\omega : Pr(E|\Pi_i(\omega)) \geq p\}$. Given this definition, *mutual p-belief* and *evident p-belief* are defined in a straightforward way as generalizations to the set-theoretic definitions of mutual and evident knowledge, leading to the following definition of common p-belief:

Definition 4.2.2. An event C is *common p-belief* at ω iff there exists an evident p-belief event E (i.e., for each $i \in N$ we have $E \subseteq B_i^p(E)$) such that $\omega \in E$ and for all $i \in N$,

$$E \subseteq B_i^p(C)$$

For finite models, Definition 4.2.2 is a pure generalization of Definition 4.2.1, since common knowledge is nothing more than common p-belief with $p = 1$ (Monderer and Samet, 1989). Just as Definition 4.2.1 entails iterative notions of common knowledge,

Table 4.1. Belief states of players given uncertain communication

State	**Prior**	$H\|d_1$	$H\|d_2$	$W\|o'_1$	$W\|o'_2$
(d_1, m'_1, o'_1)	$(1-\epsilon')/2$	$1-\epsilon'$	0	$1-\epsilon'$	0
(d_1, m'_1, o'_2)	$\epsilon'/2$	ϵ'	0	0	ϵ'
(d_2, m'_2, o'_2)	$(1-\epsilon')/2$	0	$1-\epsilon'$	0	$1-\epsilon'$
(d_2, m'_2, o'_1)	$\epsilon'/2$	0	ϵ'	ϵ'	0

Definition 4.2.2 entails iterative notions of common p-belief (where i p-believes that j p-believes (etc.) that C).

To see how to apply the notion of common p-belief to the game in Figure 4.1, consider the belief state possibilities of the game summarized in Table 4.1.[10] The leftmost column again shows the possible states of the game at the point before the worker selects his action. Here we consider only states that have a non-zero chance of occurring, given the equilibrium strategy profile. The next column shows the prior probabilities for these states. For example, before the game is played the state (d_1, m'_1, o'_1) has a prior probability of $(1-\epsilon')/2$. The third and fourth columns show updated probabilities for the helper after chance selects either d_1 or d_2 (respectively) as the lock solution at the beginning of the game. The fifth and sixth columns show the updated probabilities for the worker after chance selects either observation o'_1 or o'_2 (respectively).

Now consider the state (d_1, m'_1, o'_1), where chance selects d_1 as the solution to the lock at the beginning of the game, the helper selects action m'_1, and chance then provides the worker with observation o'_1. Table 4.1 highlights the cells that have non-zero probability for each player when this is the true state of the game. In this case, the helper and

[10]These are generated with the assumption that we have the belief system as shown in Equations 4.6 and the strategies as shown in Equations 4.7.

the worker each assign the true state a probability of $1-\epsilon'$; therefore, the players have achieved mutual p-belief in state (d_1, m_1', o_1') at the p-level of $1-\epsilon'$. In fact, we can say something stronger. Letting $E = \{(d_1, m_1', o_1')\}$, we have:

$$B_h^{(1-\epsilon')}(E) = \{(d_1, m_1', o_1'), (d_1, m_1', o_2')\}$$

$$B_w^{(1-\epsilon')}(E) = \{(d_1, m_1', o_1'), (d_2, m_2', o_1')\}$$

Therefore, since $E \subseteq B_h^{(1-\epsilon')}(E)$ and $E \subseteq B_w^{(1-\epsilon')}(E)$, it is also the case that the helper and the worker have achieved *common p-belief* in event $E = \{(d_1, m_1', o_1')\}$ at the p-level of $1-\epsilon'$. So the helper believes with probability $1-\epsilon'$, that the worker believes with probability $1-\epsilon'$ (etc.) that (d_1, m_1', o_1') is the true state of the game. As ϵ' approaches 0, this common p-belief can get arbitrarily close to full common knowledge.

This scenario describes a "good" case, where the message and observation "coincide" and proper coordination occurs. However, consider another possible state in Table 4.1: (d_1, m_1', o_2'). In this state, the helper still assigns probability $1-\epsilon'$ to the state (d_1, m_1', o_1'), because the helper only knows that d_1 is true and that she has sent message m_1'. She only assigns a probability ϵ' to the state that actually occurs, namely (d_1, m_1', o_2'). On the other hand, the worker assigns probability 0 to state (d_1, m_1', o_1') because this is inconsistent with the observation o_2' that he has received. Given this observation, the worker assigns probability ϵ' to the true state, and probability $1-\epsilon'$ to state (d_2, m_2', o_2'). Therefore, the worker will take action s_2, given his equilibrium strategy. Common p-belief for the true state of the game exists only at the p-level ϵ', and coordination therefore fails, even though both players have played their part of the equilibrium strategy profile.

Here we see how achieving only approximate common knowledge implies that coordination may potentially fail. Given the structure of the game in Figure 4.1, failure occurs with probability ϵ'. This is indeed the best we can do, since no other alternative strategy is available that gives a higher expected payoff to both players. Imagine, however, an extended version of this game wherein the helper has an additional two message actions at her disposal, m_1'' and m_2''. Furthermore, these actions are associated with cost C'' and error rates ϵ'', where $C < C' < C''$ and $\epsilon'' < \epsilon' < \epsilon$. In other words, these messages have the smallest possibility of error of any of the helper's messages, but they come at the highest cost. The question is: should the helper use these messages, given that they lead to higher likelihood of successful coordination?

The answer to this question was given in Section 4.1. It is a simple calculation to determine if $\frac{C''-C'}{M+L}$ is greater than, less than, or equal to $\epsilon' - \epsilon''$. If the former quantity is greater than the latter, the helper should stick with the strategy of using the lesser cost (and lesser precision) messages. The extra cost of the higher precision messages negates the increase in expected payoffs. However, if the former quantity is less than the latter, then the helper should use the higher cost (and higher precision) messages. In this case, the extra cost is compensated for by the increase in precision.

It is easy to see that this type of argument generalizes to any number of messages. Let us suppose that the helper has at her disposal a set of n messages, which can be ordered according to both cost and certainty. This situation is depicted in Figure 4.2, where m_1 is the cheapest (and least certain) message, while m_n is the most expensive (and most certain) message. Given this situation, the helper's optimal message m^* is the

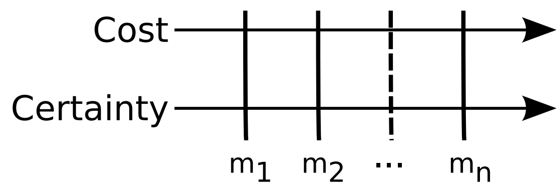

Figure 4.2. Trading off message cost and uncertainty

most certain message that obeys the following constraint:

$$\frac{(C^* - C_1)}{M + L} < (\epsilon_1 - \epsilon^*)$$

Where M is the reward received by both players for successfully coordinating a correct action, L is the penalty for taking an incorrect action, C^* and ϵ^* are the cost and error likelihood associated with m^*, and C_1 and ϵ_1 are the cost and error likelihood associated with the least costly (and least certain) message m_1.

4.2.3. Grounding Criterion

Given the discussion of common ground in Section 4.2.2, I present the following definition of the grounding criterion, as it applies to signaling games with partially observable actions:

Definition 4.2.3. The *grounding criterion for a signaling game with partially observable actions* is the p-level at which the players believe that common p-belief has been obtained, at the stage of the game before the final action is taken, given that players have belief system μ^* and play their equilibrium strategies β_1^* and β_2^*.

Definition 4.2.3 is a formal analog to grounding theory's definition of the grounding criterion as the requirement that "the contributor and his or her partners mutually believe that the partners have understood what the contributor meant to a criterion sufficient for current purposes" (Clark and Brennan, 1991, p. 129). Definition 4.2.3 doesn't literally require "mutual belief", much less common belief that the partners have understood what the contributer meant. Section 4.2.2 demonstrates that this is an impossibly high requirement in general. More reasonably, it states that the players will try to achieve the strongest possible belief that common p-belief has been achieved, given the costs that are required to reach this level of precision. This calculation falls out automatically from playing the equilibrium strategy profile.

Definition 4.2.3 follows directly from the structure of a signaling game with partially observable actions and from the general formal apparatus of game theory. Therefore, the grounding criterion does not need to be independently stipulated.

4.2.4. Least Collaborative Effort

Given the discussion of common ground in Section 4.2.2, I propose the following definition as an analog to grounding theory's principle of least collaborative effort:

Definition 4.2.4. The *optimal collaborative effort in a signaling game with partially observable actions* is the cost C^* associated with the payoffs that result from following the equilibrium strategy profile $\beta^* = (\beta_1^*, \beta_2^*)$.

In a game of pure coordination, the cost C^* is shared by both players. It can therefore be regarded as "collaborative effort". If dialog participants were to choose a strategy with higher costs than the equilibrium strategy, this would result in a lower expected payoff to both players, since the extra cost is not justified by decrease in message error. In short, no equilibrium strategy will violate the following constraint:

$$\frac{(C^* - C_{min})}{M + L} \leq (\epsilon_{max} - \epsilon^*)$$

Where C_{min} and ϵ_{max} are the cost and error likelihood associated with the least costly and least precise message. Again, Definition 4.2.4 follows directly from the structure of a signaling game with partially observable actions and from the general formal apparatus of game theory. Therefore, a formal analog to the principle of least collaborative effort does not need to be independently stipulated.

4.3. Incremental Referring Expressions

The abstract formal model developed in the previous section says nothing about how costs are to be associated with messages, and how the precision of a message should be determined. In this section, I make this model a bit less abstract by applying it to the case of incremental referring expressions, like those described in Chapter 2. Example 2.1 from that chapter is repeated here:

(4.8) **A.** And the next one is the one with the triangle to the right...

B. Okay.

A. With the square connected to it.

The question to be answered is this: why did speaker A choose to split the referring expression into two installments, thereby taking up two turns in the dialog rather than a single turn? For that matter, why did speaker A not split the referring expression into more than two installments? The answer, according to grounding theory, is that the speaker is trading off the extra cost of the turn, for a reduction in uncertainty that the meaning of the referential expression has been added to the participants' common ground. Presumably, the extra precision derived from using an extra turn was adequate compensation for the additional turn cost. Also presumably, the extra precision associated with taking more than two turns was *not* worth the additional turn cost. According to Brennan (1998, p. 13):

> Speakers trade off the costs of grounding with the benefits . . . [they] are more likely to break a contribution up into installments when they have a high grounding criterion (such as when it is important that an addressee understand an utterance verbatim) or when they use media in which utterances are ephemeral and the costs of changing speakers are relatively low (such as telephone conversations, as compared to email).

The game-theoretic model presented in Section 4.1, and Inequality 4.5 in particular, points towards a formal analysis of this intuition. In this model, there is a simple way to calculate which of two message is preferable, given the relationship between their relative costs and their relative precision. What is required to apply this model to the case of incremental referring expressions is a principled way of linking installment size with both cost and levels of precision.

More precisely, assume we have a referring expression e, whose content must be communicated by one speaker to another (e.g., the helper to the worker in a referential communication task). Abstractly, expression e can be said to carry a certain amount of information, such as the number of property values that must be specified in order for e to

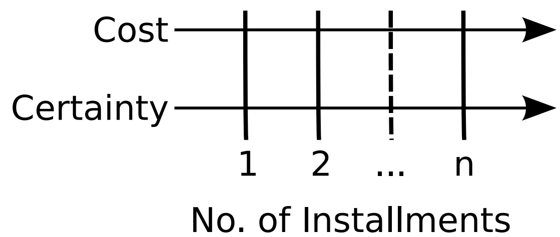

Figure 4.3. Trading off message cost and uncertainty with installments

uniquely identify its intended referent, given a certain distractor set of possible referents (Dale and Reiter, 1995). There is an inherent cost to e which is related to the number of property values it encodes, and the lexical complexity of encoding and decoding this information. However, assuming that e encodes a certain set of uniquely identifying property values, the speaker has a choice to make, and that is how many installments she wants to divide the expression into. This situation is represented in Figure 4.3 (compare this with Figure 4.2).

By hypothesis, chunking an expression into smaller and more numerous installments increases the overall cost of communicating e: it takes more turns, and therefore more time and effort to execute and coordinate turn taking. To model this, we can formulate a *cost function* that takes as argument the number of installment turns that e is divided over, and returns as a result the *cost* (above and beyond the inherent cost of e) of taking these extra turns. We can define a cost function as a simple linear function, $cost(turns, properties) = (a \times turns) + (b \times properties)$, where the constant a is the turn cost, and the constant b is the cost of encoding/decoding a property value. Figure 4.4 plots various cost functions where $b = .01$, $properties = 10$, and $turns$ ranges from 1 to 10. The slope of each cost

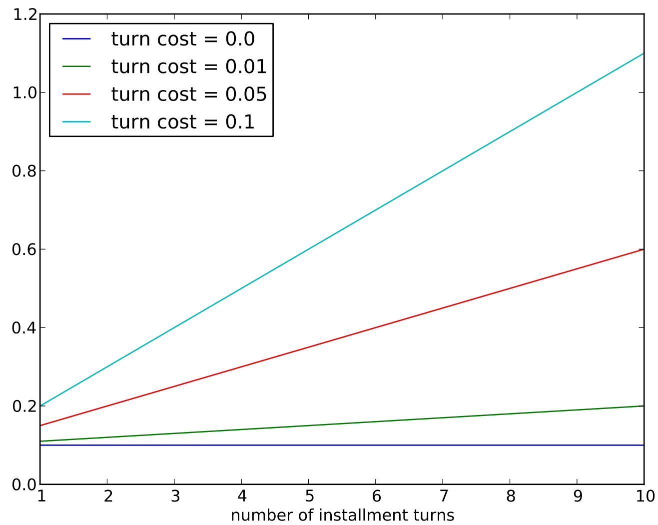

Figure 4.4. Cost functions with selected turn costs

function in Figure 4.4 is determined by the turn cost constant a, while the property cost constant b determines where the function approaches the y-intercept.[11]

Since using smaller and more numerous installments increases cost by increasing the number of turns, there must be a counter-balancing benefit for this strategy to ever make sense. By hypothesis, chunking expression e into smaller installments makes it easier for the hearer to extract the intended meaning of e, with smaller increments less error prone than larger increments. To model this, we can formulate an *error function* that takes as argument the *size of an installment*, and returns as a result the *error rate* (or uncertainty level) associated with it. Unlike the cost functions shown in Figure 4.4, I assume here

[11] The linearity assumption here is not crucial. For current purposes, the only crucial property of this function is that it monotonically increases with the number of installment turns.

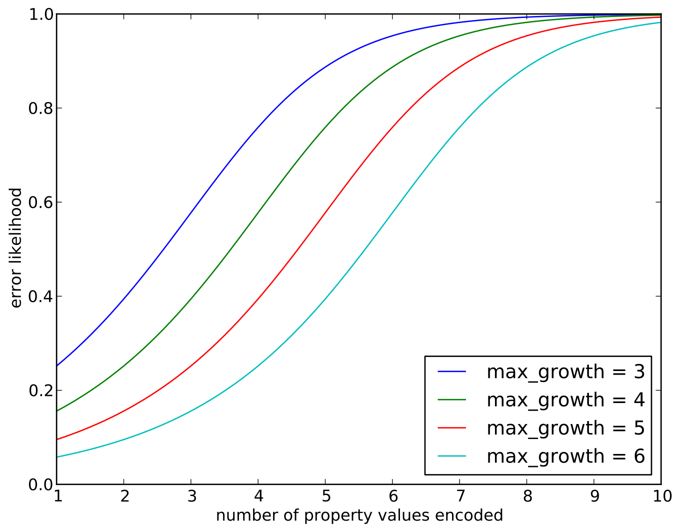

Figure 4.5. Error model with selected locations of maximum growth

that the error function is an s-shaped *sigmoid* function.[12] Figure 4.5 plots various error functions, which take size (in terms of number of property values encoded) as input.

The functions in Figure 4.5 are generated using a *generalised logistic* function (Richards, 1959), which has parameters controlling lower and upper limits, rate of growth, and location on the x-axis where maximum growth occurs. Various parameter settings for this last parameter are shown in the figure: this could be used to model such phenomena as different noise levels in the speech channel, or different levels of distraction for the listener. The intuition is that for "noisier" environments, the saturation of processing ability occurs earlier, with less information. In more ideal environments, more information can

[12]The idea behind using a sigmoid function is that there is a nonlinear relationship between the size of a message, and a hearer's ability to decode this input, due to nonlinear properties of short-term memory (see Cowan (2000)). For current purposes, the only crucial feature of this function is that it monotonically increases with the length of an installment turn.

be processed without difficulty. The lower limit for each function in Figure 4.5 is set to 0 (indicating no possibility for error), and the upper limit is set to 1 (indicating complete failure of information transmission).[13]

Figure 4.6 combines the cost and error functions into a set of expected value calculations. Each plot shows the transmission of 10 property values with a cost constant of .01, the error slope set to 2, maximum growth rate at 5 property values, $M = 1$, and $L = 0$ (with M and L the payoffs as indicated in Inequality 4.5). The red dots on each curve indicates the maximum value of the function. The values on the x-axis at which these maxima occur indicate the best choice of installment size for obtaining maximum expected value. At one extreme, if we assume a turn cost of 0, the maximum expected value is generated by taking 10 full installments, each of which transmits a single property value. Since there is no cost to taking turns, and taking extra turns drives down the error rate, it makes sense to take as many turns as possible. On the other end, with a turn cost of 0.1, the maximum expected value is obtained by taking approximately three turns, each transmitting an average of one third of the total number of property values. Turn costs in between these two values produces intermediate maxima.

Inequality 4.5 indicates that in addition to message cost, we must also look at the intrinsic payoffs of the game to determine the expected utility of a given message. If either M or L goes up, then the proportion $\frac{\Delta C}{M+L}$ goes down. Therefore, a game with higher M or L relative to cost will potentially justify more installment turns. This situation is demonstrated in Figure 4.7, which plots expected utility for several values of L, holding

[13]Note that the x-axis here is in number of property values, not turns. If we were to plot the graph as a function of number of installment turns for a particular referring expression e, we would average the number of property values that e encodes per number of turns, and the graph would be the inverse of the one in Figure 4.5.

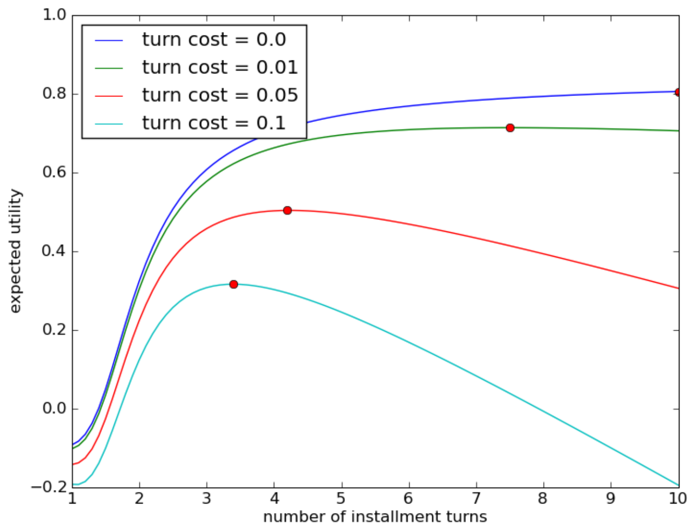

Figure 4.6. Expected utility for selected turn costs

property value cost and turn cost constant at 0.01 and 0.05, error slope constant at 2.0, maximum growth rate at 5 property values, and $M = 1$. These graphs show that when the cost of making a task error increases, the expected value of taking more installment turns also increases. Simply put, it pays to be careful when bad consequences can result from a mistake. Increasing M would have the same result: it also pays to be careful when very good consequences can result from getting it right.

Figure 4.8 demonstrates that holding the ratio $\frac{\Delta C}{M+L}$ constant at 0.1 generates expected value functions that differ in absolute value, but produce maxima at identical location on the x-axis. For various turn costs and values for M and L, as long as $\frac{\Delta C}{M+L} = 0.1$, the preferred number of installments for 10 property values remains approximately 3. What

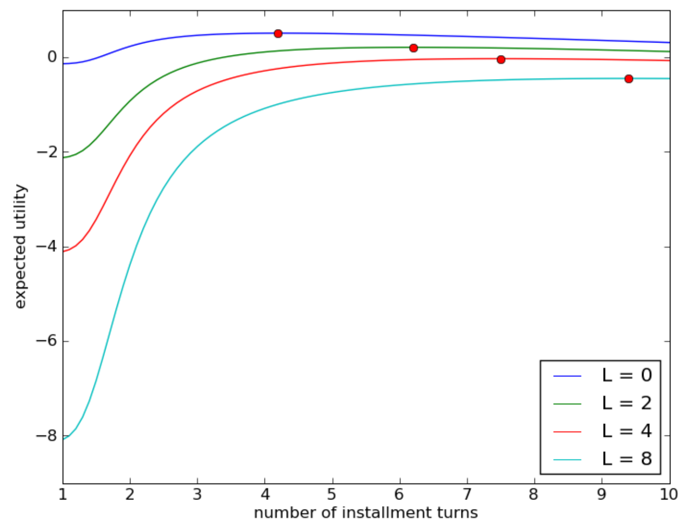

Figure 4.7. Expected utility for selected task error costs

this figure tells us is that increasing turn cost may have no effect on installment size if the payoffs are increased proportionately.

Finally, Inequality 4.5 says that changing $\Delta(\epsilon)$ should also have an effect expected utility. Figure 4.9 demonstrates that changing the properties of the error function affects the location of maxima. The plots in this figure keep all other parameters constant, only the location (in terms of number of property values) where maximum growth occurs is varied. When maximum growth occurs for smaller quantities of information, modeling a situation where there is more "noise" in the channel, the model predicts a greater number of installment turns, with smaller installments. When maximum growth occurs for higher quantities of information, modeling a less noisy channel, the model predicts a smaller number of installment turns, with larger installments.

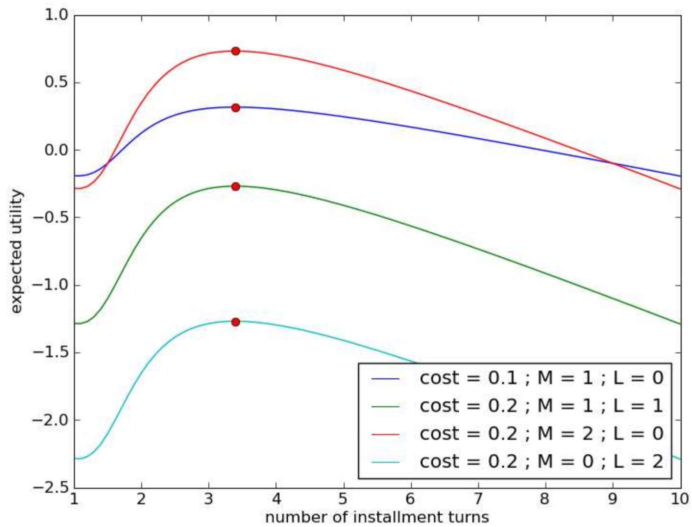

Figure 4.8. Expected utility for constant $\frac{\Delta C}{M+L}$

This model of incremental referring expression abstracts away from many details, and the properties of the cost and error functions are speculative. However, it is precise enough to make empirical predictions about dialog behaviors, and it is a step towards a theory that can be utilized by computational conversational agents to automatically chunk referring expressions into appropriately sized installments that can be reliably processed by users. To my knowledge, this issue of automatically generating appropriately sized chunks of speech or text in a principled way has not received attention in the literature on dialog systems or natural language generation (cf., Dale and Reiter, 1995; Reiter and Dale, 2000; Jurafsky and Martin, 2009). Despite this lack of attention, providing computational dialog systems with the means for automatically generating appropriately sized installments would be a step forward in naturalness and user satisfaction for such systems.

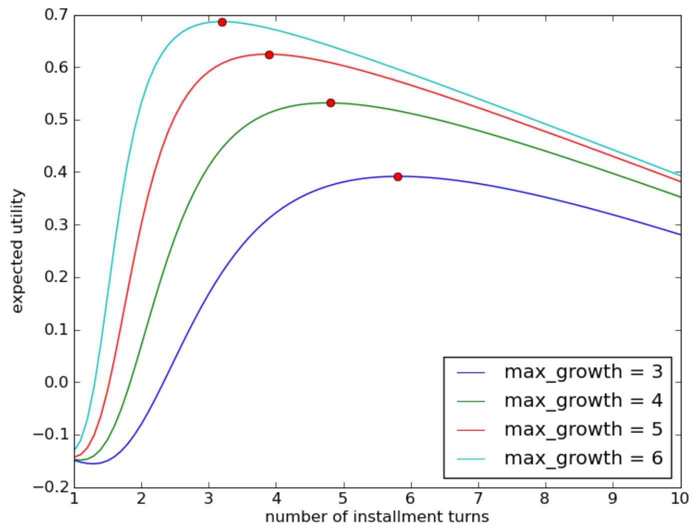

Figure 4.9. Expected utility for selected error functions

One way of viewing the model described in this section is as a complement to existing computational models of referring expression generation. For example, in an influential article Dale and Reiter (1995) describe an algorithm for generating referring expressions that is capable of generating a uniquely identifying description for a target object, given a distractor set of objects from which the target must be distinguished. The output of their algorithm is a set of property-value pairs that uniquely identify the target to a given hearer. However, the algorithm says nothing about how these property-value pairs are to be lexicalized, nor whether or not they are all to be included in a single dialog turn. This set of property-value pairs can therefore be treated as the input to the model described in this section, which can then be used to decide how many dialog turns to use in communicating this information.

4.4. Other Models of Grounding

Since the pioneering development of conversational grounding theory by Clark and colleagues, there have been several attempts to formalize it, mainly by researchers interested in building computational conversational agents. These researchers have come to realize that a practical dialogue system must deal with grounding behaviors in order to achieve any reasonable level of performance.[14] For example, Traum (1994) estimated that approximately half of all utterances in the TRAINS dialog corpus (Allen et al., 2000) dealt with coordinating understanding between participants, rather than dealing with the actual task to which the dialog was directed. Understanding how to process and generate grounding behaviors is crucial if a conversational agent is to even approximate human-like performance in dialog.

4.4.1. The Contribution Model

One of the first attempts to formalize grounding theory came from Clark and Schaefer (1987, 1989), who presented the *contribution model* of grounding. The contribution model is essentially a syntactic approach to grounding in dialog, organizing multiple dialog utterances into a single *contribution graph*. A contribution graph is composed of *contributions*, which in turn are composed of *presentations* and *acceptances*. Each presentation and acceptance is either an utterance (the basic unit from which larger scale units are composed), or another contribution that is embedded underneath it. Both cases are demonstrated in Figure 4.10, taken from Clark and Schaefer (1989). In this figure, nodes

[14]For a recent overview of computational approaches to dialog, see the relevant chapter in Jurafsky and Martin (2009).

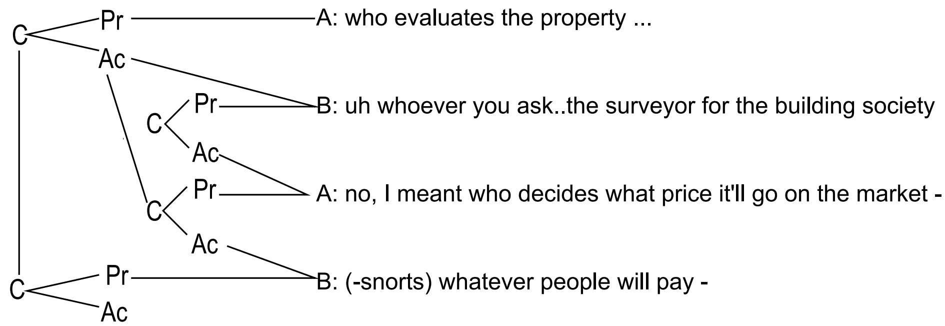

Figure 4.10. Contribution graph example (from Clark and Schaefer (1989))

labelled “C” are contributions, nodes labelled “Pr” are presentations, and nodes labelled “Ac” are acceptances.

Referring to Figure 4.10, speaker A makes an initial presentation (a question), and B follows with a relevant next utterance, an answer to this question. B’s answer to A is both an acceptance of the previous presentation, and a separate presentation of its own. At this point a complication arises, because with A’s next presentation he makes it clear that B has misunderstood A’s original question. This presentation has the result of initiating an embedded repair sequence. B subsequently responds with a new presentation, which acts as an acceptance to both the embedded repair sequence and the original top-level contribution. As a separate presentation, it furthermore initiates a new contribution. The contribution model therefore enforces the view that conversational acts are *joint* actions. A conversational act is not executed by virtue of a single participant’s utterance. It is only complete when it has been accepted as such by the addressee.

However, Traum (1999) highlights several conceptual problems with the contribution model, the effects of which are to limit its usefulness as a computational model of grounding. From his point of view, the most critical problem is that it is not possible to evaluate the grounding state of a contribution until after the dialog of which it is part is completed, since each new utterance may or may not be a complete acceptance of the previous presentation, and each new presentation must have its own acceptance before it can be committed to record as a contribution. The ever present possibility of embedded repair dialogs means that contribution graphs are always provisional, until the entire dialog is completely finished. As pointed out by Cahn and Brennan (1999), this problem is due in no small measure to the fact that Clark and Schaefer (1989) merge both participants' models of the dialog into a single contribution graph. This is an error: each participant should have their own local model of the dialog, and each participant should also be equipped with a procedure for constructing and modifying their local contribution graphs incrementally, as a dialog unfolds in time.

Even with the improvements proposed by Cahn and Brennan (1999), however, the contribution model has had very little subsequent influence on computational models of dialog. Why this is so is clear when we refer back to the core ideas of grounding theory in Chapter 2. The notion of a conversational common ground refers to epistemic states such as (probabilistic) common knowledge and belief. Contribution graphs have an obscure relationship to such epistemic states, and they certainly have no means to represent degrees of uncertainty in the common ground. Furthermore, the grounding criterion and the principle of least collaborative effort both rely on the assumption that dialog participants care about the costs and payoffs of taking particular actions. Contribution graphs have no means for representing costs and payoffs. Without a representation of costs, payoffs, or uncertainty in the common ground, contribution graphs obviously can say nothing about trading them off against each other. In short, core ideas of grounding theory are simply not addressed by the contribution model. In contrast, these core ideas *are* directly addressed by the game-theoretic model of grounding proposed in this chapter.

4.4.2. The Grounding Acts Model

In seminal work, Traum (1994) presents a computational model of grounding that overcomes some of the problems of the contribution model. This model, the *grounding acts model*, was one of the first formal models of grounding to be used in a fully functional spoken dialog system (Allen et al., 2000). Unlike the contribution model, the grounding acts model tracks the grounding status of a dialog contribution at the level of individual utterances. This means that the model does not require look-ahead capability to determine the status of a contribution, and the grounding status of a dialog contribution can be determined on-line, while a dialog is unfolding in time.

The grounding acts model is situated within a larger theory of *conversation acts* (Traum, 1994; Stent, 2002). The theory of conversation acts is a generalization of classic speech act theory (Austin, 1962; Searle, 1969) to account for Clark's observation that utterances in dialog are (constitutive parts of) joint actions. Traum (1994) presents four levels of conversation acts, two of which are relevant here: (1) core speech acts, corresponding to the acts of classical speech act theory, and (2) grounding acts, the effects of which are to update the grounding status of conversational information.

The effect of execution of a core speech act, such as `Inform` or `Request`, is to add an increment of information to the conversational common ground. This increment of information is called by Traum a *discourse unit* (DU). Discourse units correspond roughly

Table 4.2. Grounding act typology (from Traum (1999))

Label	**Description**
initiate	Begin new DU, content separate from previous uncompleted DUs
continue	Same agent adds related content to an open DU
acknowledge	Demonstrate or claim understanding of previous material by other agent
repair	Correct (potential) misunderstanding of DU content
Request Repair	Signal lack of understanding
Request Ack	Signal for other to acknowledge
cancel	Stop work on DU, leaving it ungrounded and ungroundable

to the contributions of the contribution model (Section 4.4.1). Like contributions, core speech acts are not atomic – they are composed from a sequence of grounding acts. The set of grounding acts proposed by Traum (1994) is listed in Table 4.2. Referring to this table, an `initiate` act introduces a new DU into a conversation, and other grounding acts operate on this DU to ground it, repair it, or discard it if necessary.

In the computational model of grounding presented by Traum, the lifecycle of a discourse unit is described by a finite-state transition network. The states of the network represent the states of a particular DU, and state transitions are triggered by the execution of utterance level grounding acts. The finite-state transition network proposed by Traum (1994) is shown in Table 4.3. Here, grounding act types are superscripted with either "I" or "R", indicating (respectively) that the act in question was executed by either the original initiator of the DU, or by the responder to it. Empty cells in the table indicate where the action on the left is not applicable to the state in question. The

Table 4.3. Discourse unit state transition diagram (from Traum (1994))

Next Act	In State						
	S	**1**	**2**	**3**	**4**	**F**	**D**
initiate$^{\mathbf{I}}$	1						
continue$^{\mathbf{I}}$		1			4		
continue$^{\mathbf{R}}$			2	3			
repair$^{\mathbf{I}}$		1	1	1	4	1	
repair$^{\mathbf{R}}$		3	2	3	3	3	
ReqRepair$^{\mathbf{I}}$			4	4	4	4	
ReqRepair$^{\mathbf{R}}$		2	2	2	2	2	
ack$^{\mathbf{I}}$				F	1	F	
ack$^{\mathbf{R}}$		F	F			F	
ReqAck$^{\mathbf{I}}$		1				1	
ReqAck$^{\mathbf{R}}$				3		3	
cancel$^{\mathbf{I}}$		D	D	D	D	D	
cancel$^{\mathbf{R}}$			1	1		D	

boundary states are **S**, **F**, and **D**: state **S** is the start state, before a DU is introduced, state **F** represents a final, grounded state for a DU, and state **D** represents a defective final state, where work on a DU is abandoned.

As mentioned above, the grounding acts model has the great advantage over the contribution model that the state of a given dialog contribution (or DU) is determined deterministically by the history of grounding acts that has occurred – no look-ahead is required. Given a sequence of grounding acts, the state transition table in 4.3 gives us the state of a DU. If the state is **F**, the DU is added to an agent's representation of the conversational common ground. If the state is **D**, the agent realizes that failure has occurred, and the content of the DU needs to be re-introduced with an `initiate` act if it is important to do so. If the current state is something other than **F** or **D**, then the

Table 4.4. Multiple acknowledgement example (from Traum (1994))

Act Type	**Utt. #**	**Speaker:Utterance**
continue	9.2	M: move the engine
continue	9.3	M: at Avon
repair	9.4	M: engine E
continue	10.1	M: to
repair	11.1	S: engine E1
acknowledge	11.1	M: E1
acknowledge	12.1	S: okay

agent will realize that more effort is required (by the agent or by the other participant) in order to get to a final state.

Although the grounding acts model is a big improvement over contribution model, Traum (1999) highlights a few deficiencies, two of which are relevant to the topic of this thesis. First, like the contribution model, the grounding acts model assumes that grounding is binary in nature – each DU is either grounded or not grounded, and there is nothing in between. There is no notion of degree of uncertainty about the grounding status of a piece of information. This implies that the grounding acts model has no analog to the grounding criterion of Clark's grounding theory. That this is not just a theoretical problem is shown by the stretch of dialog in Table 4.4, taken from Traum (1994). According to the state transition table 4.3, the DU referred to at the beginning of this stretch of dialog is in the grounded state after utterance number 11.1. However, speaker S follows utterance 11.1 with an additional acknowledgement at utterance 12.1, despite the fact that, according to the grounding acts model, this DU is already grounded. Multiple acknowledgements like this are not forbidden by the model, but they aren't explained either. As Traum points out, this type of phenomenon can be accounted for with a model that permits multiple levels of grounding, with the level of grounding determined by the grounding criterion.

The second problem with the grounding acts model is that it has no means to represent the *cost* of a conversation act. This implies that it has no analog to grounding theory's principle of least collaborative effort. Again referring to the stretch of dialog in Table 4.4, there is no explanation for why the sequence of acknowledgements stops where it does, rather than continuing on for an indefinite number of more turns. These two problems – having no account for either degree of uncertainty in the common ground or the cost of an action – are ones that the game-theoretic model of this chapter directly addresses.

One way of looking at these two models is that they are essentially complementary to one another, in much the same way that the account of incremental referring expressions in Section 4.3 is complementary to the model of referring expression generation given in Dale and Reiter (1995). The grounding acts model fits into a larger model of conversation that provides an empirically derived ontology of conversation acts that occur in task-oriented dialog. Furthermore, the grounding acts model has a level of granularity and specificity that makes it directly implementable in a computational dialog system. The game-theoretic model as presented in this chapter lacks this level of granularity. However, there is nothing to prevent the game-theoretic model from being modified to adopt the grounding act model's ontology of conversation act types, by including them in the game players' action sets. This move would marry the strengths of the grounding acts model (an empirically derived ontology of act types) with the strengths of the game-theoretic model (a normative approach to action that relates action to uncertainty and cost). One way in which this combination could be accomplished is suggested in Chapter 6.

4.4.3. Decision-Theoretic Models

Traum followed up on his grounding acts model with the first attempt to model grounding theory using the tools of decision theory (Traum and Dillenbourg, 1998; Traum, 1999). This attempt was motivated by awareness of the deficiencies of the grounding act model noted in Section 4.4.2, as well as the desire to create a model of grounding that generalizes to other modalities besides speech. In these papers, Traum and Dillenbourg propose a class of "generalized grounding acts", which (unlike the grounding acts model of Section 4.4.2) can be either linguistic or non-linguistic in nature. They use the term $\alpha \mapsto \mu$ to indicate that generalized grounding act α is performed in order to contribute to the grounding of content μ.

Traum and Dillenbourg propose that the utility of a grounding act can be estimated by taking into account four quantities. The first quantity is the grounding criterion of content μ (referred to by the term $GC(\mu)$). The second is the cost of an action α ($C(\alpha)$). The third and fourth quantities are the "groundedness" of content μ before action α is taken ($G(\mu)$), and the groundedness of μ after action α is taken ($G_\alpha(\mu)$). These terms are all combined in the following relationship:

$$U(\alpha \mapsto \mu) \propto \frac{GC(\mu) \times (G_\alpha(\mu) - G(\mu))}{C(\alpha)} \tag{4.9}$$

In words, the utility of action α executed to ground content μ is proportional to the grounding criterion of μ and to the increase in groundedness of μ that occurs after executing α, and is inversely proportional to the cost of executing α.

Relation 4.9 is clearly reminiscent of Inequality 4.5. However, the relation is under-specified in its details, and isn't derived from a general theory of multiagent action. For example, the quantities denoted by the functions GC and C are presumably something like utilities, while the quantities denoted by G and G_α are presumably something like probabilities. Once these quantities are spelled out in more detail, the logical path to take is the one taken in this thesis – use the dominant existing decision-theoretic approach to multiagent action (i.e., game theory) to derive the relationship among these quantities from general principles.[15] Traum and Dillenbourg's first attempt at using decision theory to model grounding theory can therefore be seen as an early step down the path taken in this thesis.

The decision-theoretic approach to grounding was taken further in the work of Eric Horvitz and Tim Paek (Horvitz and Paek, 1999; Paek and Horvitz, 1999, 2000a,b). Horvitz and Paek viewed the problem of conversational grounding as an application of general principles of decision making under uncertainty. They recognized that formalizing conversational grounding as the process of reaching a sufficient level of mutual belief among dialog participants requires at least two key components: "first, the quantification of belief, and second, a consideration of how 'sufficient' belief changes depending on the stakes of pursuing various actions in different contexts" (Paek and Horvitz, 2000a, p.2). Horvitz and Peak adopted decision theory for this purpose, because in decision theory "uncertainty, actions, and utilities all come together in a unifying mathematical framework" (Ibid., p. 2).

[15]One payoff, for example, is that applying general principles means that we don't need to stipulate domain dependent constructs like the grounding criterion and the principle of least collaborative effort.

To illustrate the decision-theoretic approach, Paek and Horvitz (2000a) adopted the perspective of a listener who must decide whether or not to make a repair to a speaker's utterance. The utility of making a repair depends on whether or not the listener has comprehended the speaker's utterance. Denote comprehension of the speaker's utterance as C, and non-comprehension as $\neg C$. Furthermore, denote the event of repairing the speaker's utterance as R and not repairing as $\neg R$. The listener must then decide which action to take based on the following four utilities:

- $u(C, R)$: the utility of repairing when the listener comprehends the utterance
- $u(C, \neg R)$: the utility of not repairing when the listener comprehends the utterance
- $u(\neg C, R)$: the utility of repairing when the listener does not comprehend the utterance
- $u(\neg C, \neg R)$: the utility of not repairing when the listener does not comprehend the utterance

Intuition suggests that these four utilities can be partially ordered with respect to each other. Namely, $u(C, \neg R) \geq u(C, R)$ and $u(\neg C, R) \geq u(\neg C, \neg R)$. Given these four utilities, the *expected utility* of taking action R or $\neg R$ depends on the degree of likelihood the listener assigns to C. This is defined as $p(C|E)$, which is the probability of utterance comprehension given all the evidence E the listener has obtained so far. This situation is depicted in Figure 4.11, where the x-axis denotes the probability of listener comprehension, and the y-axis denotes utilities.

Figure 4.11 demonstrates how different levels of $p(C|E)$ affect which decision should be taken by the listener. As this quantity increases from left to right, the expected utility

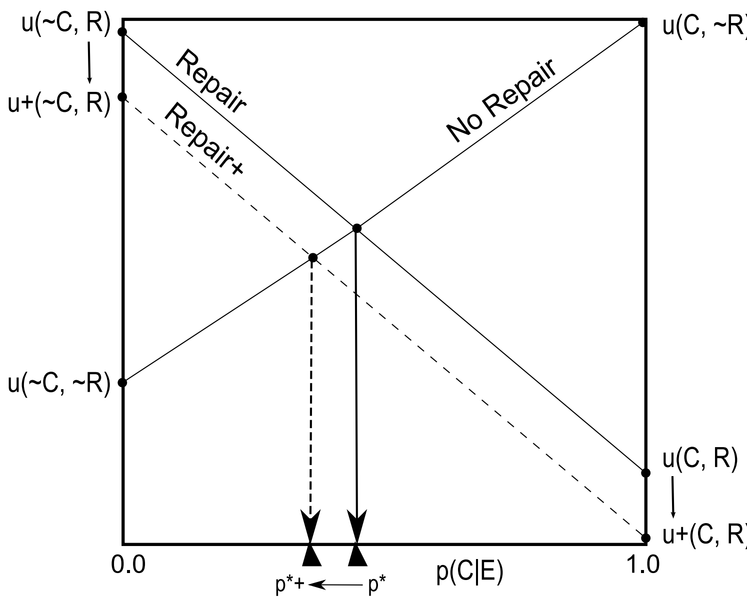

Figure 4.11. Deciding when to repair (from Paek and Horvitz (2000a))

of taking action R declines, while the expected utility of taking action $\neg R$ increases. At some point, denoted by $p^\star$, the expected utilities are equal. Above this threshold value, it is better to take action $\neg R$, while below this value it is better to take action R. The crossover point depends on the specific values of the four utilities described above. For example, in Figure 4.11 the dashed line indicates a situation where the utility function for repairing is denoted by u_+. If this latter utility function were in effect, then the crossover point (denoted in the figure by $p^\star_+$) would be shifted to the left.

The decision-theoretic approach of Horvitz and Paek comes very close to the game-theoretic approach to grounding taken in this chapter. The mathematical approaches are very similar, since game theory shares its mathematical foundations with decision theory. In this sense, I view my approach to grounding as an evolution of that expressed by Horvitz and Paek, containing many of their same key insights. However, there are at least two differences between their approach and mine.

First and foremost, game theory explicitly adopts a multiagent perspective. This is a key difference, because the core concepts of grounding theory (Section 4.4.1) are intrinsically multiagent in character. It is of course possible to bolt on aspects of multiagent reasoning to a decision-theoretic approach. For example, Paek and Horvitz (2000a) provide a brief explanation of how higher-order beliefs could be added to their model of grounding. However, game theory has already provided us with decades of work applying decision-theoretic principles to the multiagent case. This means there is no need to reinvent the wheel, and we can rely on this general work when constructing game-theoretic models of specific problem domains. For example, I was able to take advantage of the existing formal concept of common p-belief in Section 4.2.2 when formalizing the notion of common ground in game-theoretic terms.

A second key difference between the two approaches is that the game-theoretic model in this chapter takes a *sequential* approach to decision making. An extensive game models a sequence of decisions, not just a one-shot decision.[16] This is a good thing, since engaging in a dialog requires each participant to make a sequence of decisions as the dialog unfolds, where each decision impacts the utility of future decisions. It is natural, therefore, to adopt a sequential approach when formalizing decision making in dialog. Extensive games

[16]Were Horvitz and Paek to adopt a sequential decision-making perspective, they would be entering the domain of *Markov decision processes* (Puterman, 2005; Russell and Norvig, 2003). This approach to dialog has been explored in the work of Levin et al. (2000); Roy et al. (2000); Williams et al. (2005); Williams and Young (2007).

provide the means to do this. These two key differences between the approaches (single agent vs. multiagent perspective, and one-shot vs. sequential decision making) favor the adoption of game theory over decision theory as the appropriate means to formalize the core notions of grounding theory.

4.4.4. Objective and Normative Context

I conclude this section with a brief comparison to an interesting line of research on dialog agents initiated by Matthew Stone and David DeVault (Stone, 2005; Stone et al., 2006; DeVault and Stone, 2006, 2007; DeVault, 2008). Their research encompasses theoretically motivated computational approaches to natural language understanding, generation, and dialog management. Here I focus just on aspects of their proposals that deal most directly with conversational grounding and the representation of conversational state.

DeVault and Stone provide strong arguments against approaches to representing dialog context purely in terms of mutual or common belief (see especially DeVault and Stone (2006)). One of the most important objections raised against such approaches is that they have nothing to say about how interlocuters recover from misunderstandings. Misunderstandings arise frequently in conversation, yet interlocuters are also frequently able to overcome them, and it seems clear that they use some representation of their conversational state in order to do so in a targeted way. Yet if conversational state consisted solely of mutual beliefs, any misunderstanding would imply that conversational state holds no information about the topic on which the misunderstanding occurred!

As an alternative, DeVault and Stone propose that conversational state is an *objective* and *normative* product of the prior actions (both public and tacit) of the interlocuters.

Conversational state is objective, because there is a truth to the matter of what the state of a dialog is. It is also normative, because it is the task of an interlocuter to try to align her beliefs about the dialog with the true state of the dialog. DeVault and Stone (2006) compare this situation with a game of chess, where the moves that players make leave the game in a deterministic state, and where this state is separate from players' beliefs about it. Given the moves that have occurred in a particular chess game, it is the responsibility of the players to make sure they know the actual state of play, in order for them to stay within the bounds of the rules of chess. The state of the game doesn't change based on their beliefs about the game, it changes based on the moves that have actually occurred.

This view of conversational state does not rule out notions like mutual or common belief. Interlocuters certainly have beliefs about the actual state of the game, and if they coincide in the right way, these beliefs may also be mutual or common. DeVault and Stone stress, however, that these beliefs be kept separate from the actual conversational record. Keeping them distinct has several advantages, including allowing for the possibility of probabilistic beliefs *about* the conversational state (DeVault and Stone, 2006).

This general approach to dialog is fully consistent with the game-theoretic approach to dialog adopted in this chapter. An extensive game of imperfect information (Definition 3.2.1) provides an objective and normative representation of the state of play at any given point in an interaction among a set of players. Although the existence of chance nodes means that state is not a deterministic result of the players' actions, nevertheless the state of a game is always represented by a single node in the game tree. As we have seen, the players may have uncertainty about where they are in the tree, and this uncertainty is represented by probability distributions over their information sets. Consistent with the approach of DeVault and Stone, these information sets consist of the set of states that are possible at a given point in the game. Therefore, the game-theoretic approach to dialog advocated in this chapter falls squarely within the general paradigm that DeVault and Stone advocate.[17]

[17] In future work, it would be interesting to explore more deeply the connections between the line of thought developed by DeVault and Stone and the game-theoretic approach to dialog described in this chapter. In fact, DeVault and Stone (2007) themselves cite the development of decision-theoretic dialog strategies as a promising direction for future research.

CHAPTER 5

Experiment: Deriving Incrementality

Section 4.3 described how the game-theoretic model of grounding developed in this thesis could be applied to predict the incrementality of referential description installments, given values for parameters like error likelihood, turn cost, and task reward. Specifically, the model predicts that rational players will choose an optimal communicative action m^* that minimizes uncertainty, subject to the following constraint:

$$\frac{(C^* - C_{min})}{M + L} \leq (\epsilon_{max} - \epsilon^*)$$

Where M is the reward received by both players for successfully coordinating a correct action, L is the penalty for taking an incorrect action, C^* and ϵ^* are the cost and error likelihood associated with m^*, and C_{min} and ϵ_{max} are the cost and error likelihood associated with the least costly (and least precise) possible message. Given precise values for these parameters, this constraint makes precise predictions about installment frequency and length.

Even lacking precise values, however, we can still make general predictions about the behaviors of participants engaged in task-oriented dialog. First, increasing turn cost should result in fewer, less incremental (longer) installments. Second, increasing task success reward (M) or task error penalty (L) should result in more numerous, and more incremental (shorter) installments. Third, decreasing the probability that a given message

Table 5.1. Qualitative predictions for number of installment turns

	High Turn Cost		Low Turn Cost	
	High Error	Low Error	High Error	Low Error
High Task Cost	∗∗∗	∗∗	∗∗∗∗	∗∗∗
Low Task Cost	∗∗	∗	∗∗∗	∗∗

will be understood correctly will also result in more numerous, and more incremental installments. The qualitative predictions of manipulating these parameters are shown in Table 5.1, where relative number of installment turns is indicated by the number of asterisks. This chapter describes an experiment that was conducted in order to test some of these predictions.

In particular, the experiment was designed to explore the impact of turn cost and task error cost on the number and size of referential installments for participants engaged in a referential communication task. The experiment builds upon the body of work on referential communication tasks partially summarized in Section 2.2. Of this previous work, it is most closely related to and inspired by the puzzle task experiments of Gergle et al. (2004a,b,c). What differentiates the experiment described in this chapter from these previous studies is the manner in which it focuses directly on turn taking costs and task error costs. If it is correct to think of grounding behaviors in terms of costs and rewards, then it should be possible to manipulate them directly in order to derive different grounding behaviors in participants.

The referential communication task used in the experiment involved pairs of participants collaboratively solving a set of online jigsaw puzzles. For each puzzle, one of the participants was placed in the role of the helper, and the other was placed in the role of the worker. The helper had privileged access to the target puzzle solution, while the worker was responsible for selecting the target pieces from a candidate set of puzzle pieces and assembling them into the proper configuration. The asymmetric nature of the puzzle task roles forced the participants to engage in dialog in order to achieve a solution. The participants communicated with one another through a text-based medium similar to a chat system.

Turn cost was manipulated by introducing a communication delay between the participants. This was achieved by slowing down the rate at which text messages were displayed to the receiver after being sent by the sender. According to Clark and Brennan (1991), "delay costs" of this nature are one type of cost relevant to grounding behaviors, because delays (or the lack thereof) normally bear significance in dialog about a participant's degree of understanding, or desire to take dialog initiative. The manner in which communication delay was implemented in the experiment ensured that the delay cost increased in direct proportion to the number turns taken (see Section 5.1 for details). Therefore, according to the model described in Section 4.3, this turn delay cost should result in fewer turns and longer referential installments. This prediction is reflected in Table 5.1, comparing its two rightmost columns to its two leftmost columns.

In addition to a delay cost, the high turn cost condition also imposed what Clark and Brennan (1991) call an "asynchrony cost". In normal speech-based dialog, participants are able to time their utterances with great precision. For example, acknowledgements can be timed to correspond to a particular sub-part of an interlocuter's utterance (Jefferson, 1973) or physical action (Clark and Krych, 2004; Gergle, 2006). However, in the high turn cost condition of the experiment, this ability to precisely time the sending of messages was greatly reduced, because the message sender was forced to wait for the immediately preceding message to be completely viewed by the recipient before entering and sending another message (again, see Section 5.1 for details).

The model of incremental referring expressions developed in Section 4.3 only mentions a single turn cost function. This single turn cost function is to be thought of as the result of composing a set of component cost functions, one for each source of turn cost. For this experiment, I focus only on the turn costs contributed by manipulating delay and asynchrony. What matters here is that each of these components contributes to the cost of an individual turn, and this combined cost monotonically increases with the number of turns. There may be other potential sources of turn cost in the experiment (such as the lexical complexity involved in describing a particular puzzle piece), but these are not directly manipulated or accounted for.

Task cost was manipulated by introducing a penalty for puzzle piece selection errors. While solving the puzzle, the participant playing the worker role was responsible for both selecting the target puzzle pieces from a set of alternatives, and arranging them into the target configuration. In the high task error condition, selecting an incorrect piece caused the puzzle to reset to its initial state. A selection error therefore forced the participants to redo the work they had already done in order to solve the puzzle. According to the model described in Section 4.3, this high task error cost should result in more frequent turns and shorter referential installments. This is because the model assumes an error function where longer installments (and fewer turns) increases the likelihood that the receiver will be unable to recover its intended meaning. This prediction is reflected in Table 5.1, comparing the bottom row to the top row of the table. The following section explains how the predictions regarding task error cost and turn cost were operationalized in the experiment.

5.1. Method

5.1.1. Participants

Participants consisted of 12 pairs of Northwestern University students, who were paid $12 each for their participation in the experiment. All participants were native speakers of English, and proficient at entering text by keyboard onto a display screen.

5.1.2. Materials

Four online puzzles were created for this study. Each puzzle consisted of a set of twelve puzzle pieces and a target solution. Each target solution consisted of three of the puzzle's twelve pieces, arranged into a connected 2-dimensional configuration. Therefore, of the twelve puzzle pieces, three were relevant to the solution and nine were distractors. An example set of twelve puzzle pieces is shown in Figure 5.1, and a corresponding solution is shown in Figure 5.2.

As shown in Figure 5.1, each puzzle piece was composed of a set of five simpler shapes (circles, triangles, and squares) coming in a variety of colors (red, blue, yellow and green). Additionally, each component shape of the puzzle piece had an internal "decoration", e.g. a set of small black triangles, stars, or circles. The design of these stimuli was intended to elicit complex referring expressions from participants, with each referring expression consisting of a sequence of puzzle piece property values. For example, one

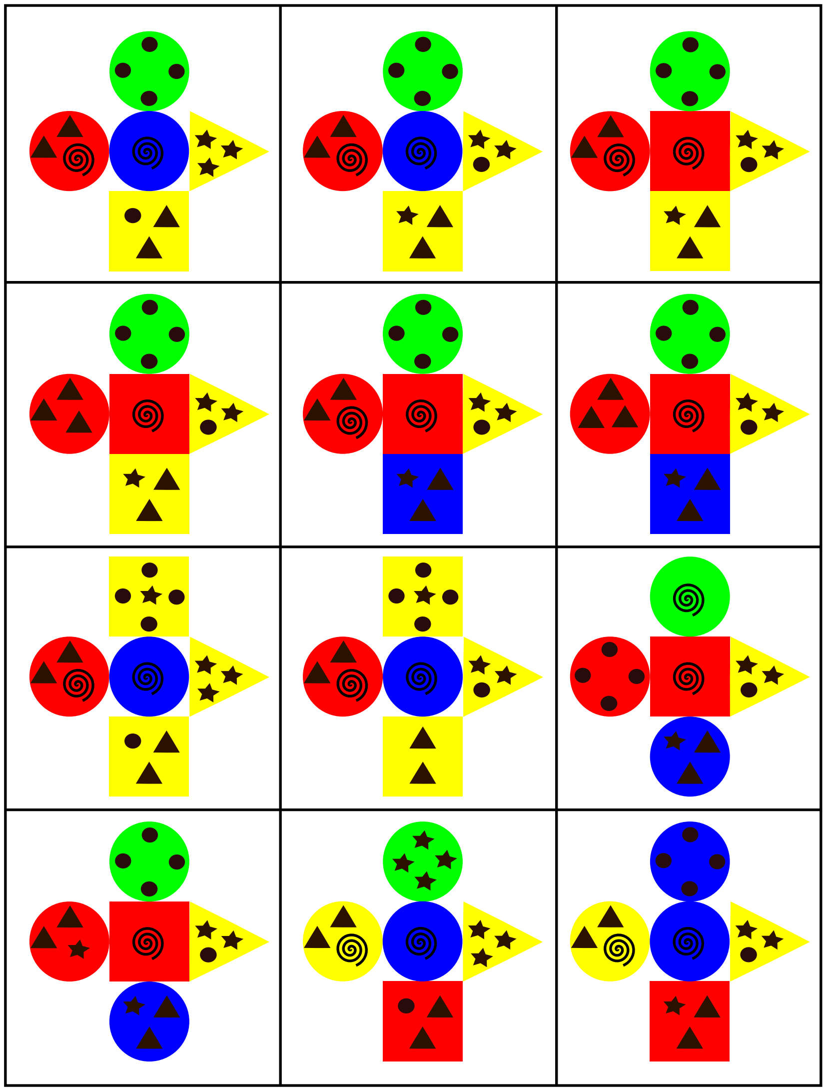

Figure 5.1. Puzzle piece stimuli

way of describing the top left puzzle piece in Figure 5.1 is as follows (where **H** signifies "helper"):

(5.1) **H:** A center blue circle containing a spiral, with a green circle on top containing 4 little circles, a red circle on the left with 2 triangles and a

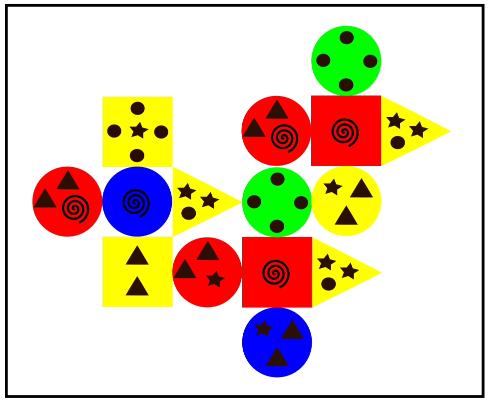

Figure 5.2. Puzzle solution

spiral, a yellow square on the bottom with 2 triangles and a circle, and a yellow triangle on the right containing 3 stars.

This example packages the entire description in a single installment turn. However, many other possibilities exist. For example, the following example packages the exact same description into five separate installments, each of which may or may not be followed by a response from the worker:

(5.2) **H:** A center blue circle containing a spiral

H: with a green circle on top containing 4 little circles

H: a red circle on the left with 2 triangles and a spiral

H: a yellow square on the bottom with 2 triangles and a circle,

H: and a yellow triangle on the right containing 3 stars.

From among a set of twelve pieces, each puzzle piece was uniquely identified by a particular set of property values of this nature. In general, each puzzle piece required more than a single property value in order for it to be uniquely distinguished from among the set of alternatives.

5.1.3. Interface

The worker and the helper each had a separate view of the evolving state of the puzzle.[1] An example worker's view is shown in Figure 5.3. To the left of the worker's view was the puzzle piece *selection area*, containing the puzzle's twelve component pieces in a 3×4 grid. In the middle of the worker's view was the *workspace area*. Clicking on a puzzle piece in the selection area caused a corresponding puzzle piece to appear in the top left corner of the workspace, from where it could be dragged by the worker to its proper location. When a puzzle piece was selected in this manner, all of the remaining puzzle pieces in the selection area were disabled until the selected piece was positioned and set in its final location. The worker did this by dragging the piece to the appropriate location and then double-clicking on it. This action brought up a dialog box giving the worker the final choice to set the piece in its current location.

On the right of the worker's view was the *target area*. For the worker this area was empty, displaying nothing until the last puzzle piece was added to the workspace and fixed in its final location. At this point, the target area displayed the puzzle solution to the worker, giving him the chance to see how closely the arrangement of his workspace pieces matched the configuration of the corresponding pieces in the solution. Also at this

[1]The participants in the study where told that their roles were "director" and "builder", rather than "helper" and "worker". I felt these alternative terms would have a slightly more positive connotation for the participants. The terms "helper" and "worker" are used in the description of the experiment in this chapter, however, to be consistent with the terminology of Gergle et al. (2004a,b,c).

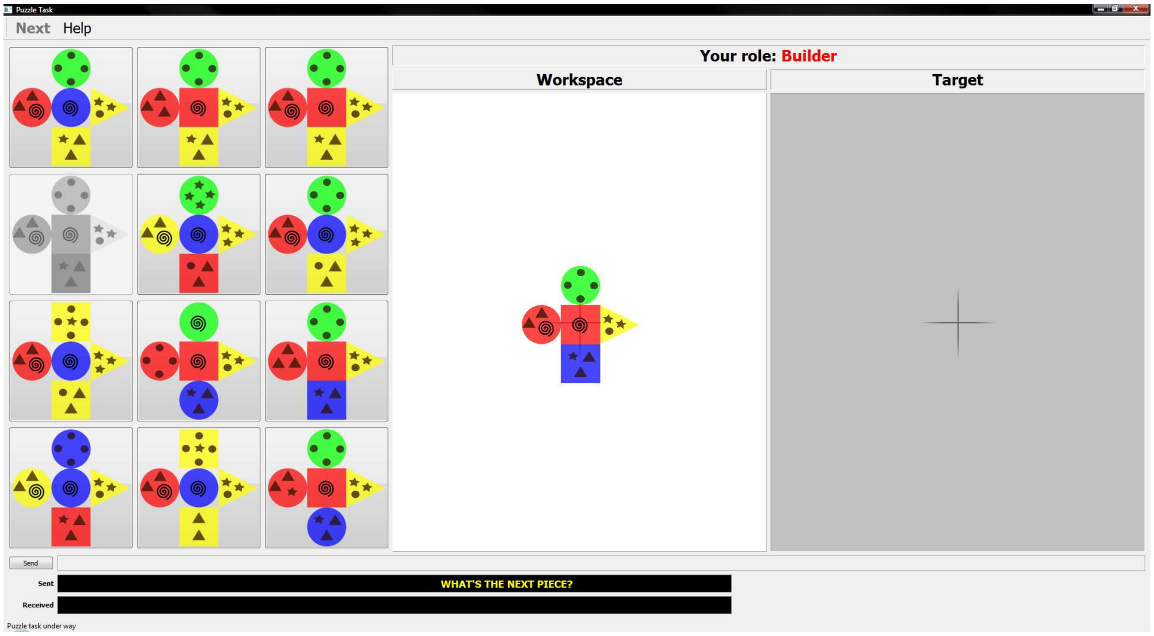

Figure 5.3. Worker interface view

point, the `Next` menu item at the top left of the screen was enabled, and a message was displayed to the worker informing him to click on it in order to start the next puzzle.

An example helper's view is shown in Figure 5.4. The helper's view contained the same components as the worker's view, but these components had different properties. The helper's view of the target area displayed the solution to the current puzzle being worked on. For the helper, the selection area was disabled, and the identity of each piece was hidden. In addition, the helper had no control over the workspace state. When the worker added a piece to the workspace, a corresponding gray box with a question mark was placed in the helper's view of the workspace. As the worker moved the piece, the box in the helper view mirrored its movements. Only when the worker set the piece in its final position was the actual identity of the piece revealed to the helper. The purpose of

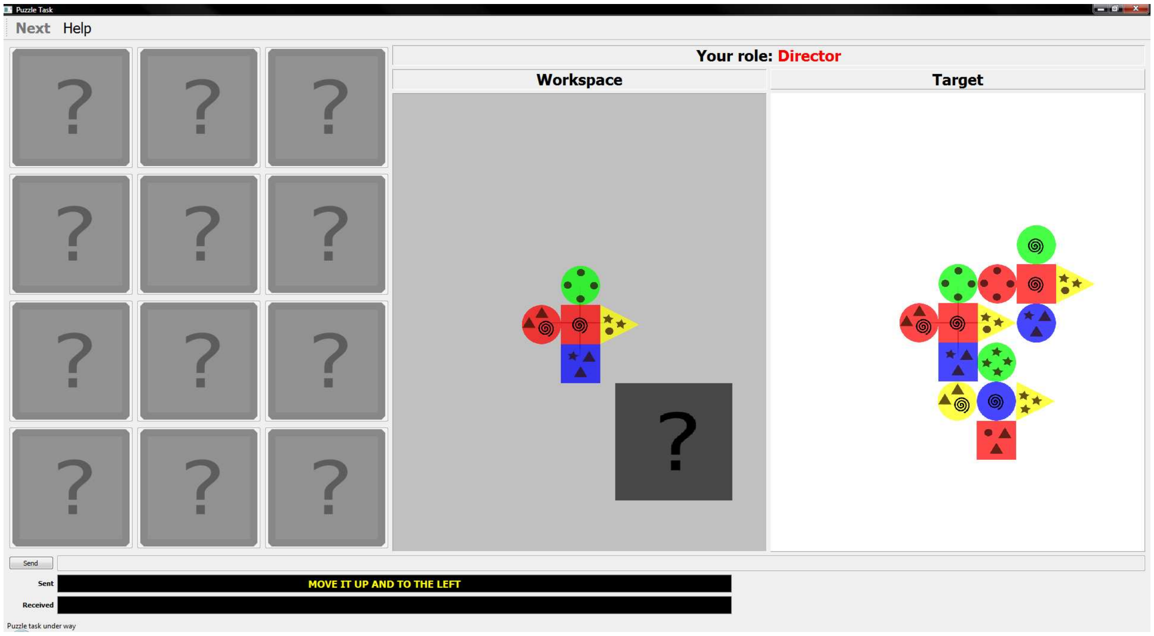

Figure 5.4. Helper interface view

this device was to prevent the helper and worker from systematically adding pieces to the workspace in a game of elimination, where an incorrect or correct piece could be simply identified as such by the helper when the piece was added to the workspace. Figure 5.4 shows a snapshot of a helper view where one piece has been added to the workspace and fixed in its final location, and another piece has been added but not yet fixed in its final location.

The text-based communication interface appeared at the bottom of both the helper's and the worker's views, and was identical in both views. The interface consisted of three boxes, which I refer to as "lines". The top line was the text input line, where the helper or worker entered text using the keyboard. The text input line extended across the length of the view, and text of arbitrary length could be entered into it. Text in the text input line line was sent to the other player by clicking on a `Send` button, or alternatively by hitting the `Enter` button on the keyboard. The second and third lines of the interface were approximately two-thirds the length of the first line. The second line was used to display to the sender the text of the messages that were sent – its intended purpose was to show the message sender a real-time view of how her message was being viewed by the other player. The third line of the interface displayed messages received from the other player.

Messages were displayed on the screen as moving text, scrolling from left to right at a constant rate. The scrolling text provided enough time for the receiver of the message to read and process the message, but not so much time that a permanent artifact was available for the receiver to rely upon while working on the task. This was crucial to the experiment, because without the ability to temporally bound the visibility of a message, it would have been very difficult to elicit incremental referring expressions.

The interface was implemented in a way that allowed for the turn cost and task error cost manipulations discussed at the beginning of this chapter. The turn delay cost was manipulated by changing the scrolling speed of the text in the chat interface. For the low turn cost condition, the text scrolled at the rate of one pixel every five milliseconds. This rate produced scrolling text that was legible, but moving fairly rapidly. Given the physical length of the display lines on the computer screens of the laptops that were used, one character took approximately three seconds to travel from the right-most edge to the left-most edge. For the high turn cost condition, the scrolling speed was one-third as fast: text scrolled at the rate of one pixel every 15 milliseconds. In this condition, one character took approximately nine seconds to travel across the display lines.

The turn asynchrony cost was manipulated by enabling or disabling the text input line after a subject sent a message. In the low turn cost condition, the text input line was always enabled, and the subject could enter and send messages at whatever time she chose. In the high turn cost condition, after sending a message the text input line was disabled until all of the text from the current sent message was finished scrolling across the screen. This forced the message sender to wait until the current message was completely done scrolling before she could enter and send another message. The result of this was that cumulative delay scaled proportionately to the number of turns that were taken; every message that was sent and displayed had a fixed lag time beginning immediately after the last word entered the display area on the right-hand side, and ending at the point where the last word exited on the left-hand side.

Task error cost was manipulated by implementing two different responses to worker puzzle piece selection errors. In the low task error cost condition, the penalty for making a selection error was that the incorrect piece was removed from the workspace. This penalty was invoked after the incorrect piece had been added to the workspace, and the worker double-clicked on it order to set it in its final location. Responding `Yes` to the dialog box that popped up at that point resulted in the incorrect puzzle piece being removed from the workspace and placed back in the selection area. Additionally, all puzzle pieces in the selection area were scrambled to random new locations in the 3×4 grid, and a message was displayed to both the worker and the helper informing them that a mistake had been made. In the high task error condition this same penalty was executed, but in addition all other puzzle pieces in the workspace were removed and put back in the selection area.

Therefore, in the high task error condition, a puzzle piece selection error always resulted in the entire puzzle starting over from scratch.

5.1.4. Procedure

The participants were placed in front of separate laptop computers in the same room, with a divider between them preventing visual contact with each other. An ethernet cable connected the two laptops to allow them to communicate with each other via TCP sockets. The participants were instructed to refrain from speaking to each other, and to use the text interface as their sole means of communication.

Before beginning the experiment, each subject was given written instructions explaining the goal of the task, and how to use the interface. For the high task error condition, the instructions informed them that a piece selection error would cause the current puzzle to start over from scratch, negating the work done so far on that puzzle. For the low task error condition, the instructions informed them that a piece selection error would result only in the incorrect piece being removed from the workspace, and all other pieces in the workspace would remain in their positions on the workspace.

After reading the instructions, the participants engaged in a brief practice session consisting of two puzzles. The first puzzle demonstrated the low turn cost condition, with the faster text scrolling speed. The second puzzle demonstrated the high turn cost condition, with the slower scrolling speed. The participants switched roles during the practice, so each of them experienced the helper role and the worker role before the actual trials began. During the practice, the players were told to intentionally make a selection error, in order to experience first hand the high task error or low task error consequences that resulted, depending on which condition they were in. The participants were also told to communicate using the text interface in any way they saw fit in order to solve the puzzles.

After the practice, one of the participants was randomly chosen to start out as the helper, and the other as the worker. The roles switched for every puzzle, so by the end of the session each subject played the helper role twice and the worker role twice. The helper's task in each trial was to generate referring expressions that uniquely identified the set of three target pieces. The helper had the opportunity to package these property values into a single installment turn, or sequence of installment turns, in any way she chose. She was free to generate a referential description packaging all of the property values into a single installment that uniquely identified a particular piece. Alternatively, she was free to package them into multiple, more incremental installments, whose sum total would uniquely identify the puzzle piece. This choice was left entirely up to the participants. Of interest was how the turn cost and task error cost manipulations would influence the decisions that the helpers made with respect to this choice.

5.1.5. Design

Each pair of participants solved four different puzzles, with the order of the four puzzles randomly determined. The pairs alternated taking turns as the helper and the worker, performing each role twice. Task error cost was manipulated between pairs of participants, with six pairs solving puzzles in the low task error condition, and six in the high task error condition. Turn cost was manipulated within each pair, with each pair solving two puzzles in the low turn cost condition, and two puzzles in the high turn cost condition.

The order of low and high turn costs puzzles was counter-balanced across the 12 pairs: six of the pairs started with two high turn cost puzzles, and six of the pairs started with two low turn cost puzzles. This resulted in a 2 (turn cost condition: delay vs. no-delay) $\times$ 2 (task error cost condition: puzzle restart vs. no-restart) mixed factorial design, with turn cost a within-pairs factor, and task error cost a between-pairs factor.

5.2. Results

All communication between the players was recorded in log files, with timestamps identifying the precise time each message was sent. Other significant events were also recorded in the logs, including when a puzzle piece was selected and added to the workspace, the locations to which it was moved, whether any selection errors were made, and how long each trial took. These logs form the basis for the analyses performed in this section. The dependent variables that are of greatest interest here are the number of helper messages that were sent, and the length of these messages, measured in number of words. As discussed at the beginning of this chapter, the model makes clear predictions regarding the relative outcomes for these variables, given manipulations to turn cost and task error cost.

Before moving on to analyses of these outcomes regarding message frequency and length, I begin by giving a high-level look at how much effort was expended by subject pairs in order to solve the puzzles. One useful metric to get a rough sense of puzzle difficulty is the amount of time it took to solve the puzzles. Figure 5.5 shows the mean trial durations across the four conditions. Duration was measured from when a puzzle was loaded on the participants' computers, to the time the puzzle's last correct puzzle piece was set in its final location in the worker's workspace. Not surprisingly, the graph

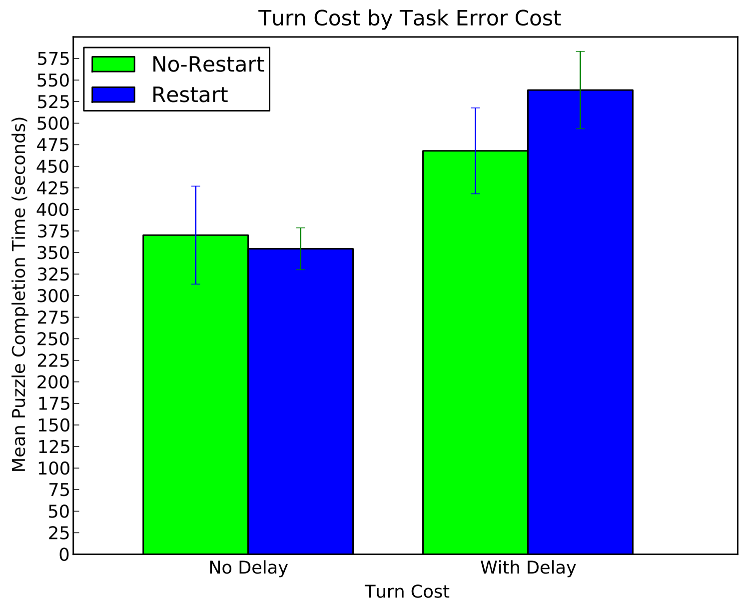

Figure 5.5. Effect of turn cost and task error cost on puzzle completion time. Error bars represent standard error of the participant mean.

shows that introducing a turn delay resulted in an increase in time to completion: for the no-delay condition, it took participants on average about six minutes to finish an individual puzzle, while for the delay condition it took on average closer to eight and a half minutes.

Another metric for measuring how difficult the puzzles were is to look at the number of puzzle piece selection errors that were made. According to this metric, the puzzles were not very difficult, as participants made relatively few puzzle piece selection errors during the experiment. Across all trials, there were only 16 piece selection errors (11 of which were in the puzzle restart condition), compared with a total of 144 correct target puzzle pieces. Additionally, the errors were not distributed evenly: four of the pairs made no errors at all, and two of the pairs made four each (half of the total number). To preview subsequent discussion, the task error cost manipulation failed to provide statistically significant support for the predictions of the model. It may be the case that this can be attributed to the fact that the puzzles weren't challenging enough to make the threat of an error relevant, as indicated by the paucity of errors that were actually made.

The remainder of this section examines results regarding the number and length of the messages that were sent by the participants. Each of the following statistical analyses is a two-way 2×2 mixed ANOVA in which turn cost (no-delay, delay) was a repeated factor, and task error cost (no-restart, restart) was a between-pair factor. Results for the helper and worker roles are reported separately, with a focus on the results for the helper role, since the helper was the one responsible for generating uniquely identifying referential descriptions for each target puzzle piece.

Figure 5.6 shows the number of helper messages across the four conditions. There was a significant main effect of turn cost on the number of helper messages, $F(1, 10) = 49.70$, $p < .001$. There were over twice as many helper messages sent in the no-delay conditions (total = 493) as in the delay conditions (total = 232). This effect tells us that adding the turn delay cost greatly reduced the number of messages that were sent, supporting a central prediction of the model. However, the main effect of task error cost on the number of helper messages was not significant, $F(1, 10) = .01$, $p > .05$. The interaction effect of turn cost and task error cost was also not significant $F(1, 10) = .61$, $p > .05$.

Since helpers were using fewer messages in the turn delay condition, this implies that they were packaging more content on average into each individual message, because the amount of information required to identify a puzzle piece was roughly constant across the

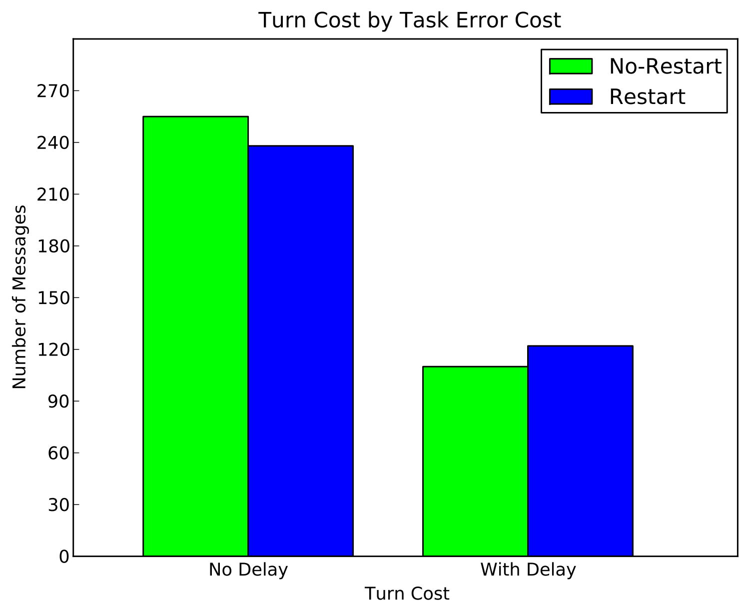

Figure 5.6. Effect of turn cost and task error cost on number of helper messages.

two conditions. To examine this possibility, I took the number of words in a message as a proxy for the amount of information it encodes, with words defined as blocks of text (tokens) separated by whitespace. Using this metric, Figure 5.7 graphs the mean message lengths of helper messages sent across the four conditions. There was a significant main effect of turn cost on the length of helper messages, $F(1, 10) = 34.53$, $p < .001$, indicating that the mean length of helper messages was significantly greater for the delay condition ($M = 17.60$) than for the no-delay condition ($M = 8.40$). In fact, the total number of helper words in the delay condition (total = 3989) was virtually identical to the total number of helper words in the no-delay condition (total = 3926). The difference between the two conditions was how these words were packaged into individual installments. Once

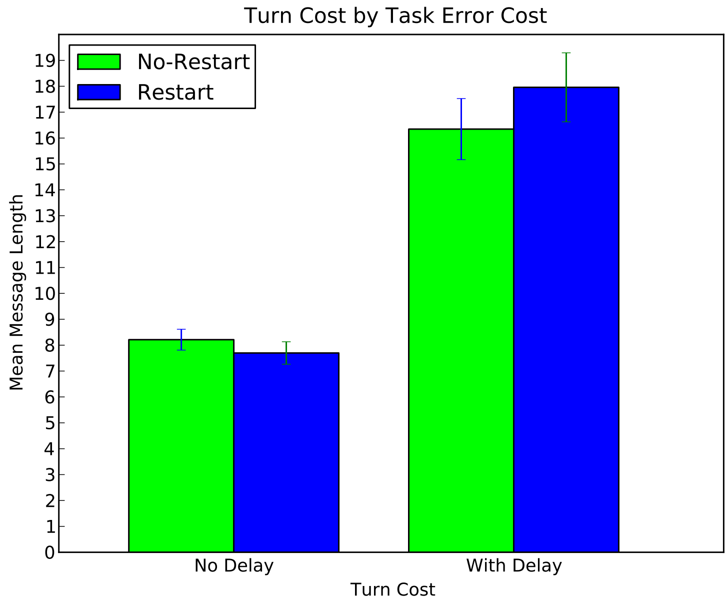

Figure 5.7. Effect of turn cost and task error cost on helper message length.

again however, the main effect of task error cost was not significant, $F(1, 10) = .41$, $p > .05$. Neither was the interaction of turn cost and task error cost, $F(1, 10) = .62$, $p > .05$.

Figure 5.8 shows a dialog excerpt of a pair in the no-delay condition, working together to add a puzzle piece to the workspace. The initial part of the excerpt illustrates the kind of incremental referring expressions that were generated in this condition: the first five messages are from the helper, and consist of a single referential description split into five installments. This sequence of five installments is followed by two brief exchanges initiated by the worker, in order to elicit more information from the helper regarding the target description. Once the worker is satisfied that he has identified the relevant puzzle

No Turn Delay

H: 1ST PIECE: RED SQUARE WITH SPIRAL IN CENTER
H: IT HAS A GREEN CIRCLE ABOVE
H: BLUE SQUARE BELOW
H: RED CIRCLE TO THE LEFT
H: YELLOW TRIANGLE TO THE RIGHT
W: WHATS IN THE RED CIRCLE?
H: 2 TRIANGLES AND A SPIRAL
W: IS THIS THE CENTER
H: YUP
W: [ADD PIECE TO WORKSPACE]
W: [MOVE PIECE]
W: [SET PIECE]
H: OK, NEXT PIECE: BLUE CIRCLE IN THE MIDDLE WITH SPIRAL
H: IT HAS A GREEN CIRCLE ABOVE

Figure 5.8. Excerpt from the no-delay condition

piece, he adds it to the workspace, moves it to where he believes it's target position is, and sets it in place. The helper then moves on to describe the next target puzzle piece.

Figure 5.9 shows a dialog excerpt from a pair working in the turn delay condition. In contrast with the excerpt shown in Figure 5.8, the referential descriptions generated by the helper are packaged into single, lengthy installments. In this excerpt, it is only after the target puzzle piece has been selected and added to the workspace that the worker sends a message. This was typical of dialog in the turn delay condition: it tended to be far less "interactive" than dialog in the no-delay condition, where the helper and the worker tended to send more frequent, overlapping, and fine-grained messages.

Turning attention to the worker role, Figure 5.10 shows the number of worker messages across the four conditions. When compared with the analogous helper numbers in Figure 5.6, it is clear that workers on average sent far fewer messages than helpers. This was

Turn Delay

H: SO THE ONE THATS ON THE CROSS HAS A BLUE SQUARE IN THE MIDDLE WITH THREE BARS. AND TO THE TOP IS A GREEN CIRCLE WITH FIVE LITTLE CIRCLES AND THE BOTTOM IS A GREEN TRIANGLE WITH THREE TRIANGLES AND TO THE RIGHT IS A TRIANGLE WITH THREE STARS.
W: [ADD PIECE TO WORKSPACE]
W: [MOVE PIECE]
W: GOOD?
H: DOWN A LITTLE BIT
W: [MOVE PIECE]
H: YEAH THAT'S GOOD
W: [SET PIECE]
W: OK NEXT?
H: OKAY SO THE NEXT PIECE HAS A RED SQUARE IN THE CENTER WITH THREE BARS. THE RIGHT IS A YELLOW TRIANGLE WITH FOUR STARS. THE LEFT IS A BLUE SQUARE WITH FOUR STARS AND THE BOTTOM IS A GREEN TRIANGLE WITH THREE TRIANGLES.
W: [ADD PIECE TO WORKSPACE]

Figure 5.9. Excerpt from the delay condition

expected given the asymmetries between the two roles. However, there was still a significant main effect of turn cost on the number of worker messages, $F(1, 10) = 7.43$, $p < .05$, with a total of 197 messages in the no-delay condition, and 135 in the delay condition. The main effect of task error cost on the number of worker messages was not significant, $F(1, 10) = 2.49$, $p = .15$, though there was a trend towards more messages in the restart condition ($M = 7.83$) than in the no-restart condition ($M = 6.00$). The interaction effect of turn cost and task error cost was not significant, $F(1, 10) = .94$, $p > .05$.

Figure 5.11 graphs the mean message lengths of worker messages sent across the four conditions. There was a significant main effect of turn cost on the length of worker messages, $F(1, 10) = 5.40$, $p < .05$, indicating that the mean length of worker messages was significantly greater for the delay condition ($M = 5.17$) than for the no-delay condition

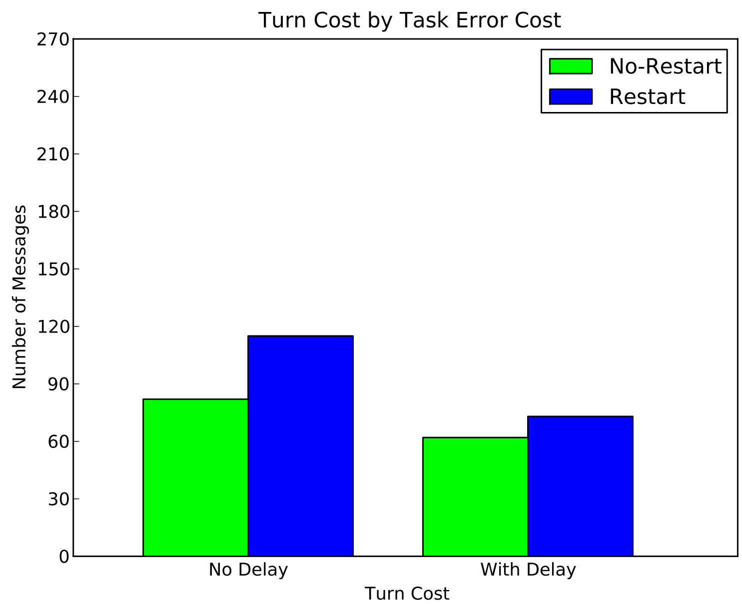

Figure 5.10. Effect of turn cost and task error cost on number of worker messages

($M = 3.84$). This shows that increasing turn cost significantly increased the mean length of worker messages, despite the relative shortness of worker messages in general when compared with helper messages. The main effect of task error cost on the length of worker messages was not significant, $F(1, 10) = 4.42$, $p = .06$, though there was a strong trend towards longer messages in the no-restart condition ($M = 5.25$) than in the restart condition ($M = 3.69$). The interaction effect of turn cost and task error cost was not significant, $F(1, 10) = 0$, $p > .05$.

5.3. Discussion

The experiment showed that the turn cost manipulation had a strong effect on the number and size of the messages that were sent. Both the helpers and the workers were

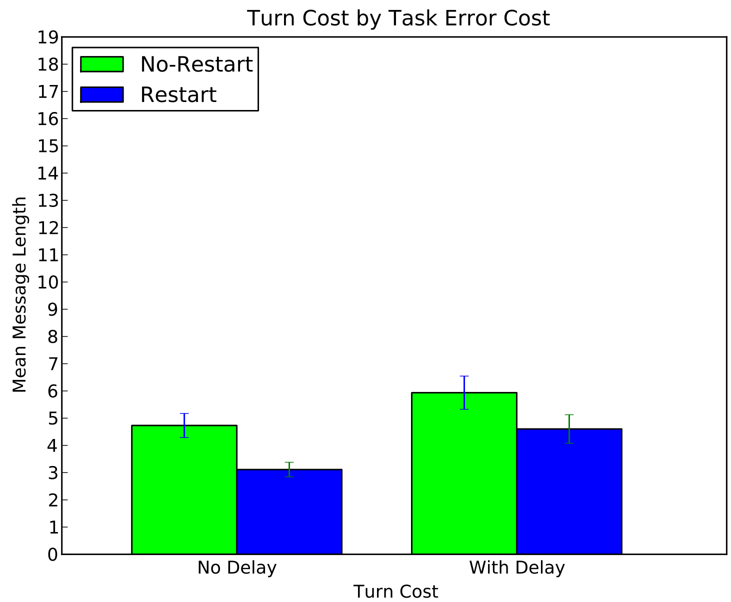

Figure 5.11. Effect of turn cost and task error cost on worker message length

affected by this manipulation. There were approximately twice as many helper messages in the no-delay condition than in the delay condition. However, in both conditions the amount of information that had to be transmitted to the worker was roughly constant. This implies that helpers were packaging more content on average into each individual message in the delay condition than in the no-delay condition, and this is what we found: messages in the delay condition were approximately twice as long as in the no-delay condition. In fact, the total numbers of helper words generated across the two conditions were virtually identical: for the no-delay condition, helpers generated 3926 words, while in the delay condition they generated 3989 words. Just as the model predicts, the difference between the two conditions lay in how these words were packaged into installments.

Turn cost affected not just helper messages, but worker messages as well. There were about 30% fewer worker messages sent in the delay condition than in the no-delay condition, and about 25% more words in the delay condition than in the no-delay condition. Again, the total numbers of words generated across the two conditions were very close: 746 in the no-delay condition, and 704 in the delay condition. Unsurprisingly given the asymmetry in responsibilities between the worker and helper roles, workers generated far fewer messages and words than helpers. It was the helpers that were responsible for generating uniquely identifying referential descriptions for each target puzzle piece, and for giving instructions on where to place each puzzle piece in the workspace. The workers' main responsibility (with respect to communication) was to ask for clarifications and provide acknowledgements.

An additional factor for why workers produced so many fewer messages than helpers was the shared visual workspace, which made it possible for the worker to substitute actions for words (Clark and Krych, 2004; Gergle et al., 2004a,b; Gergle, 2006). Figure 5.12 shows a portion of dialog where the helper is giving directions to the worker on where to place a puzzle piece, after it has been selected and added to the workspace. Several times in this excerpt, the worker moves the piece without sending a corresponding text message. The helper is able to use the worker's actions to generate an appropriate response, correcting or confirming the worker's placement of the puzzle piece. When the worker does send a message, he is able to use the fact that the helper is observing his movements by using deictics ("There?" and "Like that?"), further reducing the total number of words the worker requires in order to communicate with the helper (cf., Gergle et al., 2004b).

No Turn Delay

W: GOT THE PIECE
W: [MOVE PIECE]
H: THE TOP RIGHT CORNER OF THE YELLOW SQUARE SHOULD LINE UP WITH THE BOTTOM LEFT CORNER OF THE MIDDLE PIECES YELLOW SQUARE
W: [MOVE PIECE]
W: WAIT, AGAIN?
H: YEA, KEEP MOVING IT DOWN
W: [MOVE PIECE]
H: THE TOP RIGHT CORNER OF THIS PIECE
W: [MOVE PIECE]
H: SHOULD MEET THE BOTTOM LEFT CORNER OF THE MIDDLE PIECE
W:THERE?
H: A LITTLE LOWER I THINK
W: [MOVE PIECE]
H: THAT LOOKS GOOD
W: LIKE THAT?
W: [SET-PIECE]

Figure 5.12. Excerpt showing dialog after puzzle piece selection

Returning to the main thread of discussion, there were probably two causes for why the turn cost manipulation had such a strong effect, particularly on helper messages. First, the difference in scrolling speed across the two conditions was quite dramatic: scrolling speed in the delay condition was one third as fast as in the no-delay condition – this meant that scrolling in the delay condition was almost painfully slow. Combined with the fact that the text input line was disabled while the text scrolled across the display line, helpers were forced to sit and wait for a considerable amount of time after their messages were sent before they received any feedback from the worker. In other words, the high turn cost trials were *much* more costly than the low turn cost trials. The excerpt from Figure 5.13 shows an example of how participants explicitly reacted to the delays imposed in the delay condition. The helper begins the excerpt by telling the worker that she will

Turn Delay

H: IM GOING TO SAY THEM ALL AT ONCE SINCE ITS SO SLOW...SPEED THIS THING UP A LITTLE BIT
W: YEAH GOOD CALL
H: PIECE OVER THE CENTER CROSS: RED SQUARE(CURLY-QUE IN THE MIDDLE), GREEN CIRCLE(4 DOTS), YELLOW TRIANGLE (2 STARS, 1 DOT), BLUE CICRLE (2 TRIANGLE, 1 STAR), RED CICRLE (2 TRIANGLE, 1 STAR)
W: OKAY I'M GOING FOR IT
W: [ADD PIECE TO WORKSPACE]
W: [MOVE PIECE]
W: [SET PIECE]
H: BLUE CIRLCE(CURLY-QUE), YELLOW SQUARE(4 DOTS, 1 STAR), YELLOW TRIANGLE (2 STARS, 1 DOT), YELLOW SQUARE(2 TRIANGLE), RED CIRCLE(2 TRIANGLE, 1 C-Q)
W: GOT IT
W: [ADD PIECE TO WORKSPACE]

Figure 5.13. Excerpt from the delay condition

package descriptions into single installments in order to speed things up, and she follows this plan in the remaining portion of the transcript.

A second cause for why the cost manipulation had such a dramatic effect is that although the slower scrolling speed in the delay condition made turns more costly, it also had a positive side-effect: it allowed participants to take more time to read and process messages in a more leisurely fashion. Workers could read portions of the helper message, switch focus to the selection area to identify candidate puzzle pieces that satisfied the description so far, and then return attention to the helper message that was scrolling across the screen, without worrying that the helper message will have disappeared in the mean time. Evidence for this claim comes from observing participants' behavior during the experiment, from verbal reports they supplied afterward, and from some exchanges recorded in the transcripts. Figure 5.14 provides one example, from a trial in the no-delay

No Turn Delay

H: THE PIECE ON THE LEFT HAS A RED SQUARE WITH SPIRAL IN THE CENTER, A RED CIRCLE ON THE LEFT, A GREEN CIRCLE ABOVE, A YELLOW CIRCLE UNDER, AND A YELLOW TRIANGLE TO THE RIGHT
W: IT IS MOVING REAL FAST SO THAT WAS HARD TO READ ALL OF THAT
W: THE LEFT PIECE IS RED WITH WHAT?
H: THE LEFT PIECE HAS A RED SQUARE AT THE CENTER. TO THE LEFT OF THE SQUARE IS A RED CIRCLE
H: ABOVE THE SQUARE IS A GREEN CIRCLE
H: BELOW THE SQUARE IS A YELLOW CIRCLE
H: TO THE RIGHT OF THE SQUARE IS A YELLOW TRIANGLE
W: [ADDS PIECE TO WORKSPACE]

Figure 5.14. Excerpt from the no-delay condition

condition. The helper begins the excerpt with a lengthy description of a target puzzle piece. The worker responds to this message by informing the helper that the message was scrolling too fast for him to grasp its content. The helper responds by breaking up the original lengthy description into a sequence of four installments, which the worker then accepts by adding the appropriate puzzle piece to the workspace.

While the delay and no-delay conditions had a very strong effect on how helpers and workers packaged information into messages, the task error condition (puzzle restart vs. no-restart) had no significant effect. This may have been due to the fact that (as mentioned in Section 5.2) participants made relatively few puzzle piece selection errors during the experiment. Across all trials, there were only 16 piece selection errors, and the errors were distributed unevenly: four of the pairs made no errors at all, and two of the pairs made four each (half of the total number). Given the paucity of errors that were made during the experiment, it is plausible to suppose that high task error cost condition was ineffectual because the participants didn't feel the danger of making an error to be strongly relevant to their actions.

Another possible reason why the task error manipulation failed to have a significant effect may been the small size of the target puzzle solutions, which consisted of only three pieces. Since there were only three pieces in the solution, the low task error condition and the high task error condition may have been too similar to each other to elicit significantly different behaviors. The two conditions were in fact identical to each other for the first puzzle piece selected, since an error at this stage resulted in a de facto restart of the puzzle in both. Future work should address these issues by increasing the difference between the low and high task error cost conditions, and by making errors more likely to happen. These changes should make the task error condition more relevant to the participants' actions, providing a better test of the model's predictions.

Although the experiment was unable to generate a statistically significant result for manipulating task error cost, there was a trend towards significance in the case of the worker's messages, both in terms of frequency and length. There was a tendency ($p = .15$) towards more worker messages in the restart condition, and an even stronger tendency ($p = .06$) for shorter messages in the restart condition. If these tendencies are indicative of a real effect of task error cost on worker message frequency and length, then a possible explanation is that workers were being more careful about confirming their understanding of helpers' messages in the high task error cost condition. This result would be consistent with the model, which predicts that participants will put up with a higher cumulative turn cost in order to achieve greater certainty, in the face of higher task error cost. Exploring this possibility is a goal for future work.

5.3.1. Connecting the Results to the Model

I conclude this chapter by exploring in more detail the connection of these experimental results with the predictions of the game-theoretic model of grounding in Section 4.3. First, look again at the qualitative representation of the model's predictions as represented in Table 5.1. The turn cost prediction of the model was borne out by the experiment: high turn cost (the table's leftmost two columns) resulted in fewer installment turns, while low turn cost (the table's rightmost two columns) resulted in more frequent installment turns. On the other hand, the task error cost prediction of the model was not borne out by the experiment, with no significant contrast between the high task error cost (top row) and low task error cost (bottom row) conditions. As discussed above, this lack of an effect may plausibly be attributed to the fact that the low and high task error cost conditions ended up being too similar to each other.

Examining Table 5.1 in more detail, it shows that the predictions are subdivided further across high and low error likelihood conditions. This is relevant to the results obtained in the experiment, because (as discussed above) it appears that the turn cost manipulation had a side effect on the intelligibility of the messages that participants sent to one another. The faster scrolling speeds made it more difficult to process longer messages, because it reduced the time that participants had to process them. Conversely, the slower scrolling speeds made it easier to process longer messages, because the participants could take more time to process them. In fact, it appears that changing the scrolling speed had effects on both turn cost and error likelihood that mutually reinforced one another to produce a dramatic difference in installment frequency and length. This mutual reinforcement is completely consistent with the predictions represented in Table

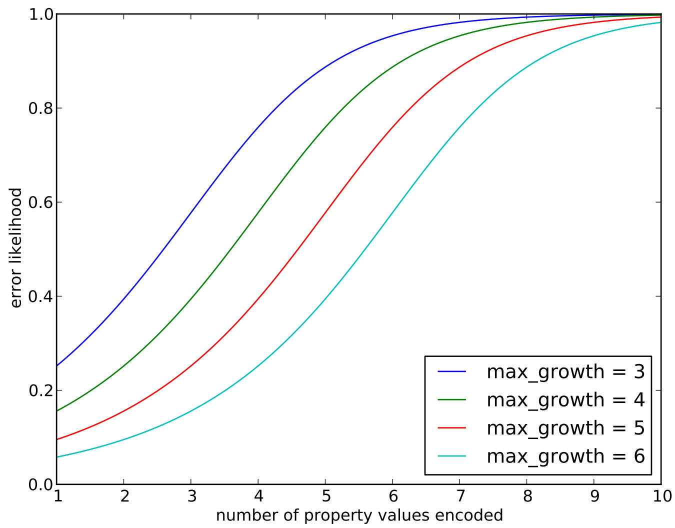

Figure 5.15. Error model with selected locations of maximum growth

5.1. Looking across the four columns, the biggest difference in the predicted number of installments is between the middle two: high turn cost and low error likelihood (column two), versus low turn cost and high error likelihood (column three).

Translating this effect into the language of the game-theoretic model, we can say that changing the scrolling speed affected not only the *turn cost function*, but also the *error likelihood function* of the puzzle task. Section 4.3 showed how different turn cost functions and error likelihood functions interact with each other to produce different predictions regarding increment frequency and length. There, turn cost functions were modeled as simple linear functions, with total turn cost directly proportional to the number of turns taken. On the other hand, error likelihood functions were generated using (nonlinear) logistic functions. Figure 4.5 (reproduced here as Figure 5.15) shows four hypothetical

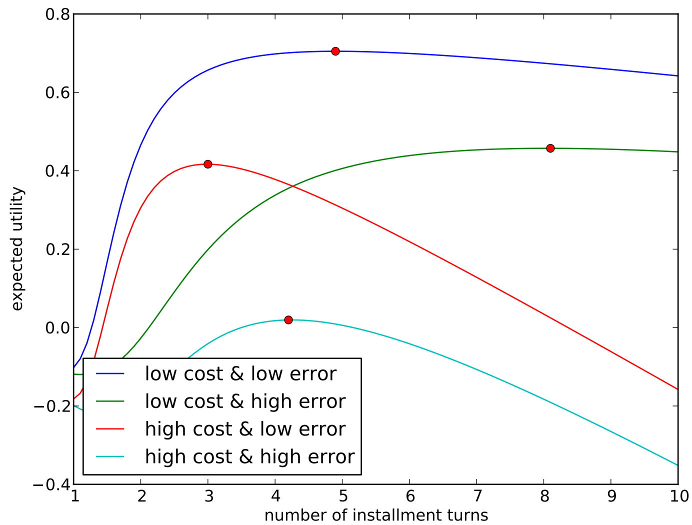

Figure 5.16. Expected utilities with high vs. low turn costs and high vs. low error growths

error likelihood functions, which differ from each other in the location on the x-axis where the curve has a maximum growth rate. The fastest growing error likelihood function is the leftmost one in the figure, while the slowest growing function is the rightmost one. Tying this back to the experiment, we can associate the faster scrolling speed with an error function that grows faster (at lower values of x), and the slower scrolling speed with a function that grows more slowly (at higher values of x).

Figure 5.16 shows the result of composing high and low error likelihood functions with high and low turn cost functions into expected value calculations. The four curves shown in this figure are the result of making the four possible combinations of the two turn cost and error likelihood functions. These expected utility functions are composed from the four possible combinations of a low turn cost (turn cost $= 0.02$) and a high turn cost (turn cost = 0.1) function, and a slow error growth (max growth = 6) and a fast error growth (max growth = 3) function. The expected utility function that predicts the fewest number of installment turns (and by implication, the longest installments) is the one with with the high turn cost and the slow error growth. This case corresponds to the slow scrolling speed manipulation in the experiment. The expected utility function that predicts the highest number of installment turns (and the shortest installments) is the one with low turn cost and the fast error growth. This case corresponds to the fast scrolling speed manipulation of the experiment. The predicted interaction between turn cost and error likelihood functions appears to hold true of the experimental results discussed in this chapter, and helps to explain the very strong effects that were found across the delay and no-delay conditions.

CHAPTER 6

Conclusion

This chapter concludes the thesis with a summary of its contributions, and by presenting ideas for future research.

6.1. Contributions

The core theoretical contribution consists of the development of a game-theoretic model of grounding, and its application to the domain of referential communication tasks. The empirical contribution consists in the experimental evaluation of predictions of the game-theoretic model, using a novel version of an online referential communication task.

6.1.1. Theoretical Contributions

The core theoretical contribution of this thesis was the development of a game-theoretic model of grounding for referential communication tasks (Chapter 4). The starting point for this model was the existing theory of *signaling games* (Lewis, 1969; Spence, 1973; Maynard Smith, 1982; Benz et al., 2005). Formally, signaling games are a restricted kind of extensive game with imperfect information, called *Bayesian extensive games with observable actions*. In this type of game, every player observes the actions of every other player, and the only uncertainty is about an initial move of chance that distributes payoff relevant information among the players in such a way that the information given to each player does not reveal the information given to other players (Osborne and Rubinstein, 1994).

This thesis shows that standard signaling games are insufficient to model grounding, because there is no means for representing imperfect information about another player's actions. Therefore, an extension to signaling games was proposed, called *signaling games with partially observable actions*. Signaling games with partially observable actions are a type of extensive game with imperfect information, and they generalize signaling games by including an *observation model* of the informed player's communicative actions. The addition of an observation model allows for the possibility of exogenously determined imperfect information with respect to communicative actions.

Figure 6.1 shows the `is-a` relationships of the various types of extensive games discussed in this thesis. For example, a signaling game is a type of signaling game with partially observable actions – one in which an observation always assigns full probability to a single message. Signaling games are therefore a corner case for signaling games with partially observable actions. Figure 6.1 highlights the fact that the theoretical contribution of this work has not been to game theory per se, but rather the identification of a particular sub-class of games that are relevant to the problem of grounding in communication. Since this class of games has a clear and simple relationship to well-known classes of games within game theory, standard solution techniques (such as the *sequential equilibrium* concept) can be brought to bear to analyze them.

Once this particular sub-class of games was identified, the next theoretical contribution of the thesis was to connect it back to the four core ideas of grounding theory that were enumerated in Chapter 2: (1) language use is *joint action*, (2) joint actions are coordinated

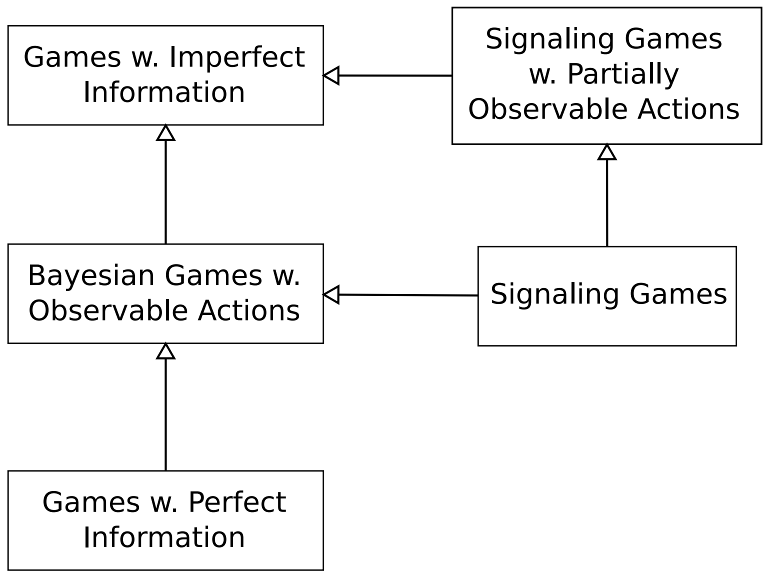

Figure 6.1. Game types and their relationships to one another

via the *common ground*, (3) the minimum amount of effort that dialog participants expend to add something to the common ground is determined by the *grounding criterion*, and (4) the maximum amount of effort that dialog participants expend to add something to the common ground is determined by the *principle of least collaborative effort*.

For (1), the conclusion of this thesis was that the game-theoretic model captures some, but not all of the Clark's notion of joint actions. It captures the coordination of the *content* of actions, but not necessarily the coordination of *process*. This latter aspect of Clark's notion of joint actions requires a modeling formalism that operates at a different level of granularity than a typical game tree can realistically provide.

For (2), the conclusion of the thesis was that common p-belief (Monderer and Samet, 1989) captures much of what is intended by Clark's notion of common ground. Common p-belief allows us to formally capture the idea that dialog participants can have degrees of certainty that common ground has been achieved. The ability to model degrees of certainty in the common ground is a prerequisite for the grounding criterion (or something analogous) to be a useful notion. It also allows for the possibility that coordination among dialog participants may sometimes fail; this possibility agrees with empirical observation.

For (3) and (4), the conclusion of the thesis is that these notions are not required as independent stipulations, but follow automatically from the game-theoretic approach. If we assume (as is standard in game theory) that agents always work to maximize their expected utility, then the grounding criterion and the principle of least collaborative effort naturally follow. Parsimoniously then, we can dispense with them in the game-theoretic approach.

The final theoretical contribution of this thesis was to apply the game-theoretic model of grounding to the area of installment noun phrases. Previous work in psycholinguistics has suggested that the size of installments is governed by principles of grounding theory (Clark and Wilkes-Gibbs, 1986; Brennan, 1998). In Section 4.3 these ideas were made precise, and the game-theoretic model was used to generate predictions about the size of referential description installments given functions for error likelihood, turn cost, and task reward. One of the primary advantages of this formal model over the informal arguments in the grounding theory literature is that the formal approach leads directly to the possibility of a computational implementation. A theoretically motivated and principled approach to the generation of appropriately sized chunks of speech or text has not received attention in existing literature on dialog modeling or natural language generation. The theoretical contribution of this thesis enables progress towards this goal.

6.1.2. Empirical Contributions

The experiment described in Chapter 5 was a step towards empirically validating the predictions of the game-theoretic model. It focused specifically on predictions with respect to the size of referential description installments. The experiment builds upon existing empirical work on referential communication tasks, and is closely related to (and inspired by) the puzzle task studies of (Gergle et al., 2004a,b,c; Gergle, 2006; Gergle et al., 2006). This previous work has shown that changing the communication modality (speech vs. text) and making the workspace shared or not shared has a significant impact on properties of the language used by the participants, including the size of referring expression installments.

The main contribution of the experiment described in this thesis was the manner in which it disaggregated turn taking costs and task error costs from modality selection. In all conditions of the experiment, the same (text-based) modality was used to communicate. The manipulations occurred directly at the level of turn cost and task error cost – which are primitives of the game-theoretic model of grounding. The turn cost manipulation was found to be highly significant, greatly affecting the size and number of the messages that the participants sent to one another. The task error cost was not found to be significant, but there is a reasonable interpretation that this was due to an inadequacy in the design of the experiment.

6.2. Future Work

Moving forward, there are three directions that immediately suggest themselves as topics for future research: (1) further theoretical development of the game-theoretic model of grounding, (2) computational implementation of a dialog agent that instantiates the model, and (3) further experimental work testing its adequacy as a descriptive model of human performance.

6.2.1. Theoretical Development

Theoretical development of the model could take many directions. Here I just mention two:

First, we can extend the model to handle games other than games of pure coordination. There are many situations where the interests of agents only partially overlap, or do not overlap at all. These situations are modeled in game theory by assigning to players non-identical preference relations over outcomes. Morris and Shin (1997) describe results from game theory that show how the strategic concerns of players interact with their belief states, in terms of the degree of coordination that results. For example, it turns out that whether or not common p-belief (at a certain level of p) is necessary for coordination partially depends on whether or not the interests of the players are aligned. It would be interesting to translate these results into the model of grounding described in this thesis, to see if different behaviors are predicted to be optimal given different arrangements of payoffs to the players.

Second, the model can be applied to additional types of grounding behaviors, other than just noun phrase installments. The literature on conversational analysis and grounding describes a wide array of dialog behaviors that are used by participants as coordination devices (Schegloff, 1968; Sacks et al., 1974; Clark and Wilkes-Gibbs, 1986; Clark and Brennan, 1991; Traum, 1994). These devices include acknowledgements of understanding, requests for clarification, repairs, try markers, and verbatim repetitions. An example of this last device is shown in Example 6.1, taken from Clark and Brennan (1991):

(6.1) **A.** It's Cambridge 12345

B. 12345

A. That's right.

B. Thank you very much.

In this example, participant B is calling a directory service, and participant A is providing information regarding a phone number. Participant B repeats verbatim the sequence of digits that has been provided to him by A, in order to confirm his understanding of this information. Such verbatim repetition is presumably more costly than a simple acknowledgement (such as "That's right" or "okay"), but results in a higher degree of certainty in the common ground. There is nothing to prevent the game-theoretic model developed in this thesis from being applied to this type of grounding behavior. In general, the model should be applicable in many situations where there is a trade-off of cost and uncertainty regarding a communicative act.

6.2.2. Computational Implementation

As stated in the introduction, one of the ultimate goals of this work is to make the core ideas and intuitions of grounding theory precise enough to be used in a computational implementation of a conversational agent. My intentions here are to leverage work on existing implementations of conversational agents, especially those that take theoretically motivated approaches to conversational grounding (Traum and Hinkelman, 1992; Traum and Allen, 1992; Traum, 1994; Paek and Horvitz, 1999, 2000b,a; DeVault et al., 2005;

DeVault and Stone, 2006, 2007; DeVault, 2008). In particular, I intend to take advantage of the empirically motivated ontology of conversation acts developed by Traum and Hinkelman (1992); Traum (1994). In Section 4.4.2 I pointed out that there is nothing to prevent the game-theoretic model from using the grounding act model's ontology of conversation act types, by including them in the game players' action sets. This move would marry the strengths of the grounding acts model (an empirically derived ontology of act types) with the strengths of the game-theoretic model (a normative approach to action that relates action to uncertainty and cost).

I also intend to take advantage of another strand of recent research in dialog systems, which has taken the approach of modeling dialog as a (partially observable) Markov decision process (MDP) (Levin et al., 2000; Roy et al., 2000; Williams et al., 2005; Williams and Young, 2007). MDPs extend decision theory to the domain of sequential problem solving (Russell and Norvig, 2003; Puterman, 2005). An MDP models a discrete process in which the state of the world evolves stochastically at each time step according to the actions of an agent. A solution to an MDP is a policy, which is defined as a function that assigns an action to every possible world state. An optimal policy is a policy that maximizes expected reward over a specified time horizon. A *partially observable* MDP (POMDP) is one in which the agent does not have direct access to the true state of the world, but receives an *observation* that gives probabilistic information about it (Kaelbling et al., 1998). Thus, the agent must maintain a belief distribution over the possible states of the world, given a world model and his observation history.

There are clear connections between the game-theoretic approach to grounding described in this thesis, and the POMDP approach to dialogs taken by Williams et al.

(2005); Williams and Young (2007). The crucial difference is that the game-theoretic approach is multiagent, and POMDPs are single agent. However, there has been recent work generalizing POMPDs to the multiagent case, in the form of *decentralized POMDPs* (Dec-POMDPs) (Bernstein et al., 2002; Spaan et al., 2006; Seuken and Zilberstein, 2008). It should be relatively straightforward to translate a signaling game with partially observable actions into the Dec-POMDP formalism. This move would preserve the essential insights of the game-theoretic model, but at the same time take advantage of existing computational algorithms for solving Dec-POMDPs. In particular, Spaan and Oliehoek (2008) have created an open-source toolbox for constructing and solving Dec-POMPDs which could serve as the basis for a computational implementation of the game-theoretic model. I intend to explore this possibility in future work.

6.2.3. Experimental Work

The experiment described in Chapter 5 should be viewed as a first step towards empirically validating the game-theoretic model of grounding. The strong main result obtained for manipulating turn cost is a positive sign that the model can be successfully applied to describe actual dialog behaviors, and together with the body of work on referential communication tasks summarized in Section 2.2 provides general support for the predictions of conversational grounding theory.

However, much more work remains to be done in this area. With respect to the experimental results obtained in this thesis, the obvious next step is to correct potential flaws in the task error cost manipulation in order to see if significant results can be obtained with an improved experimental design. I also hope to perform future experimental work to more realistically estimate the error likelihood and cost parameters of the puzzle task, in order to make the predictions of the model with respect to installment size more quantitatively precise. Finally, given a computational implementation of the game-theoretic model, future empirical work will be required in order to evaluate the implementation, by comparing its performance with the performance of baseline implementations.

6.3. Final Thoughts

Despite the early pioneering work of Lewis (1969), the application of game theory to linguistics is still in its early stages. In this thesis, I have made some progress extending the game-theoretic approach to the domain of conversational grounding, thereby formalizing and clarifying the core ideas of conversational grounding theory. Game theory makes it possible to be mathematically precise about what is meant by terms like *common ground*, *cost*, and *reward*, and to be precise about their relationships to one other. Game theory also makes it clear that notions like the *grounding criterion* and the *principle of least collaborative effort* are not necessary as independent stipulations, but follow naturally from general principles of rational behavior. Moving forward, the real test of the utility of the game-theoretic approach is the degree to which it can be used to elucidate empirically observed human behavior in dialog, and the degree to which it proves useful for implementing practical computational conversational agents that can interact with humans in natural settings. These are promising avenues for further exploration.

References

James Allen, George Ferguson, Bradford W. Miller, Eric K. Ringger, and Teresa Sikorski Zollo. *Handbook of Natural Language Processing*, chapter Dialogue systems: From theory to practice in TRAINS-96, pages 347–376. Marcel Dekker, 2000.

James Allen, George Ferguson, and Amanda Stent. An architecture for more realistic conversational systems. In *Proceedings of Intelligent User Interfaces 2001 (IUI-01), Santa Fe, NM*, pages 1–8, January 2001.

Robert Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

Robert Aumann and Adam Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):11611180, 1995.

J.L. Austin. *How to do things with words.* Oxford: Oxford University Press, 1962.

Anton Benz, Gerhard Jaeger, and Robert Van Rooij. *Game Theory and Pragmatics*, chapter An introduction to game theory for linguists. Palgrave, 2005.

D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, 2002.

Susan Brennan. *Seeking and providing evidence for mutual understanding.* PhD thesis, Stanford University, 1990.

Susan Brennan. The grounding problem in conversation with and through computers. In S. R. Fussell and R. J. Kreuz, editors, *Social and cognitive psychological approaches to* *interpersonal communication*, pages 201–225. Lawrence Erlbaum, Hillsdale, NJ, 1998.

Susan E. Brennan. How conversation is shaped by visual and spoken evidence. In John Trueswell and Michael Tanenhaus, editors, *Approaches to Studying World-Situated Language Use:Bridging the Language-as-Product andLanguage-as-Action Traditions*, pages 95–129. MIT Press, Cambridge, MA, 2005.

Janet Cahn and Susan Brennan. A psychological model of grounding and repair in dialog. In *Proceedings, AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*, pages 25–33, North Falmouth, MA, 1999. American Association for Artificial Intelligence.

Herbert Clark. *Arenas of Language Use*. University of Chicago Press, 1992.

Herbert Clark. *Using Language*. Cambridge University Press, 1996.

Herbert Clark and Susan Brennan. Grounding in communication. In L.B. Resnick, J.M. Levine, and S.D. Teasley, editors, *Perspectives on socially shared cognition*, chapter 7. APA Books, 1991.

Herbert Clark and Meredyth Krych. Speaking while monitoring addressees for understanding. *Journal of Memory and Language*, 50(1):62–81, 2004.

Herbert Clark and Catherine Marshall. Definite reference and mutual knowledge. In Aravind Joshi, Bonnie Webber, and Ivan Sag, editors, *Elements of discourse understanding*, pages 10–63. Cambridge University Press, 1981.

Herbert Clark and Edward Schaefer. Collaborating on contributions to conversation. *Language and Cognitive Processes*, 2:19–41, 1987.

Herbert Clark and Edward Schaefer. Contributing to discourse. *Cognitive Science*, 13: 259–294, 1989.

Herbert Clark and Deanna Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22:1–39, 1986.

Philip Cohen. The pragmatics of referring and the modality of communication. *Computational Linguistics*, 10:97–146, 1984.

Nelson Cowan. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24:87–185, 2000.

Vincent Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 6: 1431–1451, 1982.

Robert Dale and Ehud Reiter. Computational interpretations of the Gricean Maxims in the generation of referring expressions. *Cognitive Science*, 19:233–263, 1995.

Kris de Jaegher. *Game Theory and Pragmatics*, chapter Game-theoretic grounding. Palgrave, 2005.

Kris de Jaegher. Efficient communication in the electronic mail game. *Games and Economic Behavior*, 63:468–497, 2008.

David DeVault. *Contribution Tracking: Participating in Task-Oriented Dialogue under Uncertainty*. PhD thesis, Department of Computer Science, Rutgers, The State University of New Jersey, New Brunswick, NJ, October 2008.

David DeVault and Matthew Stone. Scorekeeping in an uncertain language game. In David Schlangen and Raquel Fernandez, editors, *BRANDIAL 06: Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, pages 139–146, 2006.

David DeVault and Matthew Stone. Managing ambiguities across utterances in dialogue. In *The 2007 Workshop on the Semantics and Pragmatics of Dialogue (DECALOG 2007)*, 2007.

David DeVault, Natalia Kariaeva, Anubha Kothari, Iris Oved, and Matthew Stone. An information-state approach to collaborative reference. In *ACL 2005 Proceedings Companion Volume*, pages 1–4. Association for Computational Linguistics, 2005.

Ronald Fagin, Joseph Halpern, Yoram Moses, and Moshe Vardi. *Reasoning about knowledge*. MIT Press, 1995.

Joseph Farrell and Matthew Rabin. Cheap talk. *Journal of Economic Perspectives*, 10 (3):103–118, 1996.

Drew Fudenberg and Jean Tirole. *Game Theory*. The MIT Press, 1991.

John Geanakoplos. *Handbook of Game Theory*, volume 2, chapter Common Knowledge, pages 1447–1496. North Holland, 1994.

Jelle Gerbrandy. Communication strategies in games. *Journal of Applied Non-Classical Logics*, 17(2):197–211, 2007.

Darren Gergle. *The Value of Shared Visual Information for Task-Oriented Collaboration*. PhD thesis, Carnegie Mellon University, Human-Computer Interaction Institute, 2006.

Darren Gergle, Robert Kraut, and Susan Fussell. Language efficiency and visual technology: Minimizing collaborative effort with visual information. *Journal of Language and Social Psychology*, 23:491–517, 2004a.

Darren Gergle, Robert Kraut, and Susan Fussell. Action as language in a shared visual space. In *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW 04)*, pages 487–496, 2004b.

Darren Gergle, David Millen, Robert Kraut, and Susan Fussell. Persistence matters: Making the most of chat in tightly-coupled work. In *Proceedings of CHI 2004*, 2004c.

Darren Gergle, Robert E. Kraut, and Susan R. Fussell. The impact of delayed visual feedback on collaborative performance. In *Proceedings of CHI 2006*, pages 1303–1312, New York, 2006. ACM Press.

H. P. Grice. *Studies in the Way of Words*. Harvard University Press, Cambridge, MA, 1989.

Joseph Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.

Joseph Y. Halpern and Mark Tuttle. Knowledge, probability, and adversaries. *Journal of the ACM*, 40(4):917–962, 1993.

John Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part i. the basic model. *Management Science*, 14(3), 1967.

Laurence Horn and Gregory Ward, editors. *Handbook of Pragmatics*. Blackwell, 2004.

William S. Horton and B. Keysar. When do speakers take into account common ground? *Cognition*, 59:91–117, 1996.

E. Horvitz and T. Paek. A computational architecture for conversation. In *Proceedings of the Seventh International Conference on User Modeling, Banff, Canada*, pages 201–210. New York: Springer Wien, 1999.

Gerhard Jaeger. Applications of game theory in linguistics. *Language and Linguistics Compass*, 2/3:406421, 2008.

Gail Jefferson. A case of precision timing in ordinary conversation: overlapped tag-positioned address terms in closing sequences. *Semiotica*, 9:47–96, 1973.

Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics.*

Prentice-Hall, 2 edition, 2009.

Leslie Kaelbling, Michael Littman, and Anthony Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.

B. Keysar and D. J. Barr. Coordination of action and belief in communication. In J. C. Trueswell and M. K. Tanenhaus, editors, *Approaches to Studying World Situated Language Use: Bridging the Language-as-Product and Language-as-Action.* MIT Press, 2005.

R. Krauss and S. Weinheimer. Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4: 343–346, 1966.

R. M. Krauss and S. Weinheimer. Effects of referent similarity and communication mode on verbal encoding. *Journal of Verbal Learning and Verbal Behavior*, 6:359–363, 1967.

R.M. Krauss and S. Weinheimer. Changes in reference phrases as a function of frequency of usage in social interaction: A preliminary study. *Psychonomic Science*, 1:113–114, 1964.

Robert Kraut, Susan Fussell, and Jane Siegel. Visual information as a conversational resource in collaborative physical tasks. *Human-Computer Interaction*, 18:13–49, 2003.

Roger Kraut, Darren Gergle, and Susan Fussell. The use of visual information in shared visual spaces: Informing the development of virtual co-presence. In *Proceedings of Computer Supported Cooperative Work*, pages 31–40. ACM Press, 2002.

Benny Lee. Mutual knowledge, background knowledge and shared beliefs: Their roles in establishing common ground. *Journal of Pragmatics*, 33:21–44, 2001.

Esther Levin, Roberto Peraccini, and Wieland Eckert. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on Speech and Audio Processing*, 8(1):11–23, 2000.

Stephen C. Levinson. *Pragmatics*. Cambridge University Press, Cambridge, UK, 1983.

David Lewis. *Convention*. Harvard University Press, Cambridge, MA, 1969.

John Maynard Smith. *Evolution and the theory of games*. Cambridge University Press, 1982.

Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190, 1989.

Stephen Morris and Hyun Song Shin. Approximate common knowledge and co-ordination: Recent lessons from game theory. *Journal of Logic Language and Information*, 6:171–190, 1997.

Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

Martin Osborne and Ariel Rubinstein. *A course in game theory*. MIT Press, 1994.

T. Paek and E. Horvitz. Grounding criterion: Toward a formal theory of grounding. Technical Report MSR-TR-2000-40, Microsoft, April 2000a.

Tim Paek and Eric Horvitz. Uncertainty, utility, and misunderstanding: A decision-theoretic perspective on grounding in conversational systems. In *AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems, Cape Cod, MA.*, 1999.

Tim Paek and Eric Horvitz. Conversation as action under uncertainty. In *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence (UAI-00)*, San Francisco, CA, 2000b. Morgan Kaufmann.

Tim Paek and Roberto Pieraccini. Automating spoken dialogue management design using machine learning: An industry perspective. *Speech Communication*, 50:716729, 2008.

Prashant Parikh. *The Use of Language*. CSLI Publications, Stanford, CA, 2001.

Martin Puterman. *Markov Decision Processes, 2nd revised edition*. Wiley-Interscience, 2005.

Ehud Reiter and Robert Dale. *Building natural language generation systems*. Cambridge University Press, 2000.

F. J. Richards. A flexible growth function for empirical use. *Journal of Experimental Botany*, 10:290–300, 1959.

N. Roy, J. Pineau, and S. Thrun. Spoken dialogue management using probabilistic reasoning. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, 2000.

Ariel Rubinstein. The electronic mail game: strategic behavior under “almost common knowledge”. *American Economic Review*, 79:385–391, 1989.

Stuart Russell and Peter Norvig. *Artificial Intelligence, 2nd Edition*. Prentice Hall, 2003.

H. Sacks, E. A. Schegloff, and G. Jefferson. A simplest systematics for the organization of turntaking for conversation. *Language*, 50:696–735, 1974.

Emanuel Schegloff. Sequencing in conversational openings. *American Anthropologist*, 70: 1075–1095, 1968.

Stephen Schiffer. *Meaning*. Oxford University Press, 1972.

John Searle. *Speech acts: An essay in the philosophy of language*. Cambridge University Press, 1969.

Sven Seuken and Shlomo Zilberstein. Formal models and algorithms for decentralized decision making under uncertainty. *Journal of Autonomous Agents and Multi-Agent Systems*, 17(2):190–250, 2008.

Yoav Shoham and Kevin Leyton-Brown. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008.

Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.

Matthijs Spaan, Geoffrey Gordon, and Nikos Vlassis. Decentralized planning under uncertainty for teams of communicating agents. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi Agent Systems*, page 249256, 2006.

Matthijs T. J. Spaan and Frans A. Oliehoek. The multiagent decision process toolbox: software for decision-theoretic planning in multiagent systems. In *Multi-agent Sequential Decision Making in Uncertain Domains, Workshop at AAMAS08*, 2008.

A. M. Spence. Job market signaling. *Quarterly Journal of Economics*, 87:355–374, 1973.

Robert Stalnaker. Assertion. In *Syntax and Semantics*, volume 9, pages 315–332. University of Chicago Press, 1978.

Robert Stalnaker. Saying and meaning, cheap talk and credibility. In *Game Theory and Pragmatics*. Palgrave, 2005.

Amanda Stent. A conversation acts model for generating spoken dialogue contributions. *Computer Speech and Language special issue on Spoken Language Generation*, 16:313–352, 2002.

Matthew Stone. Communicative intentions and conversational processes in human-human and human-computer dialogue. In John Trueswell and Michael Tanenhaus, editors, *Approaches to world-situated language use: Bridging the language-as-product and language-as-action traditions*, pages 39–70. MIT Press, 2005.

Matthew Stone, Richmond Thomason, and David Devault. Enlightened update: A computational architecture for presupposition and other pragmatic phenomena. Draft paper for Ohio State Pragmatics Initiative, 2006.

David Traum. *A Computational Theory of Grounding in Natural Language Conversation*. PhD thesis, Computer Science Dept., U. Rochester, December 1994.

David Traum. Computational models of grounding in collaborative systems. In *Working Notes of the AAAI Fall Symposium on Psychological Models of Communication*, pages 124 – 131. AAAI, November 1999.

David Traum and Peter Heeman. Utterance units in spoken dialogue. In Elisabeth Maier, Marion Mast, and Susann LuperFoy, editors, *Processing in Spoken Language Systems*, Lecture Notes in Artificial Intelligence, pages 125–140. Springer-Verlag, 1997.

David Traum and Elizabeth Hinkelman. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3):575–599, 1992.

David R. Traum and James F. Allen. A speech acts approach to grounding in conversation. In *Proceedings 2nd International Conference on Spoken Language Processing (ICSLP-92)*, pages 137–140, October 1992.

David R. Traum and Pierre Dillenbourg. Towards a normative model of grounding in collaboration. In *ESSLLI-98 workshop on Mutual Knowledge, Common Ground and Public Information*, pages 52–55, August 1998.

Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, 2007.

Robert van Rooij. Signalling games select Horn strategies. *Linguistics and Philosophy*, 27:493–527, 2004.

Jason Williams and Steve Young. Partially observable Markov decision processes for spoken dialog systems. *Computer Speech and Language*, 21:393–422, 2007.

Jason Williams, Pascal Poupart, and Steve Young. Partially observable Markov decision processes with continuous observations for dialogue management. In *Proceedings of SIGDial*, 2005.