The Use of Language

Prashant Parikh

ISBN: 1-57586-354-5

Contents

List of Figures

List of Tables

Acknowledgments

I am indebted to Noël Carroll, Aravind Joshi, Avani Parikh, Rohit Parikh, Stanley Peters, and Stephen Schiffer for many stimulating conversations over the years. I would also like to thank Rohit Parikh and Stanley Peters for their comments on portions of the text.

1 Introduction

Man is a tool-making animal. When he designs a tool, its form partly fits his own make-up, his natural endowment, and partly fits that part of the world where it will be used. A spoon has a handle that allows people to grasp it and has a shape at its front that allows us to scoop up food. Moreover, because man is rational, it is an efficient tool, by which I mean that its shape fits man and the world better than most other shapes. Of course, there are many types of spoons. This only means that many different shapes are efficient. But we also know that when there are conflicting purposes, like utility and aesthetics, spoons are sometimes less than optimal for certain ends.

Language is also a tool fashioned by man. Given the nature of his innate endowment, it is one of his most sophisticated tools. We use it to do various things, like giving commands and asking questions and expressing feelings, but we use it especially to communicate information about the world. Man, as homo faber, has shaped this tool to suit his ends and as such, it bears his marks and the marks of the world which it is about. As with the spoon, the form of language is such that it fits his endowment. It also fits the world.[1] Again, because man is rational, language is efficient; in fact, given its complexity, it is so in a myriad ways.

As with spoons, there are many languages, implying that many forms are efficient—English, Gujarati,[2] and French are only three. There are also many specialized languages like calculus and linear algebra. Of course, here too, there can be conflicting purposes. A language with limited resources (say, a fixed vocabulary) that may be good for one

[1] This fit between language and man and language and the world enables us to study certain general features of man and the world indirectly, through a study of language.

[2] An Indian language from the state of Gujarat in India.

aspect of the world may not be good for another. This tension is most markedly seen in specialized languages like calculus or linear algebra which were invented for quite different purposes.[3]

1.1 Efficiency

In this book, I will confine myself to those aspects of efficiency which have to do with its use in the world.[4]

Unlike mathematical languages (e.g. set theory or general topology) and programming languages (e.g. Lisp and Prolog), the central fact about most natural languages is that they are situated.[5] That is, different propositions can be communicated with the same sentence when it is used in different circumstances. I might say “It’s 4 p.m.” on January 7, 2002 in Bombay. You might utter the same words on July 11, 2002 in New York and communicate something quite different. Natural languages are context-dependent and the contents of utterances depend on the situations in which they are uttered.

It is also not necessary to say “It is 4 p.m. here” or “It is 4 p.m. in Bombay on January 7, 2002” or elaborate the sentence in some other way. Its meaning is clear from the context. Rather than being a defect, it is in fact another major source of efficiency, making language an extremely flexible tool for communicating propositions and also doing things like making statements or suggesting that it is time to go. Indeed, almost all sentences are incomplete in one way or another, making some disambiguation or resolution necessary in every utterance.[6]

Two types of efficiency can occur in the same sentence. If I say “The bank is nearby,” I could mean the financial institution or the river bank. If I meant the river bank in a particular situation, I could still mean either the Hudson or East River, so two levels of efficiency and disambiguation are involved. Of course, in a different situation, I could be referring to Citibank via the same two levels of disambiguation.

Yet another example of efficiency is the possibility of implied meanings (what Grice called implicature). I might say “He put on his shirt and put on his tie” and convey that he put on his shirt and *then* put on his tie. This additional information also comes from the situation rather than the sentence.

---

[3]Sometimes, these languages can be combined, as in vector calculus, in which two or more different purposes can be served equally well.

[4]Chomsky’s recent Minimalist program can be viewed as an attempt to deal with efficiency as it relates to syntax.

[5]The theory of situations was developed by Barwise and Perry in the eighties.

[6]I will be using the concept of ambiguity in an extended way to talk about any situation where there is an utterance with multiple possible meanings or interpretations.

All these examples derive from the situatedness or context-dependence of language. They all have to do with its use.

What makes communication with a situated language possible? It is man's rationality that makes it possible. Rationality allows us to disambiguate and so allows us to consider bare, unsituated sentences with more than one meaning. When such multivalent sentences are situated, they can be rationally disambiguated to yield the content of the utterance. Of course, it may happen that an utterance (i.e. a situated sentence) is intended to have multiple meanings (e.g. a pun), and then disambiguation would yield a set of contents. Or it could be genuinely ambiguous and yield any of its possible contents, a disjunction of contents.

Why aren't all languages situated? While situated languages afford flexibility, they also extract a cost. They are not precise (vagueness is another type of efficiency) and they are not unambiguous. While situated languages can be disambiguated, there is often an element of doubt about the final result. Error is always possible. When clarity and precision are essential, as it is with computers or in mathematics, a formal vocabulary and language free of ambiguity are required. The same is true of scientific terms like "neutrino," medical terms like "cholera," and legal terms like "tort," as in all specialized languages where such terms are tagged onto the ordinary vocabulary of a language like English.

The efficiency of language is the principal underlying theme of this book. Language is efficient in many ways. I try to show how rationality brings about efficiency at the level of use.

1.2 Use

What seems to be common to most uses of language is *intended information flow* between agents via language. It has three aspects: intention, information, and flow. The main type of intended flow is communication.

By information, I actually mean in this context the relation of aboutness between language and world, its having to do with an utterance's being *about* the world. That is, it has to do with how language refers to or connects with the world.

When we use language, we typically use it to communicate information. The two dimensions of communication and aboutness correspond to two constraints, communicative and informational. These two constraints interconnect and jointly enable us to use language to communicate information. In this book, I focus on the communicative constraint.

Austin (1975, 1979b) was perhaps the first philosopher to focus on the use of language, and his ideas and influence underlie the book.

1.3 Communication

Communication is the main type of use. Intended information flow between agents can be of many other types. A speaker might, for example, suggest something to an addressee without actually communicating it. A waiter in a restaurant might say "Would you like to order anything else, sir?" and suggest that it is time to leave. If he were to *communicate* this additional information, it may be taken as a sign of rudeness.[7] Communication implies a relative openness of information flow. In particular, the speaker's intention is revealed in communication.

We all have an intuitive sense of what communication is, a sense that Grice set out to sharpen and capture in a definition.[8] I show in this book how this intuitive sense can be fixed in a mathematically precise way. The starting point for this formalization is the concept of a rational agent.

Communication typically involves two people (although more than two is certainly not uncommon). Occasionally, we write or talk to ourselves, but we can think of such situations as having the same speaker and addressee. Thus, all communication can be analyzed into situations with speakers and addressees. If we are to to start with the idea of rationality and its axiomatic version found in choice theory, this involvement of more than one role/person in communication suggests that we need to consider not single-person choice theory, but interactive or multiperson choice theory. Another better known name for this is game theory.[9]

Choice theory and game theory have become the dominant sources of method and reasoning in economics in the last two decades. They have made major inroads into other disciplines like political science, biology, and law. They are used in philosophy and computer science, and also in psychology, sociology, and anthropology. And, of course, they have been around in logic and mathematics for a long time. Game theory has, however, not yet been used in natural language semantics and pragmatics.[10] I believe it has the resources to put the subfield of communication and information flow on a totally new foundation, one

[7] And the waiter doesn't want to lose his tip, so his is a delicate task.

[8] Grice actually tried to define speaker meaning, a notion that is closely related but different. This concept was his starting point for a program that Schiffer (1972) has called *intention-based semantics*.

[9] This is the first difference between traditional approaches to speaker meaning, where there is an almost exclusive focus on a single person, the speaker, and the approach this book takes, where the addressee plays a crucial role from the outset. This is also why I start with communication rather than speaker meaning.

[10] There has been some effort in this direction by Hintikka and Lauri Carlson, but their approach is very different from mine and they tackle somewhat different issues as well.

that is mathematically solid, philosophically sound, and computationally tractable.[11]

In single-person choice theory, one rational agent confronts a set of choices. His task is to choose the best option from this set. In game theory, which is a theory of social interactions, two or more rational agents have to act simultaneously or sequentially and try to choose their best options taking the other agents' possible actions into account. When speaker and addressee communicate, they must each take the other's intentions and choices into account, if they are to do so effectively and efficiently. Thus, game theory seems like a suitable framework for formalizing the theory of communication. As I said above, most utterances require disambiguation, and this is where game theory has much cutting power, apart from its basic role in the definition of communication and speaker meaning. Games make disambiguation possible. I will be arguing this in detail throughout the book.

The central person underpinning this dimension of use is undoubtedly Grice (1989). I do not attempt to summarize his ideas here, though their development in the book is more or less self-contained. In particular, I do not try to address his various attempts at definition. For this, it is best to refer to Schiffer 1972. Grice initiated this inquiry into communication, and Grice, Strawson, and Schiffer developed it in considerable detail. It is now time to analyze it afresh with the tools of game theory.

1.3.1 Applications

Utterances are actions, of course, but some actions also function as utterances. One might give someone something in order to communicate that one is generous. Indeed, this sort of communication pervades social life. One can offer a warranty on a product to signal the product's quality (economics), form an alliance with someone to indicate one's solidarity (political science), give a gift to someone to communicate one's status (anthropology), and push someone around to imply one's power (sociology).

In somewhat more detail, a salesman can offer a warranty for a used car to signal its high quality. Here, he is offering two things, the warranty itself and some indirect indication of its quality. This kind of action is necessary in a context where mere talk is not likely to be believed. We will see later that this is an example of implied meaning or what Grice called implicature. We considered an example of implicature above (the shirt and tie example).

[11] However, this is not a book on game theory as such. I have applied the game theory I have found useful, and largely developed things from first principles. The book is self-contained as far as game theory is concerned.

One political party may form an alliance with another and thereby also signal its solidarity with that party. Here again, indirection is necessary either because talk is cheap or because talk may be too costly. That depends on the context.

Gift-giving is a well-known way to communicate one's status in the relevant community. It is not enough simply to proclaim one's status because, once again, talk is cheap.

Lastly, a bully may push someone around partly just to show his power. Here he is again doing two things, actually pushing someone around and also conveying his power.

We could speculate that we use actions to communicate when talk is either cheap or not possible for some reason. In most types of dance, for example, there is a self-imposed restriction to using gestures and other movements to communicate.

All these examples are also, of course, context-dependent; they can communicate different contents in different situations. For example, gift-giving can convey generosity or status or both. So communication is integral to social life. The theory of communication I develop in this book applies to all such actions, not just to ordinary utterances. In common parlance, it applies to both verbal and nonverbal communication. I will, of course, model mainly ordinary language utterances to keep things simple and make the model clear. I invite the reader to apply the theory to his favorite domain.

Communication and information flow occur also in the arts. Literature, painting, sculpture, drama, film, photography, dance, music, and architecture all involve actions which are utterances or function as utterances and thereby often communicate information and ideas. For example, a novel or painting may convey information about a (real or fictional) situation. In part II, I take up the concepts of visual representation and also visual implicature. They can also be approached via game theory, and may be one of the first detailed applications of mathematics to the visual arts. This analysis can quite easily be extended to the other arts.

If we are willing to impute intentions, beliefs, and desires to computers, then it is arguable that we also communicate with computers via programs. Programs are unambiguous and so the games involved are in some sense trivial, but they still satisfy the conditions for communication. One day we may possibly communicate with computers and robots in English and other languages, and robots will also communicate with each other. While we would certainly be able to use game theory to analyze such communication, it may also happen that robots will actually be programmed to be game-theoretic agents themselves. So this book also has applications to artificial intelligence in general and natural language processing in particular.

It is interesting that essentially one theory or model applies to such a wide range of utterances and actions, and consequently to such a wide range of disciplines. This is because the model we will construct plays a central role in the *definition* of communication, and communication is a very wide concept and occurs wherever there are agents interacting with one another. In other words, if you have agents in the situations you are looking at, you also invariably have communication. As more researchers take up the methods of choice and game theory, extensions of this theory and new models will no doubt arise, not only of the same phenomena, but also of other related phenomena. The field is wide open.

Of course, we don't have to *communicate* when we use language. As I said above, we can hint, suggest, adumbrate things which go beyond the more stringent requirements of communication. All these possibilities fall under the broad rubric of information flow or information transfer. Communication is only the primary type of use and the use of language is itself one type of information flow. I will have more to say about this later.[12]

1.4 Rationality

It is arguable that rational agency (in its mathematical form) is the principal missing element in contemporary studies in natural language semantics and pragmatics, whether from a philosophical or linguistic point of view. Ultimately, it is the source of the efficiency of language. Supplying this missing element is an underlying concern of this book.

1.5 Summary of Book

The book starts with chapter 1, the Introduction, and is then divided into two parts. Part I is "Communication and Information Flow" and part II is "Extensions."

In part I, I build the basic model of communication and information flow that I will use throughout the book.

Chapter 2 starts with a basic distinction between meaning and content. I then introduce game theory and set up the problem of communication and choose the main example I will work with. I also list other examples to show the wide range of phenomena the model can cover.

In chapter 3, I build the game-theoretic model by identifying one set of sufficient conditions for communication. I do this from first principles and it is completely self-contained, so readers new to game theory can

[12] I give a definition in chapter 6 of what it is for one agent to make information flow to another agent and definitions of a whole host of related notions.

follow its development step by step. I also develop the model with some attention to the assumptions that justify the various steps.

In chapter 4, I show how such a model is to be "solved." I use the notion of a Pareto-Nash equilibrium as a solution concept. The solution yields the desired interpretation of the utterance in the main example. It gives us its literal content.

In chapter 5, I examine the assumptions I have made, including the fundamental assumption of common knowledge. This points to more general models than games that I call strategic interactions.

In chapter 6, I formally define communication, and give a number of related definitions of "nonnatural" meaning, interpretation, joint act, vague communication, and what it is to make information flow. I also give a new picture of the entire domain of information flow, whether natural or nonnatural, in terms of three infinite lattices.

Part II contains extensions of the basic models of part I to a number of important semantical phenomena.

In chapter 7, on conversational implicatures, I take up a further dimension of efficient use by modelling how more can be communicated than is literally said. This is the chapter that applies also to actions that function as utterances, of the type that occur in the social sciences. The communication in such actions is not literal (as in part I) but implied or, more precisely, "implicated." It seems possible to develop many applications here. This chapter is relatively brisk because we have already developed the tools we need and have only to extend them to apply to implicated meanings.

In chapter 8, I look at the topic of illocutionary force and the addressee's response to it. This brings in Austin's ideas in a direct way. It turns out that figuring out the force of an utterance is similar to the task of figuring out its content. Modelling the response of the addressee involves extending the game in a natural way. This sort of game would be directly useful in programming computers to understand natural language communication.

In chapter 9, I consider miscommunication, mainly from the point of view of understanding communication better, but also because it has interesting properties that the game theory enables us to elucidate. I take an example from Deborah Tannen's popular book *You Just Don't Understand* (1990, from the field of sociolinguistics) and show how it can be analyzed game-theoretically. This is a short chapter even though it is possible to apply game theory to many of Tannen's examples.

In chapter 10, I extend the foregoing to visual representation. It is interesting that visual representation turns out to be isomorphic to linguistic and other communication. As an example, I take up a famous painting by the Dutch painter Gerard ter Borch, which has two well-known interpretations.

In chapter 11, I discuss visual implicature, a new topic. I use a well-known article on Picasso by Krauss (1986) as a foil to make the discussion concrete and also, incidentally, to argue against semiology. This analysis has many applications to the visual and other arts and to art criticism.

I then apply, in chapter 12, these ideas to jokes, puzzles, and problems, all considered as instances of discourse. Apart from being amusing, jokes are no less efficient than other types of communication. I even give a definition of a subclass of jokes. Puzzles and problems do not fall directly within the ambit of communication, but nevertheless do involve pragmatic implication, and I consider them briefly as instances of such. What is more important perhaps is the implicit application and extension of the model to discourse. This is a large topic with many problems and the game theory could be developed extensively here. The definition of a subclass of jokes can also be seen as a definition of certain types of discourse.

I conclude in chapter 13, pointing out some directions for the future.

The Appendix sets out the formal model of communication and derives some of its important properties.

The book is primarily intended for researchers in philosophy, linguistics, artificial intelligence, and cognitive science generally. It should be of interest to game theorists, and economists and others in the social sciences. It may also induce a few in the arts to learn some game theory.

More generally, the world is changing. With this change, the fault lines dividing disciplines are also shifting. Words like "information" and "communication" are already central to many disciplines, not to mention everyday life. They have moved from their relatively narrow origins in engineering to form some of the core concepts of many fields. Both are even major words in the field of business, perhaps the dominant activity in the world today. The many popular books on communication also signal its importance to everyday life. In other words, "information" and "communication" are multidisciplinary and multi-perspectival concepts in an increasingly globalized world. The advent of computers and networks of computers over the last fifty years have in part made this change possible and in part have speeded it up.

Information science can become the foundational discipline for a number of disciplines dealing with social life, if taken beyond its narrower forms in computer science today. It could play a role analogous to physics in the natural sciences.[13] Communication and its many variants would be basic concepts in such a broader science of information. And game theory would be one of its basic tools.

Such foundational concepts therefore demand a theory today.[14] No one book can fulfil all these demands of course. But this book tackles the core issue, which is to understand what communication and its many variants are, what its many aspects are, and why it is efficient. The book also shows how these foundational concepts can connect with many disciplines along the way.

[13] In fact, many also consider information as a foundational concept of physics.

[14] See Barwise and Seligman's (1997) recent book *Information Flow*.

Part I

Communication and Information Flow

2

Setting Things Up

Communication is one dimension of use. I start by developing a theory of communication based on Grice's insights, using game theory and situation theory in an informal manner. I will then give a definition of several fundamental concepts, including communication, speaker meaning, and addressee interpretation.

I first describe the basic ideas that undergird communication, the ideas of meaning and content and certain aspects of game theory and strategic inference. I then introduce our two main characters $\mathcal{A}$ and $\mathcal{B}$, and try to fix the parameters of the problem of communication in somewhat precise terms. Greater precision will come as we move along, but this will serve as a starting point.

Having done this, I present a range of types of examples that my theory will be able to account for. This list is far from exhaustive, but it provides an initial indication of the scope of the model. I then take up one of these (that I happen to like for reasons that will become clear later) and analyze it in a way that yields a set of sufficient conditions for communication. As I do this, I urge the reader to keep the other examples in mind, since the analysis applies equally to them, and perhaps also test the model with their own examples. This model forms the core of the book; the rest is extension and refinement.

All this takes us up to chapter 5. I then abstract from and generalize these sufficient conditions and give necessary and sufficient conditions for communication, which enable us to formalize and define the intuitive concept. This is one of the major fruits of the theory.

Language, I said, is a special kind of tool. It is in fact a complex social institution. All social institutions are of course tools that enable us to organize different aspects of social life. Different institutions serve different functions in society and it seems plausible to say that the primary function of language is communication. Indeed, it is possible to see language, in particular, meanings, as arising from the interactions of a group of agents. This is how all social institutions emerge and language is no different, except that it arises from the *communicative* interactions of agents. Lewis (1969) and Schiffer (1972), for example, have argued for such a view. Grice's (1957, 1969) ideas on nonnatural meaning provide the best starting point for us.

There are two ways to proceed. I could try to build things from the ground up where our communicators $\mathcal{A}$ and $\mathcal{B}$ do not have a language but construct it as they communicate, or I could assume a language is available. The second option is much easier and I will adopt it. But it is interesting that essentially the same type of analysis can be extended to the first option, to the question of how a language originates and is maintained over time. Simpler things first, however.

2.1 Meaning and Content

Once we allow situations a role in the determination of content, it becomes clear that there are certain aspects of utterances that are constant across utterance situations and there are others that vary from one situation to another. One of the most salient linguistic constants is the meaning of a sentence; this is different from its content in an utterance, which varies from situation to situation.[1] Meaning is the collection of possible contents of a sentence. If $\mathcal{A}$ utters it in a situation, it allows $\mathcal{B}$ to disambiguate it and choose one or more propositions from this collection as its content. We could in fact write a simple schematic equation of the form "(meaning of) sentence $\oplus$ discourse situation (or the situation of utterance) = content."

Part of the task of a theory of communication is to explain how this equation comes about. If a language is given, its meanings are given, and then the problem is to get from meaning to content via the discourse situation. Solving this problem in a completely general way turns out to be an extremely difficult task. To carry it out, we will need some tools.

2.2 Game Theory

I bring to this problem the powerful ideas of game theory developed by von Neumann, Nash, Arrow, Debreu, Aumann and other game theorists and economists. The ideas of rational agency, strategic interaction, and equilibrium developed in this tradition provide the framework we need to solve this problem. They allow us to extract one more salient constant

[1] Many writers use "meaning" to refer to what I am calling content. This is also the colloquial use, and this is how I have used it so far. Unfortunately, there are two different concepts we need to talk about, so we need two different terms. I also use "meaning" to refer to the related meaning function.

(like meaning) from the discourse situation and refine the schematic equation above to "agent architecture $\oplus$ sentence meaning $\oplus$ situation of utterance = content."

The idea of rational agency tells us, via its axioms (see Myerson 1995), how a rational agent chooses an action from a set of actions. In our case, these actions are utterances and interpretations. The agent has a preference ordering over these actions, which directs his choice of action, and which can be translated into a numerical scale. Each action thus results in a payoff that can be measured on this utility scale. The agent then chooses the action with the highest utility. A slight wrinkle is introduced though, when we consider uncertainty. If payoffs are uncertain, as they often are (for example, if we consider buying a lottery ticket as an action), then the agent assigns a probability distribution to the possible outcomes, and a payoff to each outcome. In this case, the agent chooses the action with the highest *expected* utility.

An agent can no longer do this quite so simply when there are other rational agents around, because the actions of other agents also affect the first agent's payoffs: that is, payoffs are functions of everyone's actions. The idea of strategic interaction tells us how a rational agent takes into account another rational agent's possible actions before choosing his best option. In our case, this means how $\mathcal{A}$ and $\mathcal{B}$ take each other's possible actions into account before choosing their utterance and interpretation. Taking another agent's actions into account involves considering not only his options but also his knowledge and beliefs, especially his shared knowledge (with the other agent) of the situation. This is a generalization of the first idea to a multiperson situation. I call the reasoning of agents in a game *strategic inference.* We could of course consider more than two agents if we wanted, but we will stick to two agents to keep the logic simple and clear.

The idea of equilibrium comes from physics, and is used in the context of game theory to tell us when the combination of choices by two or more agents is in balance. No agent has an incentive to change his action. There are other possible conditions on equilibrium (and this has been an area of research in game theory for some time), but the basic idea is that optimality in the single-person case gives way to equilibrium in the multiperson case.

Grice, and subsequently, Strawson (1964), and especially Schiffer (1972), have shown how communication involves extremely complex interactions between speaker and addressee. These interactions are precisely what I have called *strategic* interactions. However, game theory as currently formulated does not provide a ready-made tool to model communication; it is necessary to develop its insights from first principles.

We will also need to generalize the framework of game theory itself.[2]

Apart from the obvious benefits of formalization, why do we need game theory? Isn't it possible to improve upon what Grice, Strawson, and Schiffer did using their methods? I think not. If we are really to give an account of natural language communication, we have to take into account its situatedness explicitly, which means that we need to be able to disambiguate between multiple contents. I claim this is not possible without mathematics, because probabilities are involved, and this leads to a vastly more complex structure than it is convenient or possible to handle in natural language. Quite apart from ambiguity, I will argue that communication, speaker meaning, and addressee interpretation involve a kind of reciprocal structure between speaker and addressee that may not be easily described, if at all, in natural language. Besides, once one employs the relevant mathematics, things actually become simpler, and this is one of the obvious benefits of formalization. A second benefit is precision. A third is the possibility of defining concepts like communication and deriving their properties rigorously.

This brings us to the parameters of our situation.

2.3 The Situation

We already have two rational agents $\mathcal{A}$ and $\mathcal{B}$. What do we need to know about them to get started? $\mathcal{A}$ and $\mathcal{B}$ have common knowledge[3] of their rationality and assume, moreover, that their interaction is a cooperative one.

Next, we have $\mathcal{A}$ uttering an indicative sentence φ assertively in discourse situation d to convey some information p to $\mathcal{B}$. $\mathcal{B}$ attempts to interpret $\mathcal{A}$'s utterance in d. When $\mathcal{A}$ utters φ, $\mathcal{B}$ uses his knowledge of the language to get at its meaning $m(\varphi)$, which is the collection of possible contents of the sentence.

Some aspects of the utterance will be public.[4] The agent architecture is public before the utterance and the sentence uttered will be publicly available to both agents after the utterance. The meaning of the sentence, being a linguistic constant, will also be assumed to be public. Other aspects will, in general, be private, like the beliefs and intentions

[2] A game is a structure where all agents have common knowledge of this structure. I generalize this notion of a game to what I call a *strategic interaction* where agents no longer have common knowledge of the structure. More about this later.

[3] Common knowledge between $\mathcal{A}$ and $\mathcal{B}$ of a fact f is the requirement that $\mathcal{A}$ knows f, $\mathcal{B}$ knows f, $\mathcal{A}$ knows $\mathcal{B}$ knows f, $\mathcal{B}$ knows $\mathcal{A}$ knows f, and so on, ad infinitum. This concept was first introduced by Lewis (1969) and Schiffer (1972). It has since become a staple of game theory.

[4] Public knowledge is interchangeable with common knowledge, more or less. See Barwise (1989a).

of the speaker and addressee.

Our initial problem is to spell out sufficient conditions for $\mathcal{A}$ to *communicate* p to $\mathcal{B}$ by uttering φ in d.

2.4 Strategic Inference

Sometimes it is helpful to embed a problem in a larger problem, either to get a better perspective on it or to solve the larger problem as a way of solving the smaller problem. We will do it for the first reason, to get a better perspective on communication. I have already hinted at the larger problem in chapter 1.

We can embed communication in the larger picture of information flow developed by Dretske (1981), Barwise and Perry (1983), and Barwise (1997).

Reality can be viewed as consisting of situations linked by constraints. It is the constraint between two situations that makes one situation carry information about (naturally or nonnaturally mean, in Grice's sense) another situation. A smoky situation *involves* a situation with fire in it. This is the constraint we describe when we say "Smoke means fire," an instance of natural meaning. An utterance situation with the sentence "There is a fire" also involves a situation with fire in it. This is the constraint we describe when we say the speaker means something is on fire, an instance of nonnatural meaning. In the first case we would write $s_1 \Longrightarrow s_2$ and in the second, $u \Longrightarrow s_2$. An agent who perceives the first situation (either smoke or the utterance) and who knows the relevant constraint (either natural or nonnatural) can infer the existence of the second situation s_2.

Though the two constraints are quite different, I will argue in chapter 6 that the terms "natural" and "nonnatural" are perhaps not the best way to capture this distinction. The distinction originates with the classical distinction between "natural" and "conventional," but Grice introduced the term "nonnatural" to accommodate nonconventional transfers of information that are not natural, like nonconventional gestures, drawings, sounds, and the like.

Given a group of agents, or distributed system as computer scientists call it, there will be all kinds of information flows. A communication is a special type of information flow between agents. Indeed, it is the type of flow that language makes relatively easy to accomplish, but that is not exclusive to language.

What makes a communicative transfer of information special? While smoke indicates fire, it doesn't *communicate* fire: that is, it doesn't communicate that there is a situation with fire in it. It doesn't do so because the smoke doesn't have the relevant *intention*, which is required because our intuitive notion of communication is that it is something only agents can do. This certainly rules out all inanimate objects, except maybe sufficiently sophisticated computers.[5] What about insects like bees, however? We do say that bees communicate even though they don't have intentions.[6] I suppose we have to admit two differing intuitions here. One intuition is that communication is the mere transmission of information, and the other is that it is something agential and more complex. The problem with the first notion is that the intuitive distinction between animate and inanimate transmission[7] also collapses, and all information flows become communicative. Besides, there is the intuition that human communication is different from mere information flow, and we can do the intuition justice by bringing in intentions to start with.

Grice brought in a lot more conditions as counterexamples to proposed definitions piled up, but the starting point was the requirement that the speaker have an intention to convey the relevant information. One important condition Grice introduced was that this intention be recognized by the addressee. This was required because if $\mathcal{A}$ were to leave a sign (e.g. someone's, say $\mathcal{C}$'s, handkerchief) for $\mathcal{B}$ at the scene of a crime to indicate that $\mathcal{C}$ had been there, $\mathcal{B}$ may not be able to infer that $\mathcal{A}$ had intended to put it there. Intuitively, this is not a case of "full" communication. Something is missing, and this, Grice suggested, is the recognition of $\mathcal{A}$'s intention. Grice, and Strawson and Schiffer after him, developed this line of reasoning considerably, adding more conditions to the definition of communication.

I will sidestep this reasoning involving definitions and counterexamples, and jump directly to building a model and definition of communication. What we need for the moment from the foregoing is that communication involves both the speaker and addressee jointly inferring various things about each other. I will call this joint two-sided inference a *strategic inference*.

My basic insight is that all intended information flows between agents involve a strategic interaction between them. When the strategic interaction is common knowledge between the agents, that is, when it is a *game* (with a unique solution), the flow will be communicative. Roughly then, $\mathcal{A}$ communicates to $\mathcal{B}$ just in case there is an appropriate game between $\mathcal{A}$ and $\mathcal{B}$. It is this insight I will make precise in my definition of communication.

[5]Though here it may be the programmer's intentions that are relevant, which allows us to impute intentions to computers. But more on this in chapter 6.

[6]This itself is perhaps a moot point.

[7]Bees occupy a middle ground.

I argue this by first developing a detailed account of one strategic inference. In my view, every utterance involves many separate acts and corresponding strategic inferences. For example, communication typically involves a referential act. Figuring out the reference will then involve a strategic inference, and in general, each bit of information communicated will require its own strategic inference. So any complete utterance involves a system of simultaneous strategic inferences. These inferences have to be simultaneous because, in general, they codetermine each other. An utterance of "Mary had a little lamb" will require inferring the designata of each of the five words in the sentence, not to mention its internal structure. Only then is it possible to get at the content of the utterance. No individual word has any priority in this determination: that is, there may be interactions among the various strategic inferences, and the embedding circumstances play a vital role in each inference. Mathematically, this amounts to a system of simultaneous equations.[8]

To keep things simple, I will focus on just one strategic inference in isolation. I will assume $\mathcal{B}$ has the partial information obtained from all the other inferences. $\mathcal{B}$'s problem is then to use this partial information together with the utterance situation to get to the intended content.

Consider as an example the sentence "Every ten minutes a man gets mugged in New York." This is a familiar type of ambiguity, typically viewed as an ambiguity between two possible quantifier orderings. One reason for this type of choice is that it is widespread in language and in the literature. There are other ambiguities as well in this sentence. For example, "minutes" is ambiguous between the temporal meaning and the minutes of a meeting; "New York" is ambiguous between the city and the state; and "Every ten minutes" is also vague because it usually indicates "about every ten minutes." But I will consider only the quantifier orderings and assume the rest of the utterance has been disambiguated.

A successful strategic inference requires a number of assumptions involving rationality, the agents' intentions, and their knowledge and beliefs. These assumptions will be our sufficient conditions. An important consequence of the analysis is that the content communicated depends not only on what was uttered, but also, crucially, on what the speaker *might* have uttered but chose not to, and on their shared information about these choices.

I will build a strategic inference step by step from the discourse situation d. This will make the role of the assumptions clearer and suggest ways in which the construction can be generalized or modified to in-

[8]It is also possible to deal with all the inferences as one big inference.

clude other complexities. The constructed structure $g(\varphi)$ turns out to be a new kind of game that I call a game of partial information. The content communicated will then be given by the Pareto-undominated Nash equilibrium of the game.[9]

2.5 Some More Examples First

Lest the reader think the model applies only to the example above or only to this type of ambiguity involving quantifier orderings, I consider in this section many different types of examples of resolution and ambiguity to which the model applies.

1. I'm going to the bank (lexical ambiguity)
2. He saw her duck (structural ambiguity)
3. It is 4 (indexical resolution)
4. He is eating (pronominal resolution)
5. Bill said to Bob that he would join him today (double anaphora)
6. The book is highly original (noun phrase resolution)

Most of the choices for interpretation in the above examples are fairly obvious. Once again, there are multiple ambiguities and resolution problems in each, and I have identified which problem I'm considering in parentheses. My list is far from exhaustive. Indeed, the game-theoretic model applies to any and every type of communication, including visual, gestural, aural, and even olfactory and tactile ambiguities. It applies to all actions. I will apply this model to visual communication in chapter 7.

The reader should keep these other examples in mind as we proceed with the main example. This will make it easy to see how the model can be adapted to these other examples and indeed, to any example of communication.

2.6 The Main Example

Suppose $\mathcal{A}$, after having picked up the information in a recent newsletter of The Muggees Association of New York (i.e. M.A.N.Y.[10]), says to $\mathcal{B}$ in d:

"Every ten minutes a man gets mugged in New York." (φ)

How does this communication take place and what does $\mathcal{A}$ communicate to $\mathcal{B}$? φ is ambiguous: $\mathcal{A}$ could mean either that *some person*

[9] I explain these terms in chapter 4.

[10] This acronym functions as a pun, where both meanings, that is, contents are intended.

or other gets mugged every ten minutes (call this p) or that *a particular man* gets mugged every ten minutes (call this p'). It must be the situation d that enables $\mathcal{B}$ to disambiguate φ.

Intuitively, given the circumstances above, we would be inclined towards the first interpretation. It is, after all, difficult to imagine a man as immune to experience as would be required for the truth of p'. But this much merely tells us that the first reading is the one more *likely* to be true. When can $\mathcal{B}$ select p over p' as the intended content with *complete* certainty?

Of course, there *are* circumstances in which p' rather than p might be the content of an utterance of φ. For example, $\mathcal{A}$ might follow up with "He was interviewed on TV last night." If part of such a discourse, $\mathcal{B}$ would have to interpret $\mathcal{A}$'s first utterance as conveying p'. (Indeed, we shall see in chapter 12 that this is how many jokes work.)

With respect to d, however, it seems plausible to say that $\mathcal{B}$ would infer p as the intended content with certainty. In fact, we could say that $\mathcal{A}$ *communicates* p to $\mathcal{B}$ in d.

I will make two sets of assumptions to explain this disambiguation. The first set applies to all situations, more or less, and has to do with the architecture (or "nature") of communicating agents generally. The second involves more specific circumstantial assumptions pertaining to d.

For the first set, called the Background Assumptions, we assume that both $\mathcal{A}$ and $\mathcal{B}$ are rational agents. (Grice also assumes rationality, but not in its choice-theoretic form.) Moreover, this is common knowledge. This is important because the agents would act differently if they didn't know they shared a common architecture.[11]

$\mathcal{L}$ is a shared language and m is its meaning function. I said earlier that meaning is constant across situations, which is why this assumption is in the background. Later we will see that the meaning function has to be generalized to a connection between an utterance of the sentence and the content. In other words, m may be contextual itself.

The second set of assumptions, called the Circumstantial Assumptions, contain in this particular example the assumption that $\mathcal{A}$ has the intention to convey p to $\mathcal{B}$. We cannot say "intends to communicate" here because that would imply a circularity later when we define communication. The word "convey" just means "transfer." This is in fact the kind of simple intention we have when we communicate. To intend to communicate is to intend something pretty complex. This latter intention results in the simpler intention to convey.

[11] They would have to consider alternative possible architectures and so on.

Background Assumptions

1. $\mathcal{A}$, $\mathcal{B}$ are rational.
2. $\mathcal{L}$ is a shared language.
3. m is a function from $\mathcal{L}$ to the power set of the collection of propositions. I call it the meaning function of $\mathcal{L}$ or just the meaning of $\mathcal{L}$.
4. The above assumptions are common knowledge between $\mathcal{A}$ and $\mathcal{B}$.

Circumstantial Assumptions

1. $\mathcal{A}$ intends to convey p.
2. $\mathcal{A}$ utters φ.
3. $\mathcal{B}$ intends to interpret φ.
4. $\mathcal{B}$ receives and interprets φ.
5. $m(\varphi) = \{p, p'\}$.
6. p' is relatively unlikely.
7. Expressing p, p' unambiguously takes greater effort than expressing them ambiguously.
8. All of the above except (1) and (3) are common knowledge.

TABLE 2.1 Summary of Assumptions

Next, $\mathcal{A}$ utters φ publicly. After all, the process has to get off the ground.

$\mathcal{B}$ must have a corresponding intention to interpret φ. Without it, he will not play his interpretive part.

$\mathcal{B}$ must also receive and interpret the utterance and this must be public.[12] Without publicity, p won't become public at the end of the communication. (The publicity of p is a theorem in the Appendix.)[13]

$m(\varphi) = \{p, p'\}$. Also, p' is relatively unlikely, and expressing p unambiguously takes greater effort than expressing it ambiguously. The

[12]These four assumptions (speaker and addressee intention, and utterance and reception/interpretation) replace Grice's principle of cooperation. That is, if agents act in the right way, communication can occur, but if they don't, communication can't occur. There is nothing that forces them to cooperate, as Grice required. We will see that this principle is also not required for implicature where it originated. However, we will see later that a more subtle type of cooperation between speaker and addressee may be called for.

[13]The case of e-mail is interesting because there is always doubt about whether $\mathcal{B}$ has received the message. E-mail may not be a case of communication, but only of high probability information flow. Of course, to some degree, all communicative flows may be like this, even face-to-face communication. Usually, a co-present addressee responds with movements of the head (nodding) and eye contact, indicating that he is attending to the conversation.

meaning and use of these assumptions will become clearer as we proceed.

Except for $\mathcal{A}$'s and $\mathcal{B}$'s intentions, the assumptions are common knowledge between $\mathcal{A}$ and $\mathcal{B}$. We will see later why common knowledge is required.

The two sets of assumptions taken together will be called the $\mathcal{BC}$ assumptions. The Background Assumptions hold in the background situation B and the Circumstantial Assumptions hold, of course, in d. B is a part of d.

$\mathcal{A}$ and $\mathcal{B}$ need not be persons; they can be suitably equipped artificial agents.

My claim then is that if all the $\mathcal{BC}$ conditions above are satisfied, $\mathcal{A}$ will *communicate* p to $\mathcal{B}$.