A Game-Theoretic Account of Implicature

Prashant Parikh*
Center for the Study of Language and Information
Stanford University
Stanford, CA 94305
parikh@csli.stanford.edu

December 9, 1991

Abstract

I use game theory, decision theory, and situation theory to model a class of implicatures. Two types of relevance are distinguished and used to construct a model of Gricean communication between speaker and addressee.

1 Introduction

Grice's[5] account of nonnatural communication includes as a corollary an account of how more can be communicated than what is said.[1] Conversely[2], the Gricean principle of cooperation is essential to the full range of nonnatural communication, not just to implicature.

Much information can be extracted from an utterance relative to constraints an addressee knows or is attuned to, but not all this incremental information[3] can be said to be communicated. Implicature is, roughly, communicated information beyond what is said. In this paper, I develop a game-theoretic account of a subclass of implicatures within the framework of situation theory[4]. A formal account of the model developed here together with a derivation of important properties of implicature and the Gricean maxims can be found in Parikh [13].

2 On Relevance

There appear to be two distinct notions of relevance. The first notion, which I will call relevance$_l$ (or logical relevance), is that of a relational property of a proposition with respect to a set of propositions that,

*I would like to thank David Israel and Rohit Parikh for helpful conversations.
[1] See Grice [7].
[2] A thorough discussion of this can be found in Parikh [14].
[3] See Israel and Perry [9].
[4] See Barwise and Perry[4] and Barwise[2].

intuitively, establishes a 'connection' between them. Sperber and Wilson [16] have attempted to capture this property by their idea of contextual implication. Something is contextually implied by a proposition with respect to a set of propositions just in case it is 'nontrivially' implied by both jointly without being implied by either singly. They go on to define degrees of relevance in terms of how many contextual implications a proposition has with respect to a set of propositions. I will adopt their first idea and drop the second one. Thus, a proposition will be relevant$_l$ with respect to a set of propositions if the two together have a 'nontrivial' implication without having it individually.

The second notion relevance$_r$ (or rational or rationality-based relevance) is also a relational property of a proposition but with respect to an agent's goals. It is nothing but the 'value of information' of a proposition with respect to an agent's goals in decision analysis. That is, a proposition is relevant$_r$ just in case it has a positive value. It is this concept that readily admits of degrees of relevance. A proposition is more relevant$_r$ than another if it has greater value. Let me illustrate with an example.

Assume that $\mathcal{A}$ and $\mathcal{B}$ have to attend a talk at 5 pm and that $\mathcal{B}$ believes the probability of it's being time for the talk is .2 and the probability of there still being time is .8. Assume further that his private decision problem is as shown in Figure 1(a). The square nodes are decision nodes, the circular nodes are chance nodes.[5]

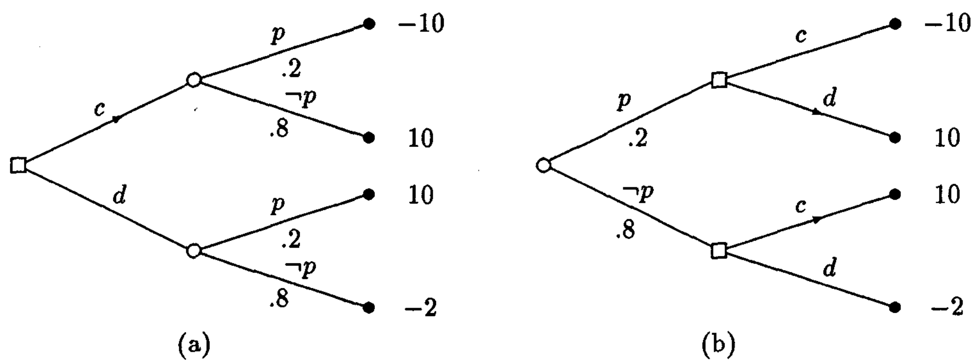

Figure 1: $\mathcal{B}$'s private decision problem

Clearly, his optimal choice is to continue what he is doing.[6] How much does he value complete information about the time in such a context? A priori, he would expect to be told that it was time 20% of the time. In each case, he would make a choice with complete foreknowledge. This would change his expected value to

[5]Key for Figure 1:
Actions:
c: continue previous actions
d: depart for the talk

Propositions:
p: it is time for the talk
$\neg p$: it is not yet time for the talk

[6]Continuing has an expected value of $.2(-10) + .8(10) = 6$ and departing has an expected value of .4.

10 as shown in Figure 1(b). And the value of this information to him would be $10 - 6 = 4$.

Knowing whether it's time for the talk is certainly relevant$_r$ to him. On the other hand, knowing *simply* that it is 5 pm is not relevant$_r$ to him. This is because it does not change his actions (or his probabilities) in any way and results in a value of $6 - 6 = 0$. He needs to infer that it is time for the talk jointly from the proposition that it is 5 pm and his belief that 5 pm is the time for the talk. However, the proposition that it is 5 pm is relevant$_l$ with respect to the larger situation he is in because it permits this inference to the relevant$_r$ proposition p.

It seems possible to construe relevance$_r$ as a special case of relevance$_l$ obtained by restricting contextual implications to goal (and action) related implications. So, in the situation above, p is also relevant$_l$.

Both types of relevance are relevant for understanding implicature. The first type, relevance$_l$, is required to constrain what is said, so that it can form a basis for selecting implicatures among relevant$_r$ propositions.

3 An Example

The game-theoretic model I use to account for implicature, called the Strategic Implicature Model (*SIM*), is similar to the Strategic Discourse Model, extensively developed in Parikh [14] and Parikh [11]. I will develop the *SIM* informally first through an example.

Assume as before that $\mathcal{A}$ and $\mathcal{B}$ have to attend a talk at 5 pm and $\mathcal{A}$ utters the sentence φ below.

'It's 5 pm' (φ)

In so doing, she succeeds in implicating something like 'It's time for the talk' to $\mathcal{B}$. This is an example of a standard implicature as opposed to a flouted implicature.[7] I will use the *SIM* to model this situation.

Assume there is a situation called the background B that contains the assumptions that it is common knowledge[8] that $\mathcal{A}$ and $\mathcal{B}$ are cooperating rational agents and share a reasonably expressive and efficient language $\mathfrak{L}$.[9] B is embedded in a larger situation d called the discourse situation.

From d, we isolate a subsituation s that supports the facts that $\mathcal{A}$ intends to communicate p to $\mathcal{B}$, that their shared goal is to attend the talk, and that they share the belief that the time for the talk is 5 pm. In s, $\mathcal{A}$ can choose to utter a range of sentences from $\mathfrak{L}$ to realize her intention. See Figure 2.

We start with $\mathcal{A}$ uttering φ. If $\mathcal{A}$ utters φ, s results in t. In t, $\mathcal{B}$ has to choose an interpretation of $\mathcal{A}$'s utterance. The obvious interpretation is the literal one, that it is 5 pm. Call it l. Since l is a possible interpretation, $\mathcal{B}$ must assume that the initial situation could be s' instead of s, where $\mathcal{A}$'s intention is to communicate l. I will presently argue that p could also be a possible interpretation, so that $\mathcal{B}$'s choice situation is partially as shown above. These possibilities are common knowledge between $\mathcal{A}$ and $\mathcal{B}$ so that both share this structure of choices.

Note that getting to l itself requires a strategic inference, that is, a game but I will abstract from that here. This is discussed in detail in Parikh [14].

As we saw above, l has no value for $\mathcal{B}$. And since $\mathcal{A}$ and $\mathcal{B}$ are cooperating, conveying l and having $\mathcal{B}$ consider it by itself has no value for $\mathcal{A}$ either. On the other hand, producing and processing φ have some cost. I will assume that each costs 1 unit. $\mathcal{A}$'s total cost will be 2, $\mathcal{B}$'s total cost will be 1.[10] Note that $\mathcal{A}$ incurs a processing cost because she has to consider $\mathcal{B}$'s processing, not because she herself needs to understand the

[7] See Levinson [10].
[8] See Barwise [1].
[9] Precise definitions for the latter terms can be found in Parikh [11].
[10] There are presumably some perception costs for $\mathcal{B}$ but I will ignore them here.

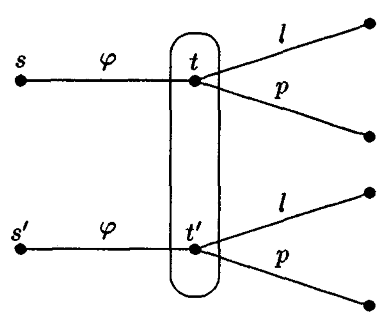

Figure 2: First Stage of Game

utterance. I will assume utilities for both agents are risk-neutral and additively separable so the net value of uttering φ is $0 - 2 = -2$ for $\mathcal{A}$ and similarly -1 for $\mathcal{B}$. But this is so only if s' is factual rather than s. That is, this is so only if $\mathcal{A}$'s intention is to communicate l, not p. The interpretation has to be correct. If the interpretation is incorrect, I will assume that its value is uniformly -2 independently of the proposition concerned. Also, I will assume these costs and values are common knowledge. The choice tree now looks as in Figure 3.

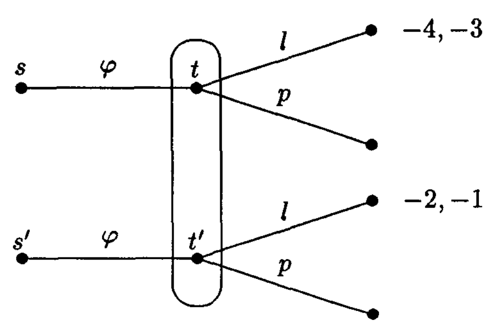

Figure 3: Second Stage of Game

It is possible to complicate this model further by providing $\mathcal{A}$ with a model of $\mathcal{B}$'s decision problem. $\mathcal{A}$ must have observed $\mathcal{B}$'s inaction (with respect to leaving for the talk) and so chose to utter φ. This means $\mathcal{B}$'s payoffs and probability distribution must have been such as to make inaction optimal. But the details of the model may differ from $\mathcal{B}$'s. I will simply assume that they have common knowledge of $\mathcal{B}$'s model.[11]

I will now argue that $\mathcal{B}$ has at least one other choice of interpretation p. The cost of communicating p for both would be increased by the additional processing involved. p has to be contextually inferred from l and the propositions known in s (more generally, the propositions accessible to them in their mutual background).

[11] It is also possible to model $\mathcal{A}$'s private decision problem based on her goals and observations. But I will also ignore this dimension in this paper.

Assume this additional cost is 1 for both. As we saw above, the value of p for both is 4. This implies that the net payoff for $\mathcal{A}$ is 1 and for $\mathcal{B}$ is 2. In other words, p is relevant$_r$ for both. This means that p is a possible interpretation of φ. Of course, the payoffs above occur when the interpretation is correct, that is, when the initial situation is s. The game tree now looks as follows.

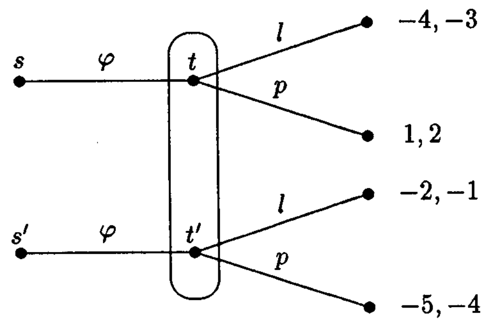

Figure 4: Third Stage of Game

In general, both speaker and addressee will need to consider all interpretations derivable from l with respect to their mutual beliefs (or more generally, their mutual background of what is taken for granted) that have a positive value.[12] However, l must also be considered as a base value against which possible implicatures are to be compared. It is quite possible for l to have a positive value and a positive payoff. This will certainly be the case when there are no implicatures, but it can also happen when there are implicatures.

Let's assume p is the only alternative for now. There are thus two choices for $\mathcal{B}$, l and p. Given what $\mathcal{B}$ can perceive, (and assuming that the relevant perceptual acts go through without a hitch), $\mathcal{B}$ could well be in situation t'. That is, given $\mathcal{A}$'s utterance of φ, $\mathcal{B}$ cannot distinguish between t and t'. (The oval enclosing t and t' represents this epistemic constraint.) From this $\mathcal{B}$ can infer that $\mathcal{A}$ intends to communicate either p or l, (or equivalently, that s or s' is factual), but has no apparent basis to decide between them. If he is in t, p is the preferred choice, and if he is in t', l is the preferred choice. As we saw above, these preferences are derived partly from the relevant intentions and partly from their private decision problems. That is, if he is in t, it is possible to infer that s is factual and that $\mathcal{A}$ intends to communicate p. So, given their net values, p is preferred to l. On the other hand, if he is in t', it is possible to infer that s' is factual and that $\mathcal{A}$ intends to communicate l. So l is the preferred choice. Note that the inference from t to s or t' to s' is not itself strategic *at this level of analysis*. In a wider model, it may be necessary to identify other aspects of the signal strategically as well. However, relative to the problem we are addressing, we can ignore these strategic aspects, just as we did with the literal content l.

Though we have already calculated payoffs, getting there actually requires the further Gricean assumption of rational cooperation between speaker and addressee. Both $\mathcal{A}$ and $\mathcal{B}$ need to have common knowledge of being cooperating rational agents to choose preference orderings that are compatible with $\mathcal{A}$'s intentions. In fact, it may be that Grice's Cooperative Principle can be made precise (in such contexts at least) by identifying it with the requirement that both agents order their preferences according to their mutual beliefs

[12] There are interesting cases to consider where the set of choices for $\mathcal{A}$ and $\mathcal{B}$ differ but we will assume they are identical for now.

about the speaker's intentions.

If this were all, the ambiguity would be ineliminable. This is perhaps the critical place (if such could be isolated) where the agents' rationality and their sharing an efficient and expressive language come in. $\mathcal{B}$ can reason that $\mathcal{A}$ could have chosen a direct but costlier sentence like μ in the case of p, or silence, represented by ν in the case of l. That is, $\mathcal{B}$ is led ineluctably to consider what $\mathcal{A}$ might have said but chose not to.

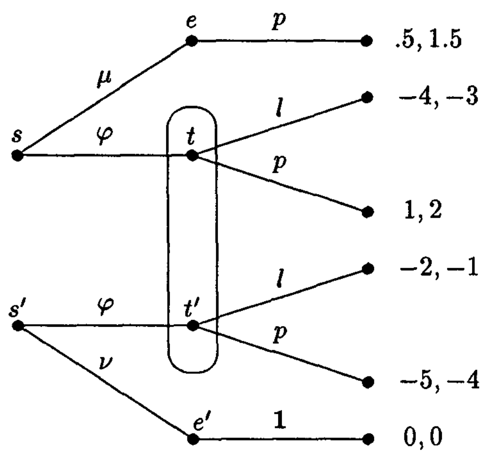

Figure 5: The Local Game $LG(\varphi)$

If $\mathcal{A}$ had chosen μ, s would have resulted in e. I will assume that μ has just one relevant$_r$ content p, its literal content. $\mathcal{B}$'s only choice is therefore p. The sequence of actions in this relatively unambiguous case is clear. The payoffs are somewhat reduced by the additional cost of producing and processing μ. In the case of silence, which is always a choice, the right interpretation is the empty one,[13] and the payoffs for both agents are 0.

Thus, $\mathcal{A}$'s initial utterance of φ leads $\mathcal{B}$ to consider the full structure $LG(\varphi)$ in Figure 5. Note that $\mathcal{A}$ has access to everything $\mathcal{B}$ used to arrive at the local game $LG(\varphi)$. In fact, it is easy to see that they come to have common knowledge of $LG(\varphi)$.

There are two further points to observe here.

First, we need to say how the game is solved. For this we need a few definitions. An information set of a player is a set of (decision) nodes that he or she cannot distinguish. In this game, $\mathcal{A}$'s information sets are $\{s\}$ and $\{s'\}$, and $\mathcal{B}$'s information sets are $\{t, t'\}$, $\{e\}$, and $\{e'\}$. A player's strategy is a function from his information sets to some action available in that set. Thus, $\mathcal{A}$ has four strategies and $\mathcal{B}$ has two. A joint strategy or strategy profile is a tuple that combines strategies of the two players. There will be eight strategy profiles.

One way to solve a game is by using the concept of a Nash equilibrium. A strategy profile is a Nash equilibrium if no player has an incentive to deviate unilaterally. Unilateral deviation by a player is deviation

[13]Here again I'm assuming this is the only relevant$_r$ content. The notation '**1**' comes from Barwise and Etchemendy [3]. '**1**' stands for the sup of the entire infon lattice under consideration, assuming it is complete. Note that I am not distinguishing between propositions and infons in this paper. Propositions just are infons here.

keeping the strategies of other players fixed. This implies that $\{(s,\varphi),(s',\nu);(\{t,t'\},p)\}$ is the unique Nash equilibrium.[14] This shows that the optimal interpretation of φ is indeed p. In other words, solving the game yields the correct interpretation of the utterance, whether literal or otherwise. Note that the way the game is set up, in particular the way the payoffs are assigned, ensures the Gricean recognition of communicative intention as a pivotal step in the disambiguation process.

Next, we must point out that this is not quite the full game that $\mathcal{A}$ and $\mathcal{B}$ play. $LG(\varphi)$ is the structure that results *after* $\mathcal{A}$ utters φ. Without this utterance, $\mathcal{B}$ cannot construct anything. This is what makes the game constructed here different from traditional games of incomplete information, where whatever knowledge of the choice sets of players is considered is available before the game gets played. We call it a game of *partial* information to mark this new feature. Before $\mathcal{A}$ utters φ, she must consider her choices. For each choice ψ, $\mathcal{B}$ can construct a corresponding local game $LG(\psi)$.

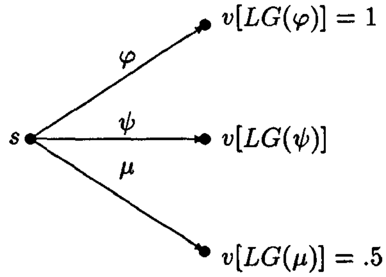

Figure 6: Global Game GG

Assume that each local game has a solution and determines a payoff. Then $\mathcal{A}$'s choice problem is first to figure out the value of each local game, and then to choose the game with the best value and perform the action prescribed by the solution to this optimal local game. For example, if $\mathcal{A}$ considers only φ and μ, then their respective values for her are 1 and .5, and $\mathcal{A}$ should choose to play $LG(\varphi)$, (rather than $LG(\mu)$) and accordingly should optimally utter φ. $\mathcal{B}$'s task is simpler. He perceives the utterance of φ and needs only to construct and play the optimal local game chosen by $\mathcal{A}$, namely $LG(\varphi)$. The larger game, that contains all the local games, we call the global game GG.[15] If LG denotes the set of all local games that result from $\mathcal{A}$'s choices, then we call the pair $\langle GG, LG\rangle$ the Strategic Implicature Model, or simply the *SIM*.

This completes my construction of the basic model. There is one more thing we need to check. We need to ensure that the model does not allow too much information to be implicated. This is an important problem that plagues most accounts of the reasoning that supports implicatures.

Suppose there is some other proposition q that is more informative than p. It will certainly be relevant,

[14] See Parikh [14] for details.

[15] Thus, local games are generated locally by each possible utterance, whereas global games combine all the local games into a single global structure. The size of the global game depends on the number of choices considered, and other contextual factors must be called upon to constrain this choice.

but it is easy to see that it cannot be more relevant$_r$ than p (in the context of B's private decision problem). Moreover, it is reasonable to assume that the greater the information the more costly it is to process. I will assume that this additional cost is one unit for both agents. This gives us the game in Figure 7 below.

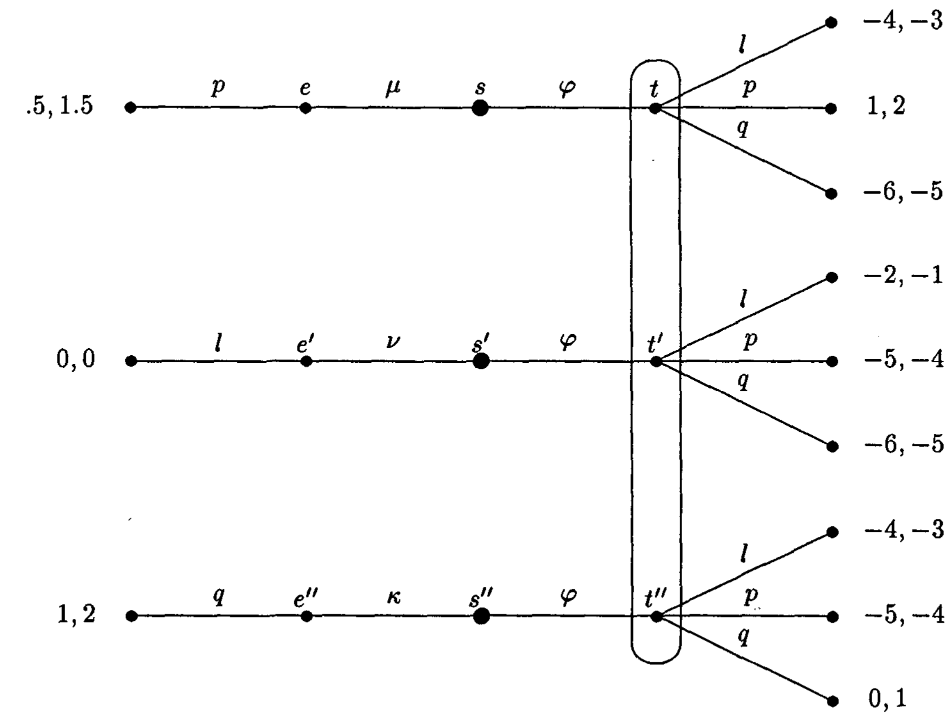

Figure 7: The Local Game $LG(\varphi)$ with q

It is easy to verify that the unique Nash equilibrium of this game is consistent with that of the smaller game in Figure 5 above. That is, both yield the implicature p.

The two-tiered public-private model accounts for the *communicative* aspect of implicatures that is also missing from current accounts of implicature. It is not enough for a proposition to yield maximal value (or net value) to the addressee for it to be the optimal interpretation of the utterance. It must be the solution of a game between speaker and addressee. That is, it must also be optimal from the speaker's perspective, or in other words, be the intended interpretation. It is this joint optimality that is characteristic of implicatures, as it is of all communication. The model also shows how and why errors often occur in interpreting implicatures. If the addressee inadvertently restricts himself to his private decision problem(s), then it becomes possible for him to infer propositions that may have large positive value for him, but that the speaker may have no clue about. Or if the speaker incorrectly models the addressee's private decision problem, then the addressee may have no way to make the right inference. No implicature is possible unless speaker and addressee share adequate knowledge of the addressee's private decision problem.

This account is able to model a large class of implicatures. As I pointed out above, the precise nature of the reasoning involved avoids the problem of inferring too much information as implicated information. In most accounts, including Grice's, the reasoning is supposed to involve a step where the implicated information has to be assumed as part of the content in order to maintain the assumption of cooperation. My model shows that such a step is unnecessary for a large class of implicatures. The implicated content falls out directly from the strategic reasoning and private decisions that the two agents participate in. Besides, there is no indication in such accounts of where the implicated content comes from, under what conditions it exists, and whether it is unique. The formal version of this model in Parikh [13] provides a neat solution to these fundamental questions. I will call the class of implicatures which the model addresses *type I* implicatures.

In type I implicatures, the addressee's private choice situation is directly affected. In what I will call type II implicatures, an implicature adds to the information or beliefs of the addressee without directly influencing any immediate choice of action. The value of the additional information does not lie at the first-order level of action but at the second-order level of belief. This type of implicature can be modelled in more or less the same way except that we need to consider preferences for information directly, rather than via action. Many generalized implicatures, including scalar implicatures, appear to form a special subclass within this type. I pursue this topic in Parikh [12].

References

[1] Jon Barwise. Three views of common knowledge. In Moshe Vardi, editor, *Theoretical Aspects of Reasoning About Knowledge II.* Morgan Kaufmann, Los Altos, CA, 1988.

[2] Jon Barwise. *The Situation in Logic.* CSLI Publications, Stanford, 1989.

[3] Jon Barwise and John Etchemendy. Information, infons, and inference. In Robin Cooper, Kuniaki Kukai, and John Perry, editors, *Situation Theory and Its Applications, I.* CSLI Publications, Stanford, 1991.

[4] Jon Barwise and John Perry. *Situations and Attitudes.* MIT Press, Cambridge, 1983.

[5] H.P. Grice. Meaning. *Philosophical Review*, 66:377–388, 1957.

[6] H.P. Grice. Utterer's meaning and intentions. *Philosophical Review*, 78:147–177, 1969.

[7] H.P. Grice. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Academic Press, New York, 1975.

[8] H.P. Grice. Logic and conversation. In *Studies in the Way of Words*, pages 1–143. Harvard University Press, Cambridge, Massachusetts, 1989.

[9] David Israel and John Perry. What is information? In P. Hanson, editor, *Information, Language, and Cognition: Vancouver Studies in Cognitive Science*, volume 1. University of British Columbia Press, 1989.

[10] Stephen C. Levinson. *Pragmatics.* Cambridge University Press, Cambridge, 1983.

[11] Prashant Parikh. Situations, games, and ambiguity. In Robin Cooper, Kuniaki Kukai, and John Perry, editors, *Situation Theory and Its Applications, I.* CSLI Publications, Stanford, 1991.

[12] Prashant Parikh. Scalar implicature. In progress, 1992.

[13] Prashant Parikh. A theory of implicature. Submitted for publication to Linguistics and Philosophy, August 1991.

[14] Prashant Parikh. Communication and strategic inference. *Linguistics and Philosophy*, 14:473–514, October 1991.

[15] D. Sperber and D. Wilson. *Relevance: Communication and Cognition.* Blackwell, 1986.

[16] D. Wilson and D. Sperber. On defining relevance. In Richard Grandy and Richard Warner, editors, *Philosophical Grounds of Rationality*, pages 243–258. Clarendon Press, Oxford, 1986.