Grice, Hoare and Nash

Pragmatics, Program Semantics and Game Theory

Rohit Parikh

Brooklyn College and CUNY Graduate Center
City University of New York

Johan's Celebration
September 27, 2014

When a diplomat says yes, he means maybe.

When a diplomat says maybe, he means no.

When a diplomat says no, he is not a diplomat.

Attributed to Voltaire.

What is Implicature?

Implicature is a component of speaker meaning that constitutes an aspect of what is meant in a speaker's utterance without being part of what is said. What a speaker intends to communicate is characteristically far richer than what she directly expresses; linguistic meaning radically underdetermines the message conveyed and understood. Speaker S tacitly exploits pragmatic principles to bridge this gap and counts on hearer H to invoke the same principles for the purposes of utterance interpretation. Horn [8].

Some Relevant Quotes

1. **Alan:** Are you going to Paul's party?
Barb: I have to work.

2. (a) He is an Englishman; he is, therefore, brave.
(b) His being an Englishman implies that he is brave.

3. (a) Some athletes smoke
(b) Not all athletes smoke.

4. **Motorist:** *My car is out of gas.*
Pedestrian: *There is a gas station around the corner.*

Cooperative Principle. Contribute what is required by the accepted purpose of the conversation.

Maxim of Quality: Make your contribution true; so do not convey what you believe false or unjustified.
Maxim of Quantity:. Be as informative as required.
Maxim of Relation: Be relevant.
Maxim of Manner: Be perspicuous; so avoid obscurity and ambiguity, and strive for brevity and order.

Theoretical Definition: S conversationally implicates p iff S implicates p when:

(i) S is presumed to be observing the Cooperative Principle (cooperative presumption);
(ii) The supposition that S believes p is required to make S's utterance consistent with the Cooperative Principle (determinacy); and
(iii) S believes (or knows), and expects H to believe that S believes, that H is able to determine that (ii) is true (mutual knowledge).

We will be concerned here with particularized conversational implicatures. These are implicatures where the implicature is not part of the meaning of an utterance, and where the implicature is not conventional in that it depends on *particular circumstances.*

Such implicatures are cancellable and non-detachable. Cancelleble in the sense that someone making a statement **s** which has implicature **p** may then go on to deny **p**. Thus the person saying "There is a garage around the corner" may go on to say, "But as far as I know, it is not open at this time." It is *non detachable* in the sense that another expression with the same truth conditions would have the same implicature.

These implicatures seem to arise in (at least) two ways. One where the purpose of the speaker is merely to convey information and no action or procedure is in the offing. Many scalar implicatures have this form. So if I say that some boys came and thereby implicate that not all boys came, there need be no contemplated action involving the presence or absence of boys.

But many other implicatures occur in the context of a goal which one person has, perhaps even a contemplated action to achieve that goal, and where the other person is volunteering information relevant to the goal. Indeed many of Grice's own examples have this form where the context is one where one person **A** wants to achieve some goal and the other person **B** makes a statement which helps to enable this goal, or, on occasion, causes him to abandon the goal.

Grice's motorist example is of this nature where the goal of the motorist is to fill his car with gas. The mention of the garage is in the context of this goal.

Motorist: *My car is out of gas.*
Pedestrian: *There is a gas station around the corner.*

The implicature is that as far as the pedestrian knows, the gas station is open.

To see how **A**'s **goal** matters, let us change the context of Grice's motorist example just a bit.

A bank robber **R** has hijacked the motorist's car and orders him to drive to some getaway point, or even to an isolated location where the motorist can be killed without much fuss.

At this point it is noticed by both the burglar and the motorist that the gas gauge reads empty. At the robber's behest the motorist **A** stops the car and says to a pedestrian **B**, "My car is out of gas."

The pedestrian has already seen the bank robber's picture on TV and understands the motorist's quandary. In this context, the motorist does *not* want to fill his tank with gas and the pedestrian realizes this. The statement "There is a gas station around the corner" no longer implicates that the gas station is open. It might well implicate that it is closed, or perhaps that there is a police car stationed near the gas station.

Information versus Action

In this talk, we shall consider the cases where the first speaker **A** has some goal in mind which is common knowledge, and the second speaker **B** makes a statement which is relevant to this goal. In that case the (common) knowledge of the goal is part of the context and is typically used to calculate the implicature.

Prashant Parikh and Benz and van Rooij ([13, 2]) do point out that some implicatures might *help* someone to *make a decision* even though other implicatures might just be informational.

It will turn out that some contributions made by the computer scientist Tony Hoare and by the economist John Nash will be relevant. The connection of implicatures with these two eminent scholars seems to have been overlooked thus far in the literature.

Hoare Semantics

Here is how Hoare semantics enters. A Hoare assertion takes the form

$$\{X\}\alpha\{G\}$$

where X, G are propositions, X is the pre-condition, G is the goal and α is a contemplated action (actions are *programs* for Hoare). Formally, S is a state space, X, G are subsets of S, and α becomes a relation R_α on S.
Then the Hoare condition is

$$(\forall s)(\forall t)(s \in X \ \& \ (s,t) \in R_\alpha \rightarrow t \in G)$$

If X holds when the program starts, then G will hold when it finishes.

Now **A** wants to reach G using α. We will assume that G is common knowledge between **A** and **B**, and **B** has some information which would affect the possibility of reaching G.

This could happen in four ways.

1. **B** has *part of* the information which implies a suitable X. Thus **B** is supporting both the goal and the action.
2. **B** has information which would cause **A** to modify α in some way (e.g. to replace it by a more specific action, or, technically, a sub-action.)
3. **B** has information which *suggests* a particular action α to **A**.
4. **B** has information which would cause **A** to *abandon* the method α. (Which may mean abandoning the goal altogether or using some completely unrelated method.)

B volunteers information which indicates whether 1) or 2) or 3) or 4) is the case and leaves **A** to

1. Either supply some *other information* which will complete the process of deducing X or
2. modify α in some way, or
3. Conclude that $\neg\{X\}\alpha\{G\}$.

This *other information* or the modification, or $\neg\{X\}\alpha\{G\}$ is the implicature.
And note that the modification is not in itself a *proposition*, it may be an imperative. Thus part of our thesis is that while an implicature may be a proposition, it might well be something else which affects **A** in some way.

The Cooperative Case

Here we assume, as Grice does, that the utilities of **A** and **B** are in accord. We do not assume that they are the same since what **A** gains from receiving the information is likely to be much more than the pleasure that **B** gets from helping out.

Consider the case of the motorist. The goal of the motorist is clear. She wants to fill her tank with gas. The statement which **B** makes points to an algorithm, *go around the corner and get gas there.* Here it is unclear whether **A** should *walk* there carrying a gas can, or has at least enough gas to *drive* around the corner. But in either case, *going around the corner* is indicated (by **B**) as the α.

Our Hoare assertion is

$$\{X_1 \& X_2 \& X_3\} \alpha \{G\}$$

Where X_1 is that there is a gas station around the corner, X_2 is that the gas station is open and X_3 is that **A** has enough money to pay for the gas.
Only **A** knows whether X_3 is true and this is not **B**'s business. **B** has volunteered X_1. It is obvious that α is useless unless X_2 is true. Since **B** has indicated that he supports the action α it follows that X_2 is true as far as **B** knows. If he is not supporting the action then he should not have said X_1. So in this case X_2 is the implicature.

The Hiring Problem

In this example **A** is the chair of the hiring committee for some college, **B** is a professor and **C** is the professor's student who is an applicant for a position at **A**'s college. **B** writes about **C**
He has excellent handwriting and he always came to class on time.
Here **A**'s goal **G** is to have a colleague who will be a good teacher and researcher, α is the action of hiring **C**, and the precondition is that **C** *is a good philosopher.*
B's statement does *not* give **A** the information which **A** needs, or even supports it in some relevant way, and the implicature is that the Hoare assertion is false. **A** should abandon α. The implicature is $\neg\{X\}\alpha\{G\}$ where X is any true condition.

But why doesn't **B** simply say, "He is not a good philosopher"? Clearly because **B** is **C**'s teacher and professional ethics preclude him from saying something negative about **C**. What he does instead is to say something positive which is not good enough.

Here is a joke which makes a similar point.

> *A tired and depressed looking man walks into a restaurant and sits down. A waiter comes over and asks what the man wants.*
>
> *"Two scrambled eggs with rye toast, and a kind word," says the man.*
>
> *After a while the waiter comes back and puts an order of eggs and rye toast before the man.*
>
> *As the waiter is walking away, the man says, "What about the kind word?"*
>
> *"Don't eat them eggs," says the waiter.*

Just like the professor, the waiter is forbidden to say something negative about the restaurant. but his "don't eat them eggs" carries the implicature that the eggs are not good, may even cause illness. And that indeed is a kind word.

Modifying α

We note that in many cases actions are not disjoint from each other. If we think of a (nondeterministic) action as a binary relation on the state space, then two actions may be disjoint, may overlap, or one may be included in the other.
One Hoare-like rule is

$$\frac{\{X\}\alpha\{G\},\ \beta \subseteq \alpha}{\{X\}\beta\{G\}}$$

If a correctness condition is satisfied by α then it is also satisfied by a subaction β but not necessarily vice versa.

Here is an example.
A to **B**, "I am thinking of going to Times square by public transport."
B, "Buses will be very slow during the rush hour."
A likely implicature is "Take the subway."
Here action α is the action of taking *some* public transport, β is the action of taking the subway and γ is the action of taking a bus. α is the union of β and γ.

Let X be the current situation, and G be the goal of getting to Times square on time. Then $\{X\}\beta\{G\}$ is true, but $\{X\}\alpha\{G\}$ is not

This is a consequence of the nondeterminism of the two actions. β is *guaranteed* to achieve the goal whereas α might but is not guaranteed to do so.. **B** is suggesting that the action be changed from α to β by eliminating γ.

When an action satisfies a Hoare condition then so does a sub-action. But this is not the case if we are trying to maximize expected utility. It is quite possible that the expected utility of α is higher than that of β even though $\beta \subseteq \alpha$.

For example, if I am betting on a horse, then it is better to choose a horse at random than to choose a specific horse which is well known to be a nag.
However, if *satisficing* is our condition then subactions would be at least as good as an action. If all outcomes of α are satisfactory, and β is a subaction of α then all outcomes of β are also going to be satisfactory.

Scalar Implicatures

It has been argued that when there is a sequence of statements $s_1, s_2, ...s_n$ which **B** could make such that s_{i+1} is stronger than s_i and **B** makes statement s_i, then s_{i+1} is not true. Thus if **B** is asked whether any girls from the class took the exam and **B** says, "*Some* girls did", then the implicature is that not all girls did since if **B** knew that they all did then the cooperative principle required him to make the stronger statement.

We do not consider for the moment the possibility that **B** does not know that all did, or for some reason considers it wiser not to say. The statement "Jill has two children" can be taken to mean that she has at least two and would not be false if she had three. Yet, the statement "Jill has two children" is conventionally taken to mean that she has exactly two. If two statements X and Y are both available, Y is stronger than X and the speaker says X, then it is often taken for granted that the stronger statement X is either false, or not known to the speaker to be true. But there are exceptions.

Suppose Professor Jones has been accused of sexual harassment and I ask, "Who is the accuser?" If I am told, "She is an undergraduate in Philosophy," it does not follow that my informant does not know who she is. For here my utilities and the utilities of my informant are not quite in accord. I want to satisfy my curiosity to the maximum extent, but my informant wants to protect the accuser's privacy. It is possible that Grice's "He is somewhere in the south of France" also follows this pattern where the utilities are not quite in accord.

But such implicatures can go in opposite directions and the direction is determined by the Hoare assertion we are considering. For instance, suppose a playground has the sign, "An adult entering must be accompanied by a child," then it is understood that an adult with two children will be allowed, but an adult with zero children will not. This is because the goal of having the playground primarily for children is served by an adult who brings two children but not by an adult who comes alone.

But consider a different example. Suppose an advertisement for a movie says, “An adult paying full price may bring two children along at half price.” Here the number of children can be reduced to one, and an adult may even come without bringing any children at half price. But if an adult wants to bring three children then she must pay full price for the third child.

So in the playground case “a child” meant “one or more children” whereas in the movie case, “two children” means two or fewer children.

The fact that the scalar implicatures go in opposite directions is a clear example of the relevance of the goal G in question.

Of course if the situation is purely informational and **B** is merely informing **A** without having a clear sense of the reason why **A** needs to know, then **B** has the obligation to convey the most accurate information available to her, excepting of course cases like that of sexual harassment referred to earlier.

Nash Bargaining

In Grice's treatment of implicature, he assumes a principle of cooperation. Thus for one of his first examples, when a motorist says, "My car is out of gas" and the pedestrian replies "There is a gas station around the corner," there is an implicature that the station is open. And this follows from the presumption that the pedestrian's desires are the same as those of the motorist, although perhaps less intense and so the pedestrian *wants* the motorist to get gas for his car. This tradition has been followed in much of the subsequent literature.

However, there are exceptions. The economics literature on cheap talk [5] no longer assumes that the utilities are aligned. What the speaker wants and what the listener wants need no longer be fully aligned, although some overlap is necessary for communication to take place at all. Stalnaker in his paper "Cheap talk and credibility..." [14] follows this tradition as well.

Yet as we noted, some degree of cooperation is requisite, for otherwise why communicate at all?
We would like to suggest a slight generalization of the Grice principle which looks like it might bridge the gap between cooperation and strategizing. This principle was originally formulated by John Nash in his paper "The Bargaining problem."

In Nash's framework two players **A** and **B** are trying to decide on a point in two space. There is a convex set S of possible solutions and each point p in S yields utilities u(p), v(p) to **A** and **B** respectively. Nash presumes that the actual bargain, i.e. the point p which is finally chosen will be Pareto optimal. That is to say, Nash assumes that there is no q in S such that $u(q) \geq u(p)$ and $v(q) > v(p)$ or that $u(q) > u(p)$ and $v(q) \geq v(p)$. There is no way to make one person better off without making the other person worse off.

Nash assumes moreover that the space S is convex.
To take an example rather like that of Nash's original example. Suppose that the two are restricted to a point in the set $\{(x, y)|2x + y \leq 3\}$. The utilities are x for **A** and y for **B**. The fallback point is (0,0). Then the Pareto optimal points will be all the points on the line $2x + y = 3$. But which particular point should be chosen? The product of the utilities is maximized at the point (.75, 1.5).
But Nash does not speak about communication and there is no guarantee even that a Pareto optimal point will be reached, let alone Nash's "ideal" point. To take a real life example, it seems highly unlikely that a Pareto optimal point will be reached in Ukraine.

Using very natural axioms on the solution concept Nash proves that the final bargain will be the unique point p such that $u(p) \times v(p)$ is maximum.
It is obvious that assuming that the players are choosing a Pareto optimal point, and there are at least two such, then there is a conflict. Neither can gain without the other losing.

The element of cooperation enters through Nash's notion of a *fallback point*. The fallback point F is the point to which they "fall back" in case they cannot arrive at a bargain, and this point is worse (for both) than any other point in S.
Thus cooperation arises through the fact that both players want to avoid the fallback point and each needs the help of the other to achieve this.
Grice's cooperative principle is a special case of Nash's. For suppose the utilities *are* aligned. I.e., if for any two points p and q we have $u(q) > u(p)$ iff $v(q) > v(p)$, then the Nash bargaining point which maximizes the product $u(p) \times v(p)$ is also the point which maximizes $u(p)$. **B** gains by helping **A** to gain. The pedestrian helps the motorist to get gas for the sake of the small pleasure of helping another[1]
But as we noted this is not the only case. The mere fact that the players both want to avoid X does not imply that their utilities are fully aligned.

[1] See for instance Tomasello, [15].

We now offer an example of how the Nash principle works.
Suppose that an American tourist is in India and wants to buy a carved wooden elephant. He has already seen such an elephant in a store for Rs. 500 but sees a hawker selling the identical elephant for Rs. 400.
It is customary to bargain with hawkers but what should the tourist offer?[2]

[2] In a similar situation, Aumann offered 200, the offer was accepted and Aumann bought the elephant, only to find that the proper price would have been Rs. 50.

In this situation, the fallback situation is that the tourist abandons the hawker and buys his elephant in the store. But the hawker himself has bought the elephant for Rs. 40 and so any price paid from 41 rupees to 499 rupees would be better for both than the fallback situation, which is no sale for the hawker and a cost of Rs. 500 for the tourist.
The element of cooperation arises because both parties want to avoid the fallback situation, but given this fact there is an element of conflict in that the hawker wants to charge more and the tourist wants to pay less.

Here we assume that the utilities of **A** and **B** are not aligned although there must be some concord for communication to take place at all.[3]
Consider the following scenario. You are in small town in India[4] You have hired a rickshaw, a small motorised vehicle and you ask the driver where you can buy some silver jewelry. Now there are in fact many places in Jaipur where you can do that, but the driver will tell you about one where the shopkeeper knows him and will give him a commission.

[3] We do not consider the important and interesting case where **A** *thinks* they are aligned but they are not.

[4] My own experience was in Jaipur.

He will say, "Go to Raja Motimal's Jewlery, they will give you the best quality. And in fact mention my name and they will give you a good price." He may even offer to take you there.
Now it may be, and probably is the case that if α is the action of going to "Raja Motimal" and β is the action of going to another place, say Rani Rukmini's shop, then β has a higher utility for you, while α has the higher utility for the driver.

If X is the condition of being where you are and G is the goal of buying silver jewelry, than both the Hoare assertions $\{X\}\alpha\{G\}$ and $\{X\}\beta\{G\}$ are correct. But α yields a higher utility for the driver and β yields the higher utility for you. But the driver may not tell you about the Rukmini shop and you are stuck with Raja Motimal. The driver may not feel a pang of conscience if the utility of going to Raja Motimal's is higher for you then the utility of simply not buying any jewelry. Perhaps Nash [11] would approve of the "bargain" that the two of you have worked out where each of you made a gain.

Stalnaker, [14] lays out the conditions which will govern the transactions when the utilities are only partially aligned and **B** has various options of what to say. Then **B** will say something which will benefit him as well, but knowing that **A** might be suspicious, he will confine himself to saying something believable.
Here it may happen that **B** and **A** have strategies for speaking and for believing which are in some sort of equilibrium.

For a last example we consider the etchings dialog discussed by Steven Pinker in a youtube video
`https://www.youtube.com/watch?v=3-son3EJTrU`
as well as in a joint *PNAS* paper [12].

A young man and his date have just had dinner at a restaurant near his apartment building.[5] After the dinner the young man says to his date "Would you like to come to my apartment and see my etchings?" Here the implicature might well be, "Would you like to come to my apartment and have sex?" It is not clear whether their utilities are aligned or if they are looking for different Pareto optimal points. Perhaps the young man actually has etchings and perhaps the young woman is interested in sex even though she has no interest in etchings.
What bargain point will they arrive at? It is a mystery.

[5] This insertion about the location of the restaurant is mine so as to make the scenario plausible.

In the full paper we will discuss this issue in more detail and propose a formalism.

Acknowledgements We thank Nicholas Allott, Anton Benz, Luciana Belotti, Michael Devitt, Stephen Neale, Prashant Parikh, Steven Pinker, Adriana Renero, Robert Stainton and Cagil Tasdemir for comments.


Nicholas Allott Game theory and communication *Game Theory and Pragmatics* (Benz, Jäger and van Rooij, editors), Palgrave Macmillan, (2006) pp. 83-100


Anton Benz and Robert van Rooij, Optimal assertions and what they implicate, *Topoi*,


Luciana Belotti and Patrick Blackburn, Conversational Implicatures 2014


Wayne Davis Implicature *Stanford Encylopedia of Philosophy - online*


Joseph Farrell and Matthew Rabin Cheap talk *Journal of Economic Perspectives* vol. 10 (1996), no. 3 , pp. 103-118.


Paul Grice, *Studies in the Way of Words*, Harvard, 1989.


C. A. R. Hoare, An axiomatic basis for computer programming, *Communications of the Association for Computing Machinery*, vol. 12 (1969), no. 10, pp. 576–580.


Laurence Horn Implicature *The Handbook of Pragmatics* (