Conversational Implicatures Via General Pragmatic Pressures*

Christopher Potts

UMass Amherst, Amherst MA 01002, USA
potts@linguist.umass.edu
http://people.umass.edu/potts/

Abstract. This paper aims to show how general pragmatic pressures, interacting with the context of utterance, can produce specific conversational implicatures — as well as the lack thereof in nonsupporting environments. Inspired by the work of Merin (1997), Blutner (1998), van Rooy (2003a), Benz et al. (2005), and others, I use probabilities to represent speakers' belief states and the content of their utterances. These values determine an utterance's *quality rating* and *quantity rating*. I adapt Roberts' (1996) view of the question under discussion to define a *relevance ranking* of utterances. These values come together in a definition of *felicitous utterance*. This definition licenses certain inferences — relevance implicatures relating to the question under discussion (section 4.1) and a variety of quantity implicatures (sections 4.2–4.3).

1 Introduction

Conversational implicatures can be exquisitely sensitive to subtle changes in the context. An utterance might conversationally implicate a meaning M very robustly in one context, but a slight change to that context — a shift in the question under discussion, a revelation about the speaker's belief state, some additional linguistic material — might cause M to vanish.

My goal for this paper is to define a system of pragmatic pressures that can capture such extreme context sensitivity. Limitations of space and understanding lead me to concentrate on a specific subclass of conversational

* I've benefitted from an amazing amount of useful feedback on this general project. I'm indebted in particular to Jan Anderssen, Ash Asudeh, Rajesh Bhatt, Daniel Büring, Donna Byron, Shai Cohen, Regine Eckhardt, Kathryn Flack, Lyn Frazier, Hans-Martin Gärtner, Yurie Hara, Larry Horn, Gerhard Jäger, Pauline Jacobson, Mark Johnson, Ed Keenan, Angelika Kratzer, Manfred Krifka, Eric McCready, Barbara Partee, Tim Roeper, Tom Roeper, David Schueler, Florian Schwarz, Raj Singh, Tim Stowell, Anna Verbuk, Mike White and audiences at the UMass Amherst Semantics Group (Sep 6, 2005), Göttingen University (Sep 21, 2005), ZAS (Sep 23, 2005), UCLA (Dec 1, 2005), Yale (Nov 7, 2005), OSU (Feb 24, 2006), and Brown (Mar 6, 2006). Craige Roberts provided some crucial insights at a crucial moment in my development of these ideas.

T. Washio et al. (Eds.): JSAI 2006, LNAI 4384, pp. 205–218, 2007.

implicatures, namely, those that we can trace back to a specific question under discussion. In this, I derive my inspiration primarily from Roberts (1996, 2004) and van Rooy (2003a), who regard the question under discussion as central both to understanding specific utterances and to following the general flow of conversation.

Section 2 gathers the requisite tools: I define a probabilistic perspective on propositions, then use the numerical values to define approximations of the Gricean maxims of quality, quantity, and relevance. In section 4, I show how this system of pragmatic pressures models conversational implicatures — and the lack thereof, where appropriate. In the closing section, I situate these results within the general theory of conversational implicatures, concentrating on the issue of whether generalized conversational implicatures are derived via the usual pragmatic mechanisms (Horn, 2005) or included in the semantico-pragmatic system as presumptive meanings that are present unless cancelled (Levinson, 2000; Chierchia, 2004).

Throughout, I have tried to give the central definitions as precisely as possible. This can make them appear somewhat complicated and involved, but I think this is just a consequence of the step-by-step format. It's my hope that the definitions are usable, more or less off the shelf, by researchers with computational[1] or mathematical aims as well as linguistic ones.

2 A System of Pragmatic Pressures

As a discourse participant, one could, in principle, say anything at all. The physical constraints are few; one could voice any number of lies and irrelevancies, at whatever length one wished. But this is not usually how things go. We (generally) limit ourselves to what is (basically) felicitous. And the set of pragmatically felicitous utterances at a given point in a discourse is highly constrained when compared with the utterance space at one's disposal.

What pressures shape the space of felicitous utterances? The Gricean maxims are a prominent, historically important answer. Grice identified pressures for truthfulness (quality), informativity (quantity), and relevance, as well as a general pressure to be clear and concise (manner). In this paper, I focus on the first three. The next section moves us from viewing propositions as sets to viewing them as probability distributions. This lays the groundwork for my particular take on the maxims of quality, quantity, and relevance.

2.1 Probabilities and Propositions

The foundational technical move of this paper is a shift in the usual perspective on propositions. We are accustomed to thinking of them as sets of intensional

[1] I have developed basic computational tools for solving problems within this theory and exploring its predictions. The CGI interface is linked from <http://people.umass.edu/potts/computation/>.

indices. But Merin (1997), van Rooy (2003b), Benz et al. (2005), and others have shown how we can fruitfully think of them instead as probability distributions when we are in the realm of utterances:

(1) For countable W, a function $P : \wp(W) \mapsto [0,1]$ is a probability distribution iff:
 a. $P(W) = 1$; and
 b. if p and q are disjoint subsets of W, then $P(p \cup q) = P(p) + P(q)$.

These general probability distributions alone are not quite the right structures, however. We have a correspondence between, e.g., the proposition $\{w_1, w_2\}$ and the probability distribution P such that $P(\{w_1\}) = P(\{w_2\}) = .5$, but we also have probability distributions like $P'(\{w_1\}) = .6; P'(\{w_2\}) = .4$, which do not treat w_1 and w_2 equally and thus do not correspond to any proposition. To bring these two classes of structure more in line, we need to build in an additional requirement:

(2) The probability distribution P mimics the proposition q (a subset of W) iff:
 a. $P(\{w\}) = 0$ iff $w \notin q$
 b. $P(\{w\}) = P(\{w'\})$ for all $w, w' \in q$

Throughout this paper, I assume a dual view. In the semantics, propositions are sets of worlds (with all the important boolean structure this provides; Keenan and Falz 1985; Kamp and Partee 1995). In the pragmatics, propositions are probability distributions (with all the measuring techniques that these facilitate).

2.2 Quality Thresholds

The primary imperative of Grice (1975) is one of truthfulness: the maxim of quality enjoins us to be truthful, to say only those things for which we have adequate evidence. For the purposes of this paper, I factor out the difficult concepts of knowledge and evidence. The approach is instead framed in terms of belief. But this is, as far as I can tell, a move of convenience rather than necessity or substance; the theory of knowledge defined by Kratzer (2002) seems well suited to combination with the present approach.

From this simplified perspective, we can read the quality maxim as an injunction to confine oneself to utterances whose content is entailed by one's belief state (Groenendijk, 1999). Given the probabilistic perspective described above, this would play out as follows. Let P_S be the probability distribution modeling speaker S's belief state. Then we associate *quality ratings* with probabilities for the speaker:

(3) **Quality ratings**
 The *quality rating* of an utterance U by speaker S in context C is $P_S([\![U]\!])$.

The Gricean imperative would then be that a speaker S should confine himself to utterances U such that $P_S([\![U]\!]) = 1$.

In practice, though, we are not nearly this strict. We can be lax on quality, as when we brainstorm new ideas or participate in bull sessions (Frankfurt, 1986). Conversely, we can be quite strict on quality, as when we maneuver to land rockets on the moon or instruct our students (perhaps). This movement towards the ideal — a maximal probability — might never quite finish; if we are extremely skeptical, we might allow that we don't fully believe (know) anything. An element of doubt creeps in, the probabilities drop. Therefore, I propose that each context comes with a *quality threshold* C_τ. This is a numerical value in the real interval $[0, 1]$:

(4) **Quality thresholds**
An utterance U by speaker S in context C satisfies quality iff its quality rating is above the quality threshold C_τ for C.

We make use of these thresholds in a very simple way: I, as a speaker in context C, am forbidden from saying anything with a probability (according to my beliefs) that is below C_τ. The practical effect of this principle is to remove from consideration a great many things that, informative and relevant though they may be, are simply not supported by my epistemic state.

2.3 Quantity Ratings

A simplified version of Grice's maxim of quantity might read: Be informative! The actual statement is more complex than this, including as it does some information about what is required in the current context. But that work is, in the present system, achieved by relevance, so we do not skimp on coverage by purifying the quantity maxim to an unqualified demand for information.

As with quality, I translate quantity into a setting in which we can take measurements, and the probabilities of section 2.1 are again the starting point. We use them to derive a standard value from information theory: informativity values, as defined in (5).

(5) **Information value of p for individual a**

$$\inf_a(p) = -\log_2 P_a(p)$$

I've chosen this particular scaling of probability values simply because it meets the two central demands in this area: (i) high probabilities correspond to low inf values, and (ii) low probabilities correspond to high inf values:

(6) **As probabilities drop,** inf $(-\log_2)$ **values rise**

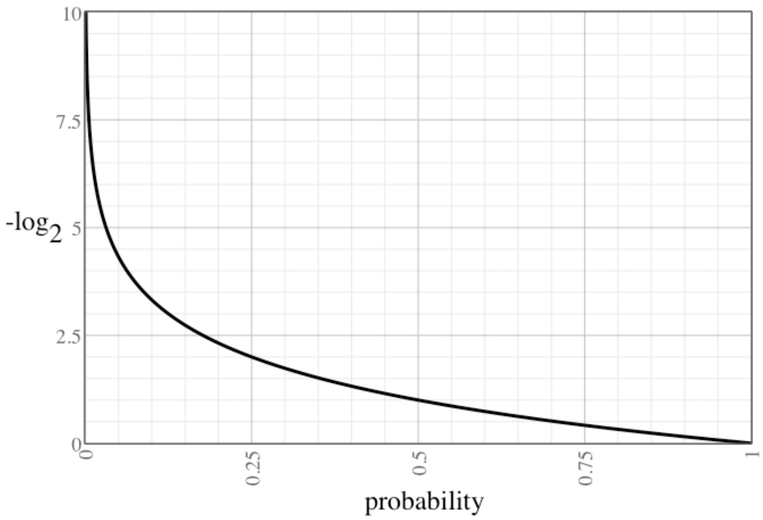

The translation into pragmatic theory is direct: the inf values derived from the hearer's probabilities are the *quantity ratings*.

(7) **Quantity rating**
The *quantity rating* of an utterance U by speaker S to hearer H in context C is

$$\text{Quantity}_C\, U = \text{inf}_H(\llbracket U \rrbracket)$$

Quantity measurements are taken relative to the hearer's belief state, P_H. We cannot define these values directly in terms of the speaker's information state, as this would place quantity in unwanted conflict with the quality maxim: quantity values would favor low probabilities while quality favored high probabilities. The result would likely be silence.

In practice, this hearer orientation means that a speaker must guess about the belief state of his addressee. But quantity values should be measured in terms of the actual belief state of the hearer, not in terms of the speaker's guesses about that belief state. Suppose I guess that you do not know the finer points of training for a marathon, but my guess is wrong: you are an expert on the topic. I begin to tell you about long training runs, proper hydration, and the misery of mile 18, and you grow agitated. What I am saying has very low quantity in this context. You are understandably annoyed: my behavior is infelicitous, and (7) tells us why.

2.4 Relevance Ranking

My approach to relevance brings to the fore two salient features of relevance as an intuitive notion: it is a matter of degree, and it is highly dependent upon the current context. The current system is built up from a gradient foundation (the probabilities), and all its values are derived by integrating information from the context. So it is well suited to providing a formal notion of relevance. (Here, we needn't worry about slighting Grice, as his maxim of relevance simply says "Be relevant".)

In this system, relevance is given as a ranking. Arriving at the ranking turns out to be a bit complex, but the procedure is derived from the insights of Groenendijk and Stokhof (1982) and van Rooy (2003b), and its motivations are intuitive (despite its involved sorting and selection procedure).

The notion of relevance defined here is relevance to a question. I assume a partition semantics for questions: a question contains a set of mutually exclusive alternatives, and asking a question can be construed as asking for an identification of the true alternative. So we get a good measure of the degree to which a proposition p answers a question Q by counting the number of cells in Q with which p has a non-null intersection. This count is the Ans value, as defined in (8).

(8) a. $p_Q = \{q \in Q \mid q \cap p \neq \emptyset\}$ (for p an answer to question Q)
b. $\mathrm{Ans}(p, Q) = |p_Q|$

This is a very partial ordering. One fact about the ordering is particularly important for getting at relevance: if A is a proper subset of B, then A and B have the same Ans values for any question that contains B or one of its supersets as a member. In practice, this means that we cannot yet distinguish a complete answer from a complete and overly informative answer. But overly informative answers are reliably judged to have degraded relevance values (Sperber and Wilson, 1995).

But the system has within it the potential to distinguish A from B in this example: the quantity rating of A will be at least as high as (probably higher than) the quantity rating of B. So we can use quantity to tighten up the ordering of answers relative to a given question:

(9) **Relevance-ranking**

i. Sort the space of utterances with quality ratings above the threshold into equivalence classes based on Ans-values.
ii. For each Ans-equivalence class, get the utterances with the lowest quantity ratings in that class. Keep them, and throw out the rest.
iii. The Ans ordering of the remaining set is the relevance ranking.

The representation in (10) might help convey the sense of (9). (For any number n, $\mathrm{Ans}\,n = \{p \mid \mathrm{Ans}(p) = n\}$.)

(10) **Strikeouts indicate relevance-based eliminations.**

Question Q

$$\left\{\begin{matrix}\{w_1, w_2\} \\ \{w_3, w_4\}\end{matrix}\right\} \Longrightarrow$$

Ans ordering for Q

$$\text{Ans}\,1 = \left\{\begin{matrix}\{w_1, w_2\} & \{w_3, w_4\} \\ \cancel{\{w_1\}}\ \cancel{\{w_2\}} & \cancel{\{w_3\}}\ \cancel{\{w_4\}}\end{matrix}\right\}$$

$$\text{Ans}\,2 = \left\{\begin{matrix} & \{w_1, w_2, w_3, w_4\} & \\ \cancel{\{w_1, w_2, w_3\}} & \cancel{\{w_1, w_2, w_4\}} & \cancel{\{w_2, w_3, w_4\}} \\ \cancel{\{w_1, w_3\}}\ \ \cancel{\{w_1, w_4\}} & \cancel{\{w_2, w_3\}} & \cancel{\{w_2, w_4\}}\end{matrix}\right\} \Longrightarrow$$

Relevance ranking for Q

$$\begin{matrix}\{w_1, w_2\} & \{w_3, w_4\} \\ \searrow & \swarrow \\ \multicolumn{2}{c}{\{w_1, w_2, w_3, w_4\}}\end{matrix}$$

So if a speaker knows that $\{w_1\}$ is the actual world, then he will answer Q with $\{w_1, w_2\}$ — a relevance ranked utterance that scores a 1 on quality. Of course, $\{w_1\}$ is also a 1 on quality. But this overly informative answer is not relevance ranked. (Section 4.1 contains a fuller discussion of overly informative answers.)

3 Felicitous Utterances

It's time now to pool together the above concepts into a single definition of *felicitous utterance*. The backdrop for it is a view of context as tuples:

(11) A *context* is a tuple

$$\left\langle P_S, P_H, Q, C_\tau, U \right\rangle$$

where P_S is the speaker's belief state (probability function), P_H is the hearer's belief state, Q is a question under discussion, C_τ is a quality threshold, and U is an utterance.

The definition of utterance felicity takes the form of a ranking of utterances:

(12) **Felicitous utterances**
The set of felicitous utterances for a partial context $C = \langle P_S, P_H, Q, C_\tau \rangle$ is obtained as follows:

i. From the set of all propositions, eliminate those that have quality ratings at or below C_τ. (See (4).)
ii. With the resulting set, determine relevance rankings and throw out all utterances without such rankings. (That is, throw out every utterance that is not among the least informative members of its Ans-equivalence class. See (9).)

iii. From the resulting set of relevance-ranked utterances, extract the utterances with the lowest Ans values.
iv. From the resulting set, select the utterances with the highest quantity ratings. These are the felicitous utterances for C.

Once again, a slightly more visual representation might help to articulate the workings of this definition:

(13) Let W be the set of all possible worlds, and let $P = \wp(W)$ be the set of all propositions. The partial context is $C = \langle P_S, P_H, Q, C_\tau \rangle$, and we are determining which utterances U are felicitous in C.

Quality elimination
$P_Q = P - \{p \mid P_S(p) \leqslant C_\tau\}$ $\Longrightarrow$

Relevance elimination and Ans-**minimization**
$P_R = \min_{\text{Ans}}(P_Q - \{p \mid p \text{ is not relevance ranked}\})$ $\Longrightarrow$

Quantity maximization
$P_{\text{felicitous}} = \max_{\text{Quantity}}(P_R)$

So P_Q is the set of utterance-contents that satisfy quality. P_R is the set of utterance-contents that satisfy relevance. And $\max_{\text{Quantity}}(P_R)$ is the set of all maximally informative things in the remaining set.

The next section is devoted to showing how we can use this definition to understand a range of conversational implicatures.

4 Conversational Implicatures

The algorithm described in section 3 is fruitfully thought of as a way of solving for the utterance element in our view of contexts:

(14) $P_S, P_H, Q, C_\tau \Rightarrow$ a ranking of utterances

It is illuminating to ask what happens when we fix the utterance U and instead solve for one (or more) of the other elements. The next few subsections do just that, pausing along the way to highlight the positive work that is done at each turn.

4.1 Relevance-Based Implicatures

What happens when we fix the elements P_S, P_H, C_τ, and U, and then ask about the question under discussion Q? The definition in (12) licenses at least one important inference about Q:

(15) Suppose a speaker S has uttered U in context C. Then the set of potential questions under discussion is not larger than the set of all Q such that U has a relevance ranking according to Q.

To make this concrete, imagine the following situation: you are a cooperative speaker. You say, "I live in New York". Then I know that the question under discussion (at least in your mind) is not Q_{country} = *Which country do you live in?*. For suppose it were. Then "I live in New York" would share an Ans value with "I live in the U.S.". By (12ii), it would be ruled out. But you are, by assumption, obeying this restriction. Hence the supposition is false — Q_{country} is not the question under discussion.

Though limited, (15) can make sense of the following dialogue, based on an example from Grice 1975:

(16) A: Is Smith happy at his new job?
B: Yes, and he hasn't even been to jail yet.

B's answer here is not relevance ranked if the question under discussion is given by A's utterance. It is eliminated in virtue of the less informative but complete answer "Yes". Let's assume that B is being truthful, and that his utterance is informative to A (which is basically guaranteed by A's question; see section 4.2). Then B has failed only on relevance. In Grice's terms, he has *flouted* it.

If B is in fact playing by the rules, then it must be because he is speaking relative to a different question under discussion. Some questions that do allow B's reply to have a relevance ranking are "Is Smith both happy at his new job and a free man?", as well as "Is Smith happy at his new job and is his job illegal". And so forth. B's utterance demands that we change the question away from A's. It does not tell us which question to switch to, so there is some indeterminacy. In practice, this seems right: it is clear that B intends something extra with his utterance. But what exactly that something extra is — that is typically unclear.

4.2 Quantity-Based Implicatures of Questions

Quite generally, if a speaker knows the answer to question Q, then it is highly marked for him to ask Q. The present system captures this infelicity by looking to the addressee — the answerer:

(17) If the speaker already knows the answer to his question, then the quantity rating of any felicitous answer, as defined in (12), will be disastrously low.

Quantity ratings are defined in terms of the addressee's belief state. If the addressee assigns p the probability 1, then the quantity rating for p (its inf value, as in (5)) is 0. So if I know p and yet persist in asking you whether-p, then I put you in a very bad spot indeed.

Of course, the answerer could in principle boost the quantity rating of his answer: he could give a false answer, which would have a pathological (∞) quantity rating, or he could say something informative but irrelevant. However, neither move is permitted, by (12). Unsupported or unsupportable claims are eliminated at (12i), and irrelevancies are eliminated at (12ii).

The system has a bit more to say on the matter. At step (12iii), we extract the elements with the lowest relevance ranking. The quantity maximization step, (12iv), is then taken in terms of this small set of elements with equivalent Ans values. One might wonder whether we could simplify the system by maximizing on the full set of relevance-ranked utterances, not just those with the lowest Ans-values, thereby removing step (12iii). However, situations in which the questioner knows the answer to his question help justify the multi-step process.

For instance, suppose A knows that Barbara lives in Moscow but nonetheless asks me "Which city does Barbara live in?". The utterances in (18) are all relevance ranked and have identical quantity values in this context.

(18)
a. "Barbara lives in Moscow." quantity $= 0$
b. "Barbara lives in Russia." quantity $= 0$
c. "Barbara lives on earth." quantity $= 0$

But, intuitively, my felicitous move is (18a) (assuming I assign it a probability above the quality threshold). This follows from (12) and the fact that it has the lowest Ans value of the three.

Potts (2006) offers a fuller discussion of (17) and addresses some prima facie counterexamples to it.

4.3 Quantity-Based Implicatures of Answers

Using (12), we can infer a lot about the speaker's belief state. When we do this, we are essentially asking how much we can infer about the output set in (19).

(19) $Q, C_\tau, U \Rightarrow$ a set of potential belief states for the speaker

That is, we know what the question is, we know what the speaker said, and we know where the quality threshold is. What can we infer about the speaker's beliefs in such a situation? The following is the primary tool for answering this question:

(20) Suppose the speaker S uttered U in a context C with question Q. The set of potential belief states for S is the set of all P_S such that
a. the speaker's utterance is above the quality threshold according to P_S (i.e., $P_S(\llbracket U \rrbracket) > C_\tau$); and
b. the speaker could not have answered A more completely with P_S (i.e., there is no utterance U' such that $P_S(\llbracket U' \rrbracket) > C_\tau$ and $\mathrm{Ans}_Q\llbracket U \rrbracket > \mathrm{Ans}_Q\llbracket U' \rrbracket$).

This is just to say that the hearer can assume that the speaker will do his best to answer the question under discussion — up to the quality threshold. It determines a set of probability distributions — potential belief states for the speaker as far as the hearer knows based on what the speaker said and what the question under discussion is. These sets encode epistemic indeterminacy, and thus they tell us about the extent to which we draw clausal conversational implicatures.

To illustrate, I again turn to a variation on one of Grice's (1975) examples:

(21) A: Which city does Barbara live in?
B: Barbara lives in Russia.

As Grice observes, we typically draw the implicature from B's utterance that B does not know which city Barbara lives in. To show how the system predicts this, let's fix some details, as in (22) and (23).

(22) $W = \{w_1 \ldots w_4\}$

$$\begin{aligned} [\![\textit{Barbara lives in Russia}]\!] &= \{w_1, w_2\} \\ [\![\textit{Barbara lives in Moscow}]\!] &= \{w_1\} \\ [\![\textit{Barbara lives in Petersburg}]\!] &= \{w_2\} \end{aligned}$$

$$\begin{aligned} [\![\textit{Barbara lives in Germany}]\!] &= \{w_3, w_4\} \\ [\![\textit{Barbara lives in Berlin}]\!] &= \{w_3\} \\ [\![\textit{Barbara lives in Cologne}]\!] &= \{w_4\} \end{aligned}$$

(23) a. $C_\tau = .9$
b. $Q = \{\{w_1\}, \{w_2\}, \{w_3\}, \{w_4\}\}$
c. U (said by B) = "Barbara lives in Russia" relevance ranking = 2

Principle (20b) licenses A to infer from B's utterance that the quality threshold prevented B from saying any of the utterances with lower relevance rankings ("Barbara lives in Moscow", "Barbara lives in Petersburg", "Barbara lives in Berlin", and "Barbara lives in Cologne", which each have relevance rankings of 1). Thus, given this limited space, the only belief state B could be in is the following:

(24) a. $P_B([\![\textit{Barbara lives in Moscow}]\!]) = .5$
b. $P_B([\![\textit{Barbara lives in Petersburg}]\!]) = .5$

This is exactly the quantity implicature that Grice identified.

Minimally different contexts fail to support this inference, though. For instance, consider the following scenario, also based on the model in (22):

(25) a. $P_A(\{w_i\}) = .25$ for all $1 \leqslant i \leqslant 4$
b. A: "Does Barbara live in Russia?" $\{\{w_1, w_2\}, \{w_3, w_4\}\}$

In this scenario, according to the above system, B will answer the polarity question with "Yes, (she lives in Russia)". Here is a summary of the measurements:

(26)

utterance	quality	relevance	quantity
$[\![\textit{In Russia}]\!] = \{w_1, w_2\}$	1	1	1
$[\![\textit{On earth}]\!] = \{w_1 \ldots w_4\}$	1	2	0
$[\![\textit{In Moscow}]\!] = \{w_1\}$	1	1	2

Though "In Moscow" is more informative, it is not even relevance ranked, hence not even a contender: it shares a relevance ranking with the less-informative "In Russia", and thus it is eliminated at step ii of (12).

And, in turn, we do not make the standard quantity implicature that Grice articulated for (21). That is, we do not infer that B lacks more specific knowledge than he offers here with his "Yes" answer. We can use (20) to see that the context underdetermines B's belief state. As far as A is concerned, B might be in any of the following and still be playing by the rules of the pragmatic game:

(27) Contenders for B's belief state = $\{P_a, P_b, P_c\}$
a. $P_a(w_1) = P_a(w_2) = .5$
b. $P_b(w_1) = 1$
c. $P_c(w_2) = 1$

And this is just to say that we do not infer that B lacks more specific knowledge. We are open to the possibility that he doesn't have it (= P_a) and we are open the possibility that he does (= P_b or P_c).

5 Summary and Prospects

The foundational move of this paper is the shift to probabilities in the pragmatics (section 2.1). The resulting values form the basis for statements of the pressures that arrive via versions of quality, quantity, and relevance (section 2). With these values integrated (section 3), we are able to derive a range of conversational implicatures, as well as the lack of them in nonsupporting contexts (section 4).

In a sense, the results reduce to the following basic equations:

(28) a. In (12), we solve for U:

$$P_S, P_H, Q, C_\tau \Rightarrow \text{ a set of ranked utterances}$$

b. In (15), we learn something about Q:

$$P_H, U \Rightarrow \text{ a set of potential questions under discussion}$$

c. In (20), we solve for P_S:

$$Q, C_\tau, U \Rightarrow \text{ a set of potential belief states for the speaker}$$

It might be that I have not yet teased out the full set of implications of the theory for the size and composition of the output sets in (28). And it might be that a different set of pragmatic values, or a different method for combining them, will lead to even more nuanced inferences. It's my hope that this paper has at least provided a useful perspective on these matters.

I close by drawing out another consequence of this approach. A central question in pragmatic research these days is the nature of generalized conversational implicatures. Horn (2005) argues that they are implicatures like any other, and thus that they must be derived like any other — via appeal to semantic denotations, contextual information, and pragmatic pressures. They appear to have the status of default inferences only because they are supported by a wide range of contexts.

However, in the hands of Levinson (2000) and Chierchia (2004), generalized conversational implicatures are effectively lexicalized. They are, in Levinson's terms, presumptive meanings — default inferences, always present unless cancelled.

The present theory cannot decide the issue. It is consistent with both perspectives. But I hope it does show that Horn's very purely pragmatic approach is feasible. General pragmatic pressures can derive the relevant inferences. It is not necessary to stipulate them as part of the lexicon or semantic/pragmatic system. Moreover, I think results like those of section 4.3 point towards Horn's conception. There, we saw the conversational implicatures come and go as the question under discussion changed. There was no need to talk of cancellation or denial of implicatures. Rather, they arose where the context supported them, and they were simply absent where the context did not support them.

References

Benz, Anton, Gerhard Jäger, and Robert van Rooij. 2005. An introduction to game theory for linguists. In Anton Benz, Gerhard Jäger, and Robert van Rooij, eds., *Game Theory and Pragmatics*. Houndsmills, Basingstoke, Hampshire: Palgrave Macmillan.

Blutner, Reinhard. 1998. Lexical pragmatics. *Journal of Semantics* 15(2):115–162.

Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Adriana Belletti, ed., *Structures and Beyond: The Cartography of Syntactic Structures*, Volume 3, 39–103. New York: Oxford University Press. [The manuscript began circulating in 2001].

Frankfurt, Harry G. 1986. On bullshit. *Raritan* 6(2):81–100.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, eds., *Syntax and Semantics*, Volume 3: Speech Acts, 43–58. New York: Academic Press.

Groenendijk, Jeroen. 1999. The logic of interrogation. In Tanya Matthews and Devon Strolovitch, eds., *Proceedings of SALT IX*, 109–126. Ithaca, NY: Cornell University.

Groenendijk, Jeroen and Martin Stokhof. 1982. Semantic analysis of wh-complements. *Linguistics and Philosophy* 5(2):175–233.

Horn, Laurence R. 2005. The border wars. In Klaus von Heusinger and Ken P. Turner, eds., *Where Semantics Meets Pragmatics*. Oxford: Elsevier.

Kamp, Hans and Barbara H. Partee. 1995. Prototype theory and compositionality. *Cognition* 57(2):129–191.

Keenan, Edward L and Leonard M. Falz. 1985. *Boolean Semantics for Natural Language*. Dordrecht: D. Reidel.

Kratzer, Angelika. 2002. Facts: Particulars or information units? *Linguistics and Philosophy* 25(5–6):655–670.

Levinson, Stephen C. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.

Merin, Arthur. 1997. If all our arguments had to be conclusive, there would be few of them. Arbeitspapiere SFB 340 101, University of Stuttgart, Stuttgart. URL `http://semanticsarchive.net/Archive/jVkZDI3M/`.

Potts, Christopher. 2006. How far can pragmatic mechanisms take us? Commentary on truckenbrodt. *Theoretical Linguistics* URL `http://semanticsarchive.net/Archive/jViODAyZ/`.

Roberts, Craige. 1996. Information structure: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon and Andreas Kathol, eds., *OSU Working Papers in Linguistics*, Volume Volume 49: Papers in Semantics, 91–136. Columbus, OH: The Ohio State University Department of Linguistics. Revised 1998.

Roberts, Craige. 2004. Context in dynamic interpretation. In Laurence R. Horn and Gregory Ward, eds., *Handbook of Contemporary Pragmatic Theory*, 197–220. Oxford: Blackwell.

van Rooy, Robert. 2003a. Questioning to resolve decision problems. *Linguistics and Philosophy* 26(6):727–763.

van Rooy, Robert. 2003b. Relevance and bidirectional OT. In Reinhard Blutner and Henk Zeevat, eds., *Pragmatics in Optimality Theory*, 173–210. Oxford: Palgrave Macmillan.

Sperber, Dan and Deirdre Wilson. 1995. *Relevance: Communication and Cognition*. Oxford: Blackwell, 2 ed.