Manuscript, UMass Amherst

Integrated pragmatic values*

Christopher Potts
UMass Amherst

September 30

Abstract

I adapt recent results by Merin, Blutner, Jäger, Krifka, van Rooy, and others to obtain integrated pragmatic values for utterances, thereby moving towards a precise definition of pragmatic felicity. I model speakers' perspectives with probability distributions over the set of possible worlds. The *quality-rating* of an utterance is an exponent of the speaker's probability value for its propositional content. An utterance's *quantity-rating* is the informativity of its content relative to the addressee's probabilities. I ensure that quality-ratings act as a check on quantity-ratings by taking the product of the two. I employ the notion of *relevance to a question* to further articulate these pragmatic values and arrive at a notion of *maximally felicitous utterance* (in context). I apply the ideas to a range of simple question–answer scenarios, and I outline extensions of them into the more challenging and novel areas of litotes (negative understatement), pragmatic halos, discourse particles, and bullshit.

1 The probability of formalizing the maxims

The formally-minded pragmaticalist could be forgiven for feeling discouraged by some of the literature, which offers more than a few pronouncements that the second of the juicy questions in (1) is outside the bounds of formalization.

*For advice and helpful insights, I'm indebted to Ash Asudeh, Rajesh Bhatt, Shai Cohen, Kathryn Flack, Lyn Frazier, Hans-Martin Gärtner, Yurie Hara, Gerhard Jäger, Elke Kasimir, Manfred Krifka, Eric McCready, Barbara Partee, Tim Roeper, Tom Roeper, Uli Sauerland, Florian Schwarz, Anna Verbuk, Malte Zimmermann, and the audiences at the UMass Amherst Semantics Group, September 6, at Göttingen University, September 21, and ZAS, September 23.

(1) a. What range of meanings is accessible to linguistic pragmatics?

b. How do the pragmatic maxims work?

To be sure, we can be cheered by recent advances regarding (1a). There is now a wealth of evidence that only a proper subset of meanings can be generated by linguistic and contextual interactions (Fauconnier 1975; Horn 1989; Chierchia 2004; Sauerland 2004). That is, the range of meanings relevant to linguistic pragmatics is partly conventionalized, suggesting that the tools of linguistic semantics are useful here.

But on the topic of (1b), one often encounters pessimism. In answer to the related question, "Will we ever have a formal theory of the maxims?", some reply with skepticism, others with a firm negative, and still others assert that the question itself is misguided. For instance, Beaver (2001:29) calls formalization in this area "notoriously problematic". Bach (1999) is more decisive, offering various reasons why "it seems futile for linguists to seek a formal pragmatics". Devitt and Sterelny (1987:§7.4) strike a similar chord.

It's a harsh verdict. Pragmatic maxims are the heart of pragmatic theory. It doesn't matter if there are five or more (Grice 1989) or just one (Sperber and Wilson 1995). It doesn't matter if they operate externally to the semantics (Gazdar 1979b,a), as part of the recursive semantics (Geurts and Maier 2003; Chierchia 2004), or somewhere in between. (Levinson 2000 offers extensive discussion.) Invariably, if one wants to provide a complete description of a context-dependent meaning, one needs to lean on the rationality of discourse participants (a cooperative principle) and some additional contractual obligations of the members of the speech community (maxims).

It is true, though, that the maxims — and in turn, pragmatic meanings — do not yield easily to a treatment in the usual terms of semantic theory. One can usually be precise up to a point, and then one must tell a story that depends more on one's knowledge of pragmatics than it does on the principles of one's theory. One thinks of the famous Sidney Harris cartoon in which a scientist's mad chalkboard scribbles are bridged by the phrase "And then a miracle occurs".

Things are looking up, though. Recent work by Reinhard Blutner, Gerhard Jäger, Arthur Merin, Robert van Rooij, and others has shed new light on the situation. The chief innovation: a shift in emphasis from truth-conditions to probabilities. The aim of this paper is to synthesize some of these developments, to provide a clear picture of what theories in this vein look like and what they are capable of. Grice's maxims are my touchstone; as an organizing principle, I attempt to formalize Grice's maxims of quality, quantity, and relevance. But the formal tools employed seem general enough to work, with modifications, under various formulations of the maxims. For now, I set aside the maxim of manner, which is though discussed in roughly the present terms by Blutner (1998) and van Rooy (2003c).

The literature on probabilistic and game-theoretic approaches to pragmatics is new but growing fast. Many of the papers cover a lot of ground, technically and empirically. Here is a brief summary of the insightful works I mined for the present paper:

- The informativity measures relevant for quantity (section 6) are explicated and explored by Blutner (1998, 2000) and van Rooy (2003c) (and, indirectly, by van Rooy (2003a) and Krifka (2003)).
- Merin (1997, 2005) explains how to use probability distributions in linguistic semantics and pragmatics.
- van Rooy (2004a) explores relevance in general, and van Rooy (2003b) applies some of those notions and more to interrogative contexts.
- Benz et al.'s (2005) overview article on probabilistic and game-theoretic methods is essential. The present paper does not make an explicit connection with game theory, but Benz et al. provide the requisite details. The related conceptual link with Bidirectional OT is explored by Jäger (2002) and van Rooy (2004b).

2 A context and two scenarios

Our two discourse participants, call them player *I* and player *A*, want to visit Barbara. Player *I* has no beliefs about where Barbara lives, so he asks player *A*,

(2) "Where does Barbara live?"

The players' interpretation of this question — in particular, the level of detail requested by *where* — will be heavily conditioned by a range of contextual factors, as will player *A*'s reply. Let's assume that the players' shared goal is to buy plane tickets, and thus that (2) is interpreted by them as equivalent to the question of which *city* Barbara lives in. This is the question under discussion. We'll focus on two different utterances with this backdrop.

(3) **Scenario 1: Underinformative relative to the question**

1. **Player *A*'s knowledge**
   Player *A* knows which country Barbara lives in: Russia. But he is not sure which city in Russia she lives in.
2. **Player *A*'s reply**
   "Well, she lives in Russia".
3. **Player *I*'s calculation**
   I will conclude that *A* has limited knowledge, that *A* was not positioned to name a city.

Player *A* does not have sufficient evidence to name a city. So he offers an underinformative answer. Our pragmatic theory should identify Player *A*'s response as the preferred one in this situation, despite possible semantic deviance in the eyes of many theories of questions (Groenendijk and Stokhof 1982, 1988; Groenendijk 1999; cf. Beck and Rullmann 1999; van Rooy 2003b).

Scenario 2 involves a different kind of underinformativity, and, in turn, a different sort of pragmatic meaning.

(4) **Scenario 2: Underinformative relative to *A*'s knowledge**

1. **Player *A*'s knowledge**
   Player *A* knows Barbara's exact street address: Tallinskaja 2, Moscow.
2. **Player *A*'s reply**
   "She lives in Moscow."
3. **Player *I*'s calculation**
   I will not conclude that player *A* lacks more specific knowledge than this, because Player *A* offered the right amount of information given the shared goal of the discourse.

Here, Player *A* is positioned to say something more informative, but doing so would be, we sense, infelicitous in light of the question under discussion.

3 The model

Like many semantic theories, this one is grounded in a set W of possible worlds. The power-set of W, $\wp(W)$, is the set of all propositions. Here are the facts:

(5) $W = \{w_1 \ldots w_8\}$

$$\begin{aligned}
[\![\textit{Barbara lives in Moscow}]\!] &= \{w_1, w_2\} \\
[\![\textit{Barbara lives on Tallinskaja St, Moscow}]\!] &= \{w_1\} \\
[\![\textit{Barbara lives on Pushkinskaja St, Moscow}]\!] &= \{w_2\} \\
[\![\textit{Barbara lives in Petersburg}]\!] &= \{w_3, w_4\} \\
[\![\textit{Barbara lives on Kolomenskaja St, Petersburg}]\!] &= \{w_3\} \\
[\![\textit{Barbara lives on Nevsky Prospekt, Petersburg}]\!] &= \{w_4\} \\
[\![\textit{Barbara lives in New York}]\!] &= \{w_5, w_6\} \\
[\![\textit{Barbara lives on 2nd Avenue, New York}]\!] &= \{w_5\} \\
[\![\textit{Barbara lives on Union Square, New York}]\!] &= \{w_6\} \\
[\![\textit{Barbara lives in Northampton}]\!] &= \{w_7, w_8\} \\
[\![\textit{Barbara lives on Main St, Northampton}]\!] &= \{w_7\} \\
[\![\textit{Barbara lives on Pleasant St, Northampton}]\!] &= \{w_8\}
\end{aligned}$$

We can use this model to articulate certain background assumptions of scenarios 1 and 2 above. First, I specified that the players ask the question *Where does Barbara live?* at the city-level (i.e., as semantically equivalent to *Which city does Barbara live in?*). The partition semantics for this question is thus as follows:

(6) $[\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{cc} \{w_1, w_2\} & \{w_3, w_4\} \\ \{w_5, w_6\} & \{w_7, w_8\} \end{array} \right\}$

In addition, we can make precise the assumptions made about the players' beliefs:

1. In both scenario 1 and scenario 2, Player I is in the state of complete ignorance: he regards any of the members of w as live possibilities for the actual world; his belief state is representable with W.

2. In scenario 1, player A knows that Barbara lives in Russia, though she is unsure which city (Moscow or Petersburg). We can model player A's belief worlds with $\{w_1 \ldots w_4\}$.

3. In scenario 2, player A knows that Barbara lives in Moscow, and, moreover, that she lives on Tallinskaja Street. So, in this scenario, player A's belief state corresponds to $\{w_1\}$.

In the next section, I employ probability distributions to model these facts about the players' beliefs.

4 Knowledge of the model

4.1 Probability distributions

I adopt an additional perspective on the set of all propositions, by employing *probability distributions* over W:

(7) A function $P : \wp(W) \mapsto [0,1]$ is a probability distribution iff:

a. $P(W) = 1$

b. $P(\{w\}) \geqslant 0$, for all $w \in W$

c. Probabilities are additive: if p and q are disjoint propositions, then $P(p \cup q) = P(p) + P(q)$

In virtue of clause (iii), we can derive the probabilities of nonsingleton propositions from the probabilities of singleton propositions. If $p = \{w_1 \ldots w_n\}$, then we calculate $P(\{w_1\} \cup \ldots \cup \{w_n\})$ to find $P(p)$. A bit more formally:

$$P(p) = \sum_{w \in p} P(\{w\})$$

4.1.1 Probability distributions for propositions

Probability distributions can mimic the usual view of propositions (Merin 1997, 2005). The definition in (8) connects them.

(8) 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 $w, w' \in q$

Clause (8a) associates 0 probabilities with non-membership in the propositional counterpart. Clause (8b) ensures that all the worlds in the proposition q have identical probabilities according to P. Together with the axioms for probability distributions, these clauses ensure that if we add up all the probabilities for the worlds in q, we get a value of 1.

4.1.2 Probability distributions for conversational implicatures

With (8), we can set up a homomorphic mapping from propositions into probability distributions, thereby allowing us to translate possible-worlds talk into probabilistic talk. But the smoothing achieved by clause (8b) is not a feature of all probability distributions. For instance, the probability distribution depicted in (9) does not mimic any proposition.

(9)
$$\begin{bmatrix} P(\{w_1\}) & \mapsto & .1 \\ P(\{w_2\}) & \mapsto & .4 \\ P(\{w_3\}) & \mapsto & .4 \\ P(\{w_4\}) & \mapsto & .1 \end{bmatrix}$$

These uneven probability distributions might ultimately turn out to be the most important for pragmatic theory. Let's briefly consider the question of what a conversational implicature update should look like in a dynamic setting. The central characteristics distinguishing such updates from the updates of regular semantic content are *cancellablity* and *reinforceablity* (Levinson 2000:15). Here are two simple illustrations:

(10) **Cancellablity**

a. Ed has six fingers — in fact, he has ten.

b. #Ed has exactly six fingers — in fact he has ten.

(11) **Reinforceablity**

a. Ed has six fingers — in fact, he has exactly six.

b. #Ed has exactly six fingers — in fact he has exactly.

The corresponding semantic versions — the (b) examples — are contradictory (in (10)) or redundant (in (11)).

In dynamic semantics, in interpreting (10a), I either updated with the content that Ed has exactly six fingers, or I did not. If I made the update, then the example is identical to (10b). If I did not make the update, then the conversational implicature is not registered in the context at all. We are left with the question of how conversational implicature updates are different.

To answer this, we need a grip on the notion of a *tentative update*. Unven probability distributions like (9) might be able to accomplish this. If I update p as an implicature, the not-p worlds might end up with *reduced, but non*-0 probabilities, as in worlds w_1 and w_4 in (9). If I reinforce p, as in (11), then the probabilities for not-p worlds drop to 0, as per definition (8). If I cancel p, as in (10), then the probabilities for not-p worlds rise, again as per definition (8).

4.2 Probability distributions for our players

Probability distributions provide a means for representing our players' beliefs, or, alternatively, their degrees of commitment to certain pieces of information.

As noted above, player I is in the state of ignorance in both our scenarios. That is, $P_I(\{w\}) = \frac{1}{8}$, for all $w \in W$. Player A has sharper beliefs. In scenario 1, she believes that Barbara lives in Russia, but nothing more specific than that. That is, in scenario 1, $P_A(\{w_i\}) = .25$ for all $1 \leqslant i \leqslant 4$, and $P_A(\{w_j\}) = 0$ for all $5 \leqslant j \leqslant 8$. Here are graphic depictions of these probability distributions:

(12) **Scenario 1**

$$P_I = \begin{bmatrix} P(\{w_1\}) & \mapsto & .125 \\ P(\{w_2\}) & \mapsto & .125 \\ P(\{w_3\}) & \mapsto & .125 \\ P(\{w_4\}) & \mapsto & .125 \\ P(\{w_5\}) & \mapsto & .125 \\ P(\{w_6\}) & \mapsto & .125 \\ P(\{w_7\}) & \mapsto & .125 \\ P(\{w_8\}) & \mapsto & .125 \end{bmatrix} \qquad P_A = \begin{bmatrix} P(\{w_1\}) & \mapsto & .25 \\ P(\{w_2\}) & \mapsto & .25 \\ P(\{w_3\}) & \mapsto & .25 \\ P(\{w_4\}) & \mapsto & .25 \\ P(\{w_5\}) & \mapsto & 0 \\ P(\{w_6\}) & \mapsto & 0 \\ P(\{w_7\}) & \mapsto & 0 \\ P(\{w_8\}) & \mapsto & 0 \end{bmatrix}$$

(13) **Scenario 2**

$$P_I = \begin{bmatrix} P(\{w_1\}) & \mapsto & .125 \\ P(\{w_2\}) & \mapsto & .125 \\ P(\{w_3\}) & \mapsto & .125 \\ P(\{w_4\}) & \mapsto & .125 \\ P(\{w_5\}) & \mapsto & .125 \\ P(\{w_6\}) & \mapsto & .125 \\ P(\{w_7\}) & \mapsto & .125 \\ P(\{w_8\}) & \mapsto & .125 \end{bmatrix} \qquad P_A = \begin{bmatrix} P(\{w_1\}) & \mapsto & 1 \\ P(\{w_2\}) & \mapsto & 0 \\ P(\{w_3\}) & \mapsto & 0 \\ P(\{w_4\}) & \mapsto & 0 \\ P(\{w_5\}) & \mapsto & 0 \\ P(\{w_6\}) & \mapsto & 0 \\ P(\{w_7\}) & \mapsto & 0 \\ P(\{w_8\}) & \mapsto & 0 \end{bmatrix}$$

4.3 Utterance contexts

Our pragmatic values will of course be set relative to specific utterance contexts. In practical terms, this means that our definitions are encumbered by parameters for the speaker, the addressee, belief states for each, etc. To keep things relatively straightforward, I define here a notion of context that is just rich enough to make the distinctions we need to make:

(14) An *utterance context* is a tuple $C = \langle S, P, Q, U^{a_i \mapsto a_j} \rangle$, where

1. $S = \{a_1, a_2, \ldots\}$ is the set of players.
2. $P = \{P_1, P_2, \ldots\}$ is the set of probability distributions, where P_i represents the beliefs of player a_i.
3. $Q \subseteq \wp(W)$ is the question under discussion.
4. $U^{a_i \mapsto a_j}$ is an utterance by player $a_i \in A$ to addressee $a_j \in S$.

I often use s to pick out the speaker and a to pick out the addressee. In such cases, I am assuming that $S = \{s, a\}$.

Each of the scenarios in section 2 defines an utterance context in the above sense. Player A is the speaker, player I the addressee. The elements P_A and P_I represent their beliefs, and $U^{A \mapsto I}$ is always an utterance by A to I (in response to the immediately preceding utterance situation in which I was the speaker, A the addressee, and $U^{I \mapsto A}$ was "Where does Barbara live").

For easy reference, here are the scenarios in the mold of utterance contexts:

Scenario 1

$$S = \{A, I\}$$

$$P_I = \begin{bmatrix} P(\{w_1\}) & \mapsto & .125 \\ P(\{w_2\}) & \mapsto & .125 \\ P(\{w_3\}) & \mapsto & .125 \\ P(\{w_4\}) & \mapsto & .125 \\ P(\{w_5\}) & \mapsto & .125 \\ P(\{w_6\}) & \mapsto & .125 \\ P(\{w_7\}) & \mapsto & .125 \\ P(\{w_8\}) & \mapsto & .125 \end{bmatrix} \quad P_A = \begin{bmatrix} P(\{w_1\}) & \mapsto & .25 \\ P(\{w_2\}) & \mapsto & .25 \\ P(\{w_3\}) & \mapsto & .25 \\ P(\{w_4\}) & \mapsto & .25 \\ P(\{w_5\}) & \mapsto & 0 \\ P(\{w_6\}) & \mapsto & 0 \\ P(\{w_7\}) & \mapsto & 0 \\ P(\{w_8\}) & \mapsto & 0 \end{bmatrix}$$

$$Q = [\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{cc} \{w_1, w_2\} & \{w_3, w_4\} \\ \{w_5, w_6\} & \{w_7, w_8\} \end{array} \right\}$$

$$U^{A \mapsto I} = \text{“In Russia”}$$

Scenario 2

$$S = \{A, I\}$$

$$P_I = \begin{bmatrix} P(\{w_1\}) \mapsto .125 \\ P(\{w_2\}) \mapsto .125 \\ P(\{w_3\}) \mapsto .125 \\ P(\{w_4\}) \mapsto .125 \\ P(\{w_5\}) \mapsto .125 \\ P(\{w_6\}) \mapsto .125 \\ P(\{w_7\}) \mapsto .125 \\ P(\{w_8\}) \mapsto .125 \end{bmatrix} \quad P_A = \begin{bmatrix} P(\{w_1\}) \mapsto 1 \\ P(\{w_2\}) \mapsto 0 \\ P(\{w_3\}) \mapsto 0 \\ P(\{w_4\}) \mapsto 0 \\ P(\{w_5\}) \mapsto 0 \\ P(\{w_6\}) \mapsto 0 \\ P(\{w_7\}) \mapsto 0 \\ P(\{w_8\}) \mapsto 0 \end{bmatrix}$$

$$Q = [\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{matrix} \{w_1, w_2\} & \{w_3, w_4\} \\ \{w_5, w_6\} & \{w_7, w_8\} \end{matrix} \right\}$$

$$U^{A \mapsto I} = \text{“In Moscow”}$$

5 Quality

5.1 Quality-ratings

The above view of probability distributions is all we need for a fresh statement of Grice's quality maxim, which is given in its original form in (15).

(15) **Quality (Grice 1975)**

Say only what you know to be true. Do not say that which is false. Do not say that which you lack evidence for.

We don't do too much damage to this maxim if we reduce it to the injunction that speakers should assert only those propositions that are entailments of their beliefs. We could formalize this in nonprobabilistic terms: an utterance U by speaker s (in w) respects quality iff the semantic value of U is a superset of the set of belief worlds for s (in w).

But, in the spirit of the semantics described above, I follow a different route, relying instead on a gradient view of quality. The general strategy is to assign *quality-ratings* to utterances in context. There are numerous ways that one can do this. To start, we can imagine that we simply identify quality-ratings with the values for our probability distributions. For instance, in scenario 1, player A assigns the proposition that Barbara lives on Tallinskaja Street the probability .25. This would seem to represent only a modicum of confidence in this proposition. Player I assigns this same proposition the probability .125, an even lower value, representing even less confidence and, one might imagine, even less felicity if asserted.

However, these values do not drop off rapidly enough. If a speaker assigns *p* a probability of .5, then we want that speaker to be extremely reluctant to assert *p*. The quality-rating for *p* should be much lower than .5, a number that seems to suggest a fair degree of assertability.

To achieve a steeper drop-off in values, I employ exponents of the values we obtain from our players' probability distributions. The value of the exponent can be any *n*. The higher *n* is, the greater a speaker's confidence in a proposition will have to be before he can felicitously assert it. In (16), I provide a graph depicting five values of *n*, with exponents of probabilities along the *y*-axis and probabilities along the *x*-axis). The leftmost line is with *n* set to 1 and thus duplicates the values for the probability distribution. The next three values are for exponents 2, 3, and 4. The rightmost curve has the exponent at 50, and thus nearly mimics a very rigid 1-or-nothing view of quality.

(16) quality-rating

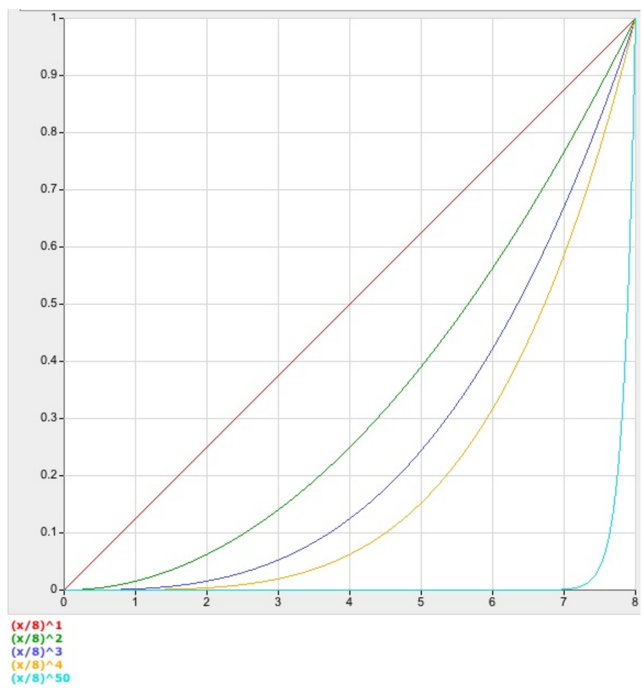

P

Emphasizing that 4 is somewhat arbitrarily chosen, I offer (17), which assigns *quality-ratings* to speakers' utterances.

(17) **Quality-rating**

The *quality-rating* of an utterance $U^{s \mapsto a}$ with content p in context C is

$$\text{Quality}_C(p) = (P_s(p))^4$$

We can reproduce the traditional, binary formulation of quality by saying that $U^{s \mapsto a}$ respects quality iff its semantic content of has a probability of 1 according to P_s. Or, almost equivalently, we could make the exponent very high, as in the rightmost curve in (16). But the truly gradient view defined in (17) is useful in studying interactions among the maxims and understanding which interpretations count as preferred.

It is often helpful to imagine what would happen if a given maxim constituted the only pressure speakers felt when choosing their utterances. If quality were the only pressure, then speakers would maximize the quality-ratings of their utterances by saying only those things that have probability of 1 for them. In reality, though, speakers are rarely this confident about the things they say. The things about which they have the highest degree of confidence are generally dull (tautological) things. Speakers take risks with contingent truths, things that they might have some doubts about. Why? Because they want their utterances to be relevant and contentful. The primary function of quality-ratings is to keep the forces of relevance and informativity from growing so powerful that they overwhelm belief.

5.2 Where there is doubt

One might be initially suspicious of the gradient view of quality that is suggested by the above account. The original Gricean account (slightly simplified) might amount to simply: Say only those things that are entailments of your beliefs. That condition is very easy to formalize, and it delivers the right results for a broad range of cases.

This classical view can of course be duplicated in the above system. If the exponent involved in deriving quality-ratings from probabilities is set high enough, then anything less than a 1 probability will deliver extremely low quality-ratings. But, again, the reply might come back: why work hard to reproduce what we can get straightaway from entailments?

The next few subsections are devoted to answering this objection from quite narrow linguistic evidence. I try to identify contexts and constructions that depend on a gradient view of quality and even suggest that the value of the exponent in (17) is context-dependent.

Before moving to those topics, it's worth pointing out that there are plenty of people who never have one-hundred percent confidence in what they say. If one has seriously considered Hilary Putnam's 'brain in a vat' thought experiment and its many variants, one will probably admit that one is never *completely* confident that what one says is true. I feel a high degree of confidence that my name is Christopher Potts, but, then again, I've never had my birth certificate validated, nor am I sure that such validation would remove all lingering doubts that my name is in fact something else. And, more generally (once one has started down this road), I might be a brain in a vat, in which case everything I say is false.

5.2.1 On pragmatic halos

Lasersohn's (1999) central observation is that we nearly always approximate the truth with our utterances. If I tell you that Kyle arrived at the office at 8:00, what I say is very likely false; the chances that he arrived right on the hour are slim. He might have arrived at 7:59:04, or 8:01:27, or 8:05:30… But what I said is, in Lasersohn's words, "close enough to the truth" to count as an accurate report on the state of things. If Kyle arrived at 8:01:27, then, in normal circumstances, the statement that he arrived at 8:00 will be in the *pragmatic halo* of the truth, and thus it will count as close enough in the technical sense. Pragmatic halos are highly context-dependent — *He arrived at 8:00* might be infelicitous if we are running a road race and he arrived (at the finish) at 8:01:00. *He arrived at 8:00* might be close enough to true if he arrived at 8:27:00 after a journey across four continents.

For our purposes, what's important is that we all recognize that we only approximate the truth with our claims. I say that Kyle arrived at 8:00, but I pick the round number only out of a general pragmatic preference for round numbers (Krifka 2005), not (necessarily) because I believe in it more strongly than the other elements in its halo. My probability distribution actually looks something like this:

(18)

$$\begin{bmatrix} [\![\textit{Kyle arrived at 7:57}]\!] & \mapsto & 0 \\ [\![\textit{Kyle arrived at 7:58}]\!] & \mapsto & .2 \\ [\![\textit{Kyle arrived at 7:59}]\!] & \mapsto & .2 \\ [\![\textit{Kyle arrived at 8:00}]\!] & \mapsto & .2 \\ [\![\textit{Kyle arrived at 8:01}]\!] & \mapsto & .2 \\ [\![\textit{Kyle arrived at 8:02}]\!] & \mapsto & .2 \\ [\![\textit{Kyle arrived at 8:03}]\!] & \mapsto & 0 \end{bmatrix}$$

The halo is reflected by, or can be reconstructed by, the fact that these probabilities sum to 1. We can increase the size of the halo by lowering the value of the exponent in (17), and we can decrease the size of the halo by raising that value.

Crucially, we see here that I have some doubt about the proposition that Kyle arrived at 8:00. This in turn leads to a degraded quality-rating for that proposition. The rigid view of quality — one-hundred percent belief or nothing — would leave me speechless or force me to extremes of approximation.

5.2.2 On bullshit

Frankfurt (1986) draws a subtle, important distinction between the liar and the bullshitter. The liar has a vested interest in the truth — if he is pathological, he aims to tell you its opposite. Thus, if you discover a liar, you have a path to the truth: just reverse the truth values of all his asserted content. The bullshitter, in contrast, is by definition indifferent to the truth. He tells you p, and p might be true, but, then again, it might be false. He might simply want to try out what it feels like to assert p. Or he might see some value in getting you to accept p, true or not. The unpredictable relationship to the truth is what makes bullshit so invidious.

Possible-worlds models provide an easy way to identify bullshit: a proposition p is bullshit for a speaker S just in case the belief state of S, construed as a set of worlds, contains both p worlds and not-p worlds. If you are a skeptical philosopher, everything is bullshit to some non-null degree.

With probability distributions, we can measure the degree to which something is bullshit for a speaker. This will be directly reflected in the quality-ratings. What is of interest here is that our *tolerance* for bullshit is highly context-dependent. If we are brainstorming new approaches to presupposition projection, our standards for belief and accountability might be low. If we are trying to land a rocket ship on the moon, our standards should be very high. The present system allows us to make such fine-grained adjustments — we simply allow the value of the exponent in (17) to value with the context. Or, to put it another way, the exponent reflects the strictness of the context.

Frankfurt (1986) observes that some conventionalized expressions are almost inherently bullshit:

(19) a. I feel like I am floating on air.

b. I am sick as a dog.

I do not know what it feels like to be a dog, sick or otherwise. And I am unsure of what the sensation of floating on air would be like. So, if I utter either of the sentences in (19), I lack the sort of evidence that we normally demand for assertions. But my audience is likely to be charitable. By convention, our quality-ratings are lenient for utterances of these sentences.

5.2.3 On discourse particles

There is great potential in quality-ratings for enforcing pragmatic conditions on modals, evidentials, and the like. For instance, it is rarely if ever made part of the semantics of *might p* that the speaker is not positioned to assert *p*. But we can nonetheless see that felicitous utterances of *might p* almost always happen in situations in which, in the present terminology, the speaker assigns a non-0, non-1 probability to *p*.

I think it would be a mistake to push the probabilities much deeper into the realm of modals. We have a semantics for modals. The probabilities are for pragmatics.

But what about modal-like discourse particles? Hara (2005) studies the Japanese discourse particle *darou*. Its meaning resembles that of the regular modal adverb *tabun* ('probably'), but the two have different distributions, and speakers intuit that their meanings are slightly different. We might capture the differences by giving a standard, possible-worlds semantics for *tabun*, while treating *darou* entirely in terms of preconditions on, or alterations to, the probability distribution of the speaker. A similar move might be appropriate for the German pair *wahrscheinlich* ('probably') and *wohl* (also somewhat like *probably*, but in the form of a discourse particle; Zimmermann 2004). *Wahrscheinlich* might be a modal operator. *Wohl* might be keyed to specific quality-ratings.

6 Quantity

6.1 Quantity-ratings

The quantity maxim calls on speakers to tailor the information content of their utterances to the present discourse situation:

(20) **Quantity (Grice 1975)**

Make your contribution as informative as is required. Do not say more than is required.

The "required" portions of this maxim are, I believe, fully duplicated by the relevance maxim, discussed in detail below. So I regard quantity as a call for speakers to maximize information content, expecting this injunction to be mitigated by quality as well as relevance.

So we need a measure of information content. There are many such measures on the market (van Rooy 2004a; Benz et al. 2005). To keep things simple, I adopt a version of Blutner's (1998) proposal to derive information content from probabilities using a logarithmic function, in the mode of Carnap (1950). Here is the basic measure:

(21) **Information value of p for player a**

$$\mathrm{inf}_a(p) = -\log_2 P_a(p)$$

By this measure, informativity values rise as probabilities fall (with the probability of 0 assigned the pathological value ∞):

(22) $-\log_2$

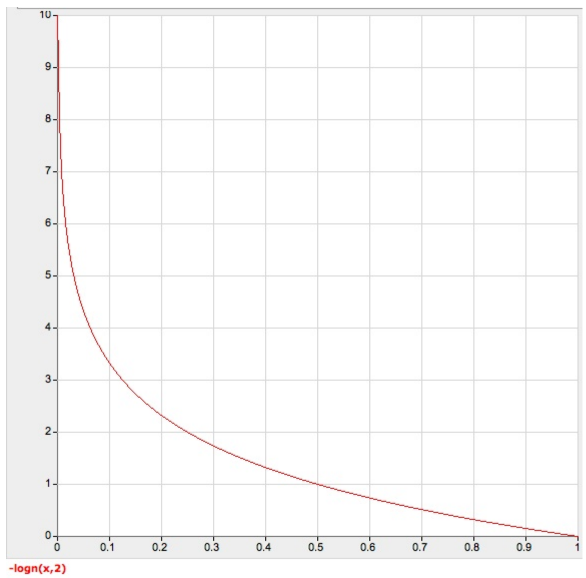

probability

As with quality, we want to assign utterances *quantity-ratings*. However, it won't do to identify these with the inf values for propositions relative to the speaker's probability distribution. On that approach, the more strongly a speaker believed a proposition p, the lower p's information content would be. This indicates that the speaker's belief state is not the one we want to use to model quantity.

I argue instead that we should use the *addressee's* probability function in calculating the inf values relevant for pragmatics, as in (23).

(23) **Quantity-rating**

The *quantity-rating* of the proposition p in context C with addressee a is

$$\mathrm{Quantity}_C(p) = \mathrm{inf}_a(p)$$

It might seem strange at first that the quantity-rating of the speaker's utterances are determined by the probability distribution of her addressee. But this seems to match well our intuitions about pragmatic values. I might accidentally tell you something you already know, on the mistaken assumption that it is new to you. In such cases, the information value of what I said is indeed very low. More generally, as a speaker, I must guess about what your probability distribution is like. If I guess wrong, my utterance is infelicitous. (You might also be insulted by my supposition about your belief state.)

6.2 Litotes

In section 5.2, I identified some grammatical phenomena that seem especially relevant to quality and quality-ratings. For quantity-ratings, I think comparable evidence is available in the form of *litotes* (Horn 1989; Majewski 2005). Some typical examples are given in (24).

(24) a. She's not exactly stupid. [$\approx$ she's a genius]

b. The picture isn't exactly even with the top of the door. [$\approx$ it's way off]

We leap to the opposite end of the scale. In (24a), we go from *stupid* to genius. And (24b) is naturally read as an amazed observation at how far out of alignment the door and picture are. What motivates this sharp change in direction?

Example (24b) is useful in answering this question. It differs from (24a) in having a non-litotic reading. This is presumably connected with the fact that *The picture is exactly even with the top of door* is well formed, whereas **Sue is exactly stupid* is ungrammatical. I refer to Majewski 2005 for discussion of this point. For now, I suggest just looking at *not exactly even*. The important thing about this predicate is that it is extremely uninformative. Of all the possible positions that the picture can be in, we learn from this predication only that it is not at the position even with the door. This leaves a huge range of possibilities open.

This near lack of information content is encoded in quantity-ratings. For any context, utterances of this phrase will have very low quantity-ratings relative to other predications of a similar form (*exactly even*, *not even*, etc.). I claim that these low ratings make for unstable predicates. They are pragmatically marked, and this makes them prime for pragmatic reanalysis. The conventional aspect of this — the aspect of special linguistic interest — is that the instability engendered by semantic uninformativeness leads to a maximization of informativeness — we leap all the way to the opposite end of the scale, thereby maximizing quantity-ratings.

We might even go so far as to say that **exactly happy* is marked because it is maximally uninformative, in virtue of the fact that we lack a standard of exactness for happiness. This paves the way for an obligatory litotic interpretation.

7 Quality–quantity interactions

Imagine what would happen if the quantity maxim were left unchecked. Let the background be some set p. Speakers in a quantity-only community would always choose utterances that expressed some singleton subset of p. These propositions have the lowest non-0 probabilities and hence the highest inf values. Speakers would offer this content regardless of whether or not they believed these propositions were true.

The quality maxim provides an important check on the content that speakers can felicitously offer. Why don't speakers actually restrict themselves to maximally specific content? Part of the answer is that they rarely if ever have the capacity (or poor judgment) to believe such things. As a result, such things receive poor quality-ratings.

So quality is a check on quantity. I propose to formalize this with the product of the relevant quality and quantity (inf) ratings:

(25) **Quality–quantity (QQ-) ratings**

Let $U^{s \mapsto a}$ in context C have propositional content p The quality–quantity (QQ) rating of p, written $\mathrm{QQ}_C(p)$, is the product of its quality and quantity ratings. That is:

$$\mathrm{QQ}_C(p) = \mathrm{Quality}_C(p) \times \mathrm{Quantity}_C(p)$$

Scenario 1, described in (3), is excellent for observing how QQ-ratings work. At the heart of that scenario is the observation that speaker A does not provide a semantically complete answer to the question "Where does Barbara live?" when she replies "In Russia". The stated goal of the discourse participants is to determine which city Barbara lives in (so that they can purchase plane tickets). The utterance "In Russia" does not narrow down the options to a single city even in our impoverished universe (5).

However, we intuit also that player A's response is the most felicitous available to her, given the goals and limitations of the context. In the technical terms laid out above, this is because all utterances of the form "Barbara lives in city C" have poor quality-ratings for A. In (26), I provide a more detailed look at the state of things. (These values are the same for "Barbara lives in Petersburg".)

(26) a. **The scenario 1 quality-rating for "Barbara lives in Moscow"**

$$(P_A([\![\textit{Barbara lives in Moscow}]\!]))^4 = .0625$$

b. **The scenario 1 quantity-rating for "Barbara lives in Moscow"**

$$\mathrm{inf}_I([\![\textit{Barbara lives in Moscow}]\!]) = -\log_2 .25 = 2$$

c. **The scenario 1 QQ-rating for "Barbara lives in Moscow"**

$$\mathrm{QQ}_{\text{scenario 1}}([\![\textit{Barbara lives in Moscow}]\!]) = .0625 \times 2 = .125$$

Let's compare this to the rating of A's actual utterance in scenario 1, which is "In Russia":

(27) a. **The scenario 1 quality-rating for "Barbara lives in Russia"**

$$(P_A([\![\textit{Barbara lives in Russia}]\!]))^4 = 1$$

b. **The scenario 1 quantity-rating for "Barbara lives in Russia"**

$$\mathrm{inf}_I([\![\textit{Barbara lives in Russia}]\!]) = -\log_2 .5 = 1$$

c. **The scenario 1 QQ-rating for "Barbara lives in Russia"**

$$\mathrm{QQ}_{\text{scenario 1}}([\![\textit{Barbara lives in Russia}]\!]) = 1 \times 1 = 1$$

As desired, A's answer "In Russia" receives a higher QQ-rating than "In Moscow" (1 vs. .125), and we are thus on our way to an account of why it is pragmatically preferred. Still missing, though, is some notion of relevance. To see that this is a pressing problem, observe that the function QQ currently delivers incorrect results for scenario 2, in which player A knows that Barbara lives on Tallinskaja Street but still says merely that she lives in Moscow. In that scenario, $\mathrm{Quality}_C([\![\textit{Barbara lives in Russia}]\!]) = 1$. Since the values for P_I are unchanged in this scenario (I remains in the state of ignorance), the QQ value here is the same as in (27), namely, 1. Here is the value for $[\![\textit{Barbara lives on Tallinskaja Street}]\!]$ in this revised state of affairs:

(28) a. **The scenario 2 quality-rating for "Barbara lives on Tallinskaja Street"**

$$(P_A([\![\textit{Barbara lives on Tallinskaja Street}]\!]))^4 = 1$$

b. **The scenario 2 quantity-rating for "Barbara lives on Tallinskaja Street"**

$$\mathrm{inf}_I([\![\textit{Barbara lives on Tallinskaja Street}]\!]) = -\log_2 .125 = 3$$

c. **The scenario 2 QQ-rating for "Barbara lives on Tallinskaja Street"**

$$\mathrm{QQ}_{\text{scenario 2}}([\![\textit{Barbara lives on Tallinskaja Street}]\!]) = 1 \times 3 = 3$$

Thus, "Barbara lives on Tallinskaja Street" has a higher pragmatic value than "Barbara lives in Russia", the reverse of our intuitions. The missing ingredient is clearly that, given the shared goal of I and A, Barbara's street address is *too much* information in scenario 2 (and scenario 1). It isn't in accord with the players' choice for how to interpret question (2).

8 Relevance

Grice's relevance (relation) maxim depends entirely on our intuitive understanding of what relevance involves:

(29) **Relevance (Grice 1975)**

Be relevant.

The concept is left unanalyzed. What we would really like is to better understand the notion of *relevance to a question*, so we need to do a bit of question semantics. But not much. All we need to see in this regard is that the question "Where does Barbara live?" can, in our model (5), be viewed as semantically ambiguous between three salient partitions of W:

(30) a. **Street-level semantics**

$$[\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{cccc} \{w_1\} & \{w_2\} & \{w_3\} & \{w_4\} \\ \{w_5\} & \{w_6\} & \{w_7\} & \{w_8\} \end{array} \right\}$$

b. **City-level semantics**

$$[\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{cc} \{w_1, w_2\} & \{w_3, w_4\} \\ \{w_5, w_6\} & \{w_7, w_8\} \end{array} \right\}$$

c. **Country-level semantics**

$$[\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{c} \{w_1, w_2, w_3, w_4\} \\ \{w_5, w_6, w_7, w_8\} \end{array} \right\}$$

Answering a question means selecting part or whole of one or more of the partitions in the question. The more complete an answer is, the closer it comes to be identical to one of the questions cell's. The following definitions, due to van Rooy (2003b), get at this notion:

(31) a. $p_Q = \{q \in Q \mid q \cap p \neq \emptyset\}$ (for p an answer to question Q)

b. $\mathrm{Ans}(p, Q) = |p_Q|$

The set p_Q is the set of propositions in the question Q that are consistent with the proposition p (our answer). $\mathrm{Ans}(p, Q)$ is simply the cardinality of the set p_Q. A complete answer to Q has cardinality 1 by this measure. Partial answers have cardinalities greater than 1. (Only the empty-set answer has a cardinality of 0, so we ignore that case.) Here is an illustration using the city-level semantics in (30):

(32) $[\![\textit{Where does Barbara live?}]\!] = \left\{ \begin{array}{cc} \{w_1, w_2\} & \{w_3, w_4\} \\ \{w_5, w_6\} & \{w_7, w_8\} \end{array} \right\}$

a. $[\![\textit{Barbara lives on earth}]\!] = \{w_1 \ldots w_8\}$

$$\mathrm{Ans}([\![\textit{Barbara lives on earth}]\!], [\![\textit{Where does Barbara live?}]\!]) = 4$$

b. $[\![\textit{Barbara lives in Russia}]\!] = \{w_1 \ldots w_4\}$

$$\mathrm{Ans}([\![\textit{Barbara lives in Russia}]\!], [\![\textit{Where does Barbara live?}]\!]) = 2$$

c. $[\![\textit{Barbara lives in Moscow}]\!] = \{w_1, w_2\}$

$$\mathrm{Ans}([\![\textit{Barbara lives in Moscow}]\!], [\![\textit{Where does Barbara live?}]\!]) = 1$$

d. $[\![\textit{Barbara lives on Tallinskaja Street}]\!] = \{w_1\}$

$$\mathrm{Ans}([\![\textit{Barbara lives on Tallinskaja Street}]\!], [\![\textit{Where does Barbara live?}]\!]) = 1$$

The set of sample calculations (32) has one particularly important property: relative to the city-level question semantics, the proposition $[\![\textit{Barbara lives in Moscow}]\!]$ and the proposition $[\![\textit{Barbara lives on Tallinskaja Street}]\!]$ have identical Ans values. If we move to the country-level semantics (30)Q-sem-country, then the propositions

$[\![\textit{Barbara lives in Russia}]\!]$,

$[\![\textit{Barbara lives in Moscow}]\!]$, and

$[\![\textit{Barbara lives on Tallinskaja Street}]\!]$

all have Ans values of 1.

These observations are important for two central reasons. First, they show that Ans values are highly dependent upon the question under discussion. Second, they provide an important clue as to how we should explicate the relevance maxim. Ans ranks propositions according to how well they answer a given question, but it cannot be the only factor in calculating relevance, since scenario 2 distinguishes, e.g., "Barbara lives in Moscow" and "Barbara lives on Tallinskaja Street".

At the heart of the problem lies the observation that *too much* information can lead to a decrease in relevance (Sperber and Wilson 1995; van Rooy 2003c). If two propositions have identical Ans values, then we want to get rid of one of them. Following van Rooy (2003c:14), I resolve these ties by eliminating all but the *least* informative member of the set. This moves us to a total ordering on the set of propositions. The definition is (33).

(33) **Relevance-ranking**

The relevance-ranking of propositions with respect to a question Q is the numerical ordering induced by Ans, except that we throw out (do not rank) a proposition p iff there is a distinct proposition q with the same Ans value as p but a lower Quantity value than p.

In words: if two propositions answer the question under discussion equally well (have identical Ans values), then the more informative one is eliminated from the pragmatic competition entirely. With these overly-informative propositions removed, we obtain a total ordering: the relevance ranking for the context C containing question Q.

This definition permits us to come to grips with what happens in scenario 2, (4). There, player A is in a position to offer a more informative answer than she did. However, her actual answer does not lead to pragmatic infelicity. On the contrary, a more informative answer would have been less felicitous.

The crucial factor is that, in scenario 2, the players have asked the question "Where does Barbara live?" at the city level. We are thus in a situation like the one defined in (32), in which the the answer "Barbara lives in Moscow" and "Barbara lives on Tallinskaja Street" share the relevance value of 1. But "Barbara lives in Moscow" is less informative than "Barbara lives on Tallinskaja Street". By (33), "Barbara lives on Tallinskaja" is eliminated from pragmatic consideration — it is not even in the relevance ranking.

If we change the goals of the players, then we change the status of A's answer. For instance, suppose that the question "Where does Barbara live?" is interpreted instead at the street-level. Then "Barbara lives in Moscow" and "Barbara lives on Tallinskaja Street" have different relevance values. The first has a value of 2, the second of 1. So the two do not compete by the metric in (33). Since player A's quality-rating for "Barbara lives on Tallinskaja Street" is 1, and since it has high information content, it is favored over "Barbara lives in Moscow" in this revised scenario.

9 Overall pragmatic values

QQ-ratings and relevance-rankings provide two measurements of pragmatic felicity. I propose that we use these two measurements to arrive at a notion of *maximally-felicitous utterance*. To do this, let's first have a broad look at the values we obtain for various utterances in versions of scenarios 1 and 2:

(34) **Scenario 1 rankings**

	Quality		Quantity		QQ	Ans	relevance-ranking
B. lives on Tallinskaja	.004	$\times$	3	=	.012	1	—
B. lives on Pushkinskaja	.004	$\times$	3	=	.012	1	—
B. lives in Moscow	.0625	$\times$	2	=	.125	1	1
B. lives in Petersburg	.0625	$\times$	2	=	.125	1	1
$\star$ B. lives in Russia	1	$\times$	1	=	1	2	2
B. lives in the U.S.	0	$\times$	1	=	0	2	2

(35) **Scenario 2 rankings**

	Quality		Quantity		QQ	Ans	relevance-ranking
B. lives on Tallinskaja	1	$\times$	3	=	3	1	—
B. lives on Pushkinskaja	0	$\times$	3	=	0	1	—
$\star$ B. lives in Moscow	1	$\times$	2	=	2	1	1
B. lives in Petersburg	0	$\times$	2	=	0	1	1
B. lives in Russia	1	$\times$	1	=	1	2	2
B. lives in the U.S.	0	$\times$	1	=	0	2	2

The stars are endorsements, but not yet the selections of the theory. But these tables provide a clear indication of how to proceed. The starred candidates represent the highest possible QQ values, on the crucial assumption that utterances without relevance-rankings do not enter into pragmatic competition. (If they did, then "Barbara lives on Tallinskaja Street" would win in scenario 2 — its QQ value of 3 is unbeaten.)

Thus, I propose the following definition of *maximally felicitous utterances*:

(36) **Maximally-felicitous utterances**

Let C be a context with question under discussion Q and players s and a. The set of maximally felicitous utterances in C is set of utterances $U^{s \mapsto a}$ whose propositional content have the highest QQ-rating of all the relevance-ranked propositions in C.

According to this definition, speakers maximize QQ ratings, even if this entails a drop in relevance. The relevance-rankings are thus required only to settle competitions among equally-relevant propositions.

Speakers are expected to choose maximally-felicitous utterances. Hearers know this, and they can use this fixed assumption to arrive at the preferred interpretations for the utterances that they hear.

10 Summary and prospects

This paper was largely about taking measurements. I proposed integrated quality–quantity values as the primary measure of pragmatic worth. Relevance rankings further regulate the amount and kind of information that felicitous utterances can carry. In the interest of keeping this paper short, I did not do much beyond a basic analysis of two simple question–answer scenarios. But the overall framework suggests a number of more complex and challenging applications. For instance, probabilities seem to provide an excellent way of distinguishing regular updates from conversational implicature updates in a dynamic setting (section 4.1.2). The measurements taken in this paper seem ideal for understanding why phrases with low information content like the *not exactly ADJ* tend to receive *litotic* readings on which they leap to the opposite end of the scale (section 6.2). The framework also provides a novel perspective on pragmatic halos. And (arguably its greatest virtue) it provides a tool for identifying and measuring the relative felicity of bullshit.

References

Bach, Kent. 1999. The semantics–pragmatics distinction: What it is and why it matters. In Ken Turner, ed., *The Semantics–Pragmatics Interface from Different Points of View*, 65–84. Oxford: Elsevier. URL `http://online.sfsu.edu/ kbach/spd.htm`.

Beaver, David Ian. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford, MA: CSLI.

Beck, Sigrid and Hotze Rullmann. 1999. A flexible approach to exhaustivity in questions. *Natural Language Semantics* 7(3):249–298.

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*, Basingstoke, Hampshire. Oxford: Palgrave McMillan.

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

Blutner, Reinhard. 2000. Some aspects of optimality in natural language interpretation. *Journal of Semantics* 17(3):189–216.

Carnap, Rudolph. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena, and the syntax/pragmatics interface. In Adriana Belletti, ed., *Structures and Beyond*. Oxford: Oxford University Press. [The manuscript began circulating in 2001].

Devitt, Michael and Kim Sterelny. 1987. *Language and Reality: An Introduction to the Philosophy of Language*. Cambridge, MA: MIT Press.

Fauconnier, Gilles. 1975. Polarity and the scale principle. In *Papers of the Eleventh Regional Meeting of the Chicago Linguistic Society*, 188–199. Chicago: Chicago Linguistic Society.

Frankfurt, Harry G. 1986. On bullshit. *Raritan* 6(2):81–100. Reprinted in Frankfurt (1988), 117–133.

Frankfurt, Harry G. 1988. *The Importance of What we Care About*. Cambridge University Press.

Gazdar, Gerald. 1979a. *Pragmatics: Implicature, Presupposition and Logical Form*. New York: Academic Press.

Gazdar, Gerald. 1979b. A solution to the projection problem. In Choon-Kyu Oh and David A. Dinneen, eds., *Syntax and Semantics*, Volume 11: Presupposition, 57–89. New York: Academic Press.

Geurts, Bart and Emar Maier. 2003. Layered DRT. URL `http://www.kun.nl/phil/tfl/bart/papers/ldrt.pdf`, Ms., University of Nijmegen.

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.

Grice, H. Paul. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University 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* 5:175–233.

Groenendijk, Jeroen and Martin Stokhof. 1988. Type-shifting rules and the semantics of interrogatives. In Gennaro Chierchia, Barbara Partee, and Raymond Turner, eds., *Properties, Types and Meaning*, Volume 2, 21–68. Dordrecht: Kluwer Academic Publishers.

Hara, Yurie. 2005. Modal-meaning as expressive, Paper to be presented at Sinn und Bedeutung 10, Berlin.

Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press. Reissued 2001 by CSLI.

Jäger, Gerhard. 2002. Some notes on the formal properties of bidirectional Optimality

Theory. *Journal of Logic, Language and Information* 11(4):427–451.

Krifka, Manfred. 2003. Polarity items in questions, Invited talk, Korean Society of Logic and Information, Inha University, Incheon, Korean, June 28.

Krifka, Manfred. 2005. Approximate interpretation of number words: A case for strategic communication. In Irene Vogel and Jost Zwarts, eds., *Cognitive Foundations of Communication*. Amsterdam: Koninklijke Nederlandse Akademie van Wetenschapen.

Lasersohn, Peter. 1999. Pragmatic halos. *Language* 75(3):522–551.

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

Majewski, Helen. 2005. Polarity-sensitive *exactly*, Ms., University of Massachusetts, Amherst.

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/`.

Merin, Arthur. 2005. Probabalistic and deterministic presuppositions. In Rainer Bäuerle, Uwe Reyle, and Thomas Ede Zimmermann, eds., *Presuppositions and Discourse: Essays Offered to Hans Kamp*. Amsterdam: Elsevier. URL `http://semanticsarchive.net/Archive/GE0MGM5N/`.

van Rooy, Robert. 2003a. Negative polarity items in questions: Strength as relevance. *Journal of Semantics* 20:239–273.

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

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

van Rooy, Robert. 2004a. Relevance of complex sentences, Ms., University of Amsterdam.

van Rooy, Robert. 2004b. Signalling games select Horn strategies. *Linguistics and Philosophy* 27(4):493–527.

Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27(3):367–391.

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

Zimmermann, Malte. 2004. Zum wohl: Diskurspartikeln als Satztypmodifikatoren. *Linguistische Berichte* 199:1–35.