Knowledge and Implicature: Modeling Language Understanding as Social Cognition

Noah D. Goodman,[a] Andreas Stuhlmüller[b]

[a]*Department of Psychology, Stanford University*
[b]*Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology*

Abstract

Is language understanding a special case of social cognition? To help evaluate this view, we can formalize it as the *rational speech-act* theory: Listeners assume that speakers choose their utterances approximately optimally, and listeners interpret an utterance by using Bayesian inference to "invert" this model of the speaker. We apply this framework to model scalar implicature ("some" implies "not all," and "N" implies "not more than N"). This model predicts an interaction between the speaker's knowledge state and the listener's interpretation. We test these predictions in two experiments and find good fit between model predictions and human judgments.

Keywords: Language; Bayesian model; Scalar implicature

To what extent does language understanding rely on extra-linguistic knowledge and processes? One view of language processing suggests that it consists of largely separate, special-purpose faculties; another view suggests that it depends critically on domain-general inference mechanisms, and even on intuitive theories that are not specific to language. Indeed, many thinkers have viewed speech as an action with communicative goals, such as informing a listener (Clark, 1996; Grice, 1975; Levinson, 2000). A listener making this assumption can make stronger inferences than an utterance would allow from its literal meaning—pragmatic effects can strengthen, or change, the interpreted meaning.

Recent work has aimed to formalize the social inference view of pragmatics using tools of Bayesian statistics and information theory (Frank & Goodman, 2012); we refer to this formal framework as the *rational speech-act* theory of language understanding. It views pragmatic competence as following naturally from an intuitive theory of speech production, which in turn is a special case of intuitive theory of mind. More precisely, listeners have an internal model that describes speakers as choosing their utterances approximately optimally

on the basis of certain social goals, such as conveying information to the listener; listeners then interpret an utterance by using Bayesian inference to "invert" this model of the speaker, drawing conclusions about the world state and speaker's intention from the utterance and any other relevant world knowledge. These two rationality assumptions, for listener and speaker, have a role similar to those made in ideal observer models of perception (Geisler, 2003): They provide a starting point for a quantitative understanding of the complex interactions involved in language understanding. Indeed, this view has provided good quantitative models for pragmatic inference in a number of simple settings (Frank & Goodman, 2012; Frank, Goodman, Lai, & Tenenbaum, 2009).

Because rational decision making predicts that action selection is related to the *expected* utility—a quantity that depends on the actor's belief distribution—a listener who views the speaker as rational should be sensitive to the speaker's belief state. The rational speech-act theory thus predicts an interaction between (shared) knowledge about a speaker's knowledge state and a listener's interpretation of his utterance. This is a very general prediction of the framework, which could easily prove to be false—if pragmatic inference is a highly modularized computation, for instance, we would not expect such general knowledge to affect it. Deriving and testing precise predictions about this interaction thus provide an important test of the rational speech-act theory.

If you hear "some of the apples are red," you will infer that not *all* of the apples are. Pragmatic effects of this sort are called scalar implicatures (Horn, 2004), and they provide a window on the interactions between the language faculty and general cognition. Consider Fig. 1: If the speaker has seen all the apples, his utterance would be interpreted as "some, but not all, of the apples are red." However, if the speaker had only looked at two of the apples, the listener might draw a different conclusion. Indeed, we show below that the implicature "not all" can be canceled by facts about the speaker's perceptual access. This interaction between language understanding and general knowledge is not predicted by strongly modular theories that place scalar implicature within a semantics module (Chierchia, Fox, & Spector, 2011). We show further that the interaction of knowledge and implicature is fine grained: The details of a speaker's belief distribution affect the details of an implicature.

This article provides a formal model of the pragmatic inference that leads to scalar implicatures, building on the rational speech-act framework. We directly model the possibility that the speaker may have incomplete knowledge and the effects this has on the listener's interpretation. We derive predictions of this model for interpretation of the quantifier "some" and the number words ("one," "two," etc.). The model both explains the standard implicatures for these words and predicts that these implicatures can be canceled, completely or partially, when the speaker has incomplete knowledge. We test these predictions in two experiments and find good fits to human judgements, both qualitatively and quantitatively.

1. A rational speech-act model

We view language comprehension as a rational inference based on an intuitive theory of language production. Our setting is illustrated in Fig. 1. The listener infers the world

state, s, given the speaker's utterance, w, and shared information about the speaker's (possibly incomplete) information access, a. By Bayes' rule:

$$P_{\text{listener}}(s \mid w, a) \propto P_{\text{speaker}}(w \mid s, a)P(s), \tag{1}$$

where $P(s)$ captures the listener's prior beliefs about the world state and $P_{\text{speaker}}(w \mid s, a)$ describes the listener's intuitive theory of how the speaker chooses words.

The speaker chooses an utterance in accord with Bayesian decision theory (Berger, 1985): She acts to approximately optimize expected utility. Imagine a speaker who makes observation o about the true state of the world. (For instance, in Fig. 1, the speaker has perceptual access to two of three apples and observes that those two are red). The speaker selects an utterance w to convey information about the world state to a listener and does so by soft-max optimizing expected utility:

$$P_{\text{speaker}}(w \mid o, a) \propto \exp(\alpha \mathbb{E}_{P(s \mid o, a)}[U(w; s)]). \tag{2}$$

The speaker's utility function, $U(w;s)$, captures the value of saying w if the world is actually s. The expectation is taken over the speaker's belief state, $P(s \mid o, a)$, because the speaker may still be uncertain about the state of the world. The parameter α controls the deviation from optimality.

So far, nothing in the model is unique to language—indeed, similar models have been used to model social cognition more generally (Baker, Saxe, & Tenenbaum, 2009; Goodman, Baker, & Tenenbaum, 2009; Ullman et al., 2009). To capture a motivation to be informative, utility must be related to the information conveyed in the utterance.[1] More specifically, utility is related to the amount of information that a *literal* listener would not yet know about state s after hearing it described by utterance w—this is the negative *surprisal* (Cover & Thomas, 1991):

$$U(w; s) = \ln(P_{\text{lex}}(s \mid w)), \tag{3}$$

where the *literal interpretation* probability $P_{lex}(s \mid w)$ is determined by the lexicon—here, we will assume that each utterance has a truth function, $F_w : s \mapsto \{0, 1\}$, and the distribution is otherwise uninformative: $P_{\text{lex}}(s \mid w) \propto \delta_{F_w(s)}$.

We assume that the speaker's access a is common knowledge of speaker and listener, but the listener still does not know what observation the speaker made, hence:

$$P_{\text{speaker}}(w \mid s, a) = \sum_{o} P_{\text{speaker}}(w \mid o, a)P(o \mid a, s). \tag{4}$$

In the experiments below (as in Fig. 1), the state of the world is always a set of objects that may have a given property and observations consist of looking at a subset of the objects. Thus, the observation probability $P(o \mid a, s)$ is a hyper-geometric distribution (i.e., the probability of drawing N balls of a given color, without replacement, from an urn containing a given set of colored balls). In this setting, it is also reasonable to assume
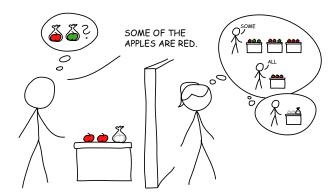
Fig. 1. How will the listener interpret the speaker's utterance? How will this change if she knows that he can see only two of the objects?

that the prior probability, $P(s)$, is a binomial distribution (i.e., draws *with* replacement); we will initially assume so and will later measure the distribution empirically.

Overall, the above equations describe the inferences that a rational listener will make to comprehend a speaker that she believes to be approximately rational and have a goal to be informative. Importantly, these inferences depend on shared knowledge about the aspects of the world to which the speaker has access. Thus, the rational speech-act theory predicts that the speaker's access affects utterance interpretation; we test this prediction below. To derive more specific predictions, we must describe the set of alternative utterances.

1.1. The alternatives

We have assumed that the interpretation of an utterance is made with respect to a set of alternative utterances. These alternatives could be all possible utterances or could be a limited set generated by replacing key words in the actual utterance with related words. The alternatives may, for instance, be generated by a grammatical mechanism as in Fox and Katzir (2011). For our results, the details of this process are unimportant; what is important is that there exists a set of alternative expressions and a (truth-functional) literal meaning for each. We make standard assumptions in both respects.

Consider the case of the quantifier words "some" and "all." Under the standard semantics, "all the balls are red" is true exactly when N of the N balls are red, while "some of the balls are red" is true when $M \geq 1$ of the N balls are red. In particular, the literal meaning of "some" allows the state where all N balls are red. We use a lexicon that consists of the standard meanings for "none," "some," and "all." Model predictions are shown in Fig. 2A for the listener's interpretation of "some of the balls are red" when there are three objects, under varying conditions of speaker access.

When the speaker has perceptual access to three of the three objects (hence complete knowledge) and says "some," there is a lower probability on three than two—this is the standard "some but not all" implicature. To understand why the model predicts this implicature, we can first simplify the speaker model: In a complete knowledge situation, the

speaker's belief distribution is concentrated on the true state, so Eqs. 2, 3, and 4 become:

$$P_{\text{speaker}}(w \,|\, s, a) \propto \exp(\alpha \ln(P_{\text{lex}}(s \,|\, w)))$$
$$\propto (P_{\text{lex}}(s \,|\, w))^{\alpha}. \tag{5}$$

Using the standard meanings described above, we see that if the state were 3, the speaker would be more likely to say "all" ($P_{\text{lex}}(s \,|\, w) = 1$) than "some" ($P_{\text{lex}}(s \,|\, w) = \frac{1}{3}$); conversely, if the state were 2, the speaker would say "some," as the probability for "all" is now $P_{\text{lex}}(s \,|\, w) = 0$. Now, consider Eq. 1 and imagine for the moment a uniform prior over states. In this case, the listener will infer each state s in proportion to how likely the speaker was to say "some" given this state. Overall, this leads to the implicature that "some" is unlikely to be interpreted as 3—"some but not all."

In contrast, when the speaker has only partial access, the calculation is more complex, involving the inferred belief distribution of the speaker. Comparing across the three panels of Fig. 2A, we see the probability of 3 is much higher when access is 1 or 2 than when it is 3. When access is 1, no implicature is predicted (the probability of 3 is approximately the same as the probability of 2); when access is 2, only a very slight implicature. Overall, we predict that incomplete speaker knowledge can *cancel* the standard "some but not all" implicature.

The case of numerals ("one," "two," ...) is similar but more subtle. It has been argued that number words have a lower bound meaning (Horn, 1972) (e.g., "two balls are red" means $M \geq 2$ of the balls are red), and the intuitive, exact, meanings arise as a pragmatic implicature—"one but not two, three, etc." In Fig. 2B, we show model predictions based on the lower bound semantics for number words, varying speaker's access. We see that exact meanings do arise as an inference when the speaker has complete access, but there is an interaction: Number words do not receive an exact interpretation when the speaker has incomplete knowledge. Of particular interest is the case where the speaker has seen two of three objects and says "one": here a *partial* implicature is predicted, with the probability of 3 low, but 1 and 2 high.

2. Experiment 1

Because a rational speaker chooses actions based on *expected* utility, the rational speech-act model predicts an effect of speaker's knowledge on the listener's interpretation of "some" statements. We tested the predictions of the model by putting participants in the role of the listener and asking them to judge the state of the world in scenarios where perceptual access (and hence knowledgeability) of the speaker varied.

2.1. Participants, materials, and methods

Fifty participants were recruited through Amazon's Mechanical Turk crowd-sourcing service and completed the experiment for a small payment.

We constructed six scenarios in which a speaker had three objects that could have (or not) a given property. The speaker then made a statement indicating the number of objects he or she had looked at and a quantified ("some") statement. We split each scenario into setup and speech-act phases. The setup phase named the speaker and described the objects and the relevant property. For example:

> Letters to Laura's company almost always have checks inside. Today Laura received 3 letters.

Because our model predicts greater effects when the a priori base rate of the property is high (otherwise it is difficult to tell an implicature from an a prior belief that it is unlikely for all objects to have the property), we describe all properties as "almost always" holding. To make sure participants attended to the setup, we asked them to report the *a priori* probability that 0, 1, 2, or 3 objects had the property:

> How many of the 3 letters do you think have checks inside?

The speech-act phase introduced a speech act in which the speaker both declared how many of the objects he or she had observed and stated that some objects had the property:

> Laura tells you on the phone: "I have looked at 2 of the 3 letters. Some of the letters have checks inside."

We then again elicited judgements about how many objects had the property:

> Now how many of the 3 letters do you think have checks inside?

Finally, because the speech act might not be a perfect manipulation of speaker's knowledgeability (for instance, the speaker may have gained knowledge by another route), we elicited this directly:

> Do you think Laura knows exactly how many of the 3 letters have checks inside?

Each response was given by a betting measure: Participants were instructed to divide "$100" among the options, betting to indicate their confidence in each option. For the first two questions, there were four options (0–3 of the objects have the property), and for the final question, there were two options (the speaker does/does not have complete knowledge). Before the experiment began, participants were given a brief warm-up, using unrelated questions, to familiarize them with the betting measure.

Each scenario existed in forms varying speaker access (the number of objects the speakers had looked at) from 1 to 3. Each participant saw each access condition once, in random order, presented using randomly chosen scenarios (with no duplicate scenarios).

In terms of our predictions, we have two partial-knowledge conditions (where we expect cancelation of the implicature) and one complete-knowledge "control" (where we expect the standard implicature).[2]

2.2. Qualitative results

There was no effect of scenario, so we collapse across this factor in all analyses. As expected based on related work (Goodman et al., 2009), the speaker's access statement (e.g., "I have looked at 2 of the 3 letters") was not a perfect manipulation of knowledge-ability: In the partial access conditions, some participants judged that the speaker was likely to know exactly how many objects had the property. (The bet that the speaker had complete knowledge when access = 1 was $M = 27.1$, $SD = 4.9$; when access = 2, $M = 34.8$, $SD = 5.7$; when access = 3, $M = 93.0$, $SD = 2.7$.) As we were interested in the effects of varying knowledgeability, we initially exclude trials in which the knowledge judgement was less than 70 in the expected direction (we come back to the full data set in the quantitative analysis below). Fig. 2C shows the mean of bets on each option, as access varied. As predicted, there was an effect of speaker's access on listener's interpretation (one-way ANOVA with bets on 3 as dependent variable, $F(2, 102) = 10.18$, $p < .001$).

We next performed our preplanned comparisons to check that an implicature was drawn when the speaker had complete knowledge, but not when the speaker had partial knowledge. In the complete access condition, bets on 3 were less than bets on 2 (paired, directional t test, $t(43) = -10.2$, $p < .001$). In the partial access conditions, the implicature was canceled: Bets on 2 did not exceed bets on 3 when speaker had access to 1 object (paired, directional t test, $t(31) = 0.77$, $p = .78$) or when access was 2 (paired, directional t test, $t(28) = -0.82$, $p = .21$). If we look at just bets on 3, we see significantly lower bets in the complete access condition than the access 1 condition (unpaired, directional t test, $t(47) = -4.0$, $p < .001$) or the access 2 condition (unpaired, directional t test, $t(43) = -3.5$, $p < .001$). While there was no significant implicature in either partial-access condition, there is a slightly greater tendency toward implicature in the access 2 condition than the access 1 condition, as predicted by the model (two-way ANOVA with bet as dependent variable, and access (1 or 2) and state (2 or 3) as independent variables, $F(2, 294) = 3.77$, $p < .05$).

Thus, the knowledge of the speaker affected listener's interpretation of "some" in the way predicted by the rational speech-act model. We examine the quantitative fit of the model below.

3. Experiment 2

In Experiment 2, we tested the predictions of the rational speech-act model for interpretation of numerals. We again expect to find an effect of speaker's knowledge, but in this case, there is a more detailed interaction: The implicature should be canceled when

the speaker says "one" after seeing only one object and when the speaker says "two" after seeing two objects, but it should only be partially canceled when the speaker says "one" after seeing two objects—this implies a fine-grained interplay between the speaker's knowledge state and the interpretation of her utterance.

3.1. Participants, materials, and methods

Fifty participants were recruited through Amazon's Mechanical Turk crowd-sourcing service and completed the experiment for a small payment.

We used the same stimuli as in Experiment 1, modifying the scenarios only in the speech act: The speaker now made a statement indicating the number of objects he or she had looked at and the number that had the property. For instance:

> Laura tells you on the phone: "I have looked at 2 of the 3 letters. 1 of the letters has checks inside."

Each scenario existed in forms varying the speaker's access, from 1 to 3, and the number word the speaker used, from 1 to 3; we limited to sensible situations where the word used was no greater than the number of objects seen. Hence, we had six conditions, with access/word: 1/1, 2/1, 2/2, 3/1, 3/2, and 3/3. In terms of our predictions, we have three partial-knowledge conditions (where we expect partial or complete cancelation of the implicature) and three complete-knowledge "controls" (where we expect the standard implicatures). The order of scenarios and the order of conditions were randomized between participants.

3.2. Qualitative results

There was again no effect of scenario, so we collapse across this factor. As in Experiment 1, the speaker's access statement was not a perfect manipulation of knowledgeability (bets that speaker had complete knowledge in partial-access conditions: $M = 42.0$, $SD = 3.4$, in complete access conditions: $M = 92.1$, $SD = 1.6$). We once again limit to trials with the expected judgements of knowledgeability (with a threshold of 70). The mean of participants' bets are shown in Fig. 2D. To evaluate the overall effect of access, we performed an ANOVA with access and word as independent measures and bet on 3 as dependent measure. We find a main effect of access ($F(2, 205) = 6.57$, $p < .01$), an interaction between word and access ($F(1,205) = 34.7$, $p < .001$), and a main effect of word ($F(2, 205) = 269.8$, $p < .001$).

We then explored the results in more detail using planned comparisons to test whether implicatures were drawn (only) when predicted. We found an implicature in the complete access conditions: When the speaker said "two," bets on state 3 were less than on state 2 (paired, directional t test, $t(43) = -10.2$, $p < .001$). When the speaker said "one," bets on state 1 were greater than on state 3 (paired, directional t test, $t(42) = -13.1$, $p < .001$) or state 2 (paired, directional t test, $t(42) = -17.1$, $p < .001$). In contrast, there was no

implicature when access was 1 and the speaker says "one"—bets on 1 were not greater than on 2 (paired, directional t test, $t(24) = 1.9$, $p = .96$) or on 3 (paired, directional t test, $t(24) = 3.2$, $p = 1.0$)—and no implicature when access is 2 and the speaker says "two"—bets on 2 were not greater than on 3 (paired, directional t test, $t(24) = 1.1$, $p = .87$). When access is 2 and the speaker says "one" we found the predicted *partial implicature*: Bets on state 1 were significantly greater than on state 3 (paired, directional t test, $t(25) = -3.9$, $p < .001$), but not on state 2 (paired, directional t test, $t(25) = 1.5$, $p = .92$).

These results again support the predictions of the rational speech-act model, showing not merely an interaction between the speaker's knowledge and the listener's tendency to draw an implicature, but a fine-grained interaction that is unlikely to result from a modular process of language understanding. In addition, these results support the standard, but controversial (Barner & Bachrach, 2010; Huang, Snedeker, & Spelke, 2004), view that number words have a "lower-bound" semantics which is only strengthened into an exact meaning by pragmatic inference.

4. Quantitative model comparison

To evaluate the quantitative model predictions, we first compare model predictions with mean human ratings for the subset of trials in which participants gave knowledgeability ratings in the expected direction (>70, as above). As in the model description, we assume a binomial prior distribution. We fit the prior base rate parameter and the α parameter by minimizing the root mean squared error (RMSE) of the model predictions and mean data for both experiments. The resulting model fit is RMSE = 9.01 and $r = .96$. The model predictions with the best-fitting parameters are shown in Fig. 2A,B and show striking correspondence with the human data in Fig. 2C,D.

To consider all responses, including those previously removed due to unexpected knowledge judgements, we extend the model by including an additional knowledge parameter: the probability that the speaker did in fact have complete knowledge (regardless of perceptual access). As we measured prior expectations and knowledgeability in each trial, we compute model predictions, trial by trial, using these values. Because the betting interface encouraged participants to round small bets to 0, but probability 0 is very different than a non-zero small probability to the model, we changed 0 and 100 bets to 1 and 99. In addition, we assumed a simple power–law relationship between subjective probability and bets (equivalent to the soft-max used to model the speaker, Eq. 2). The parameter of this power law and the speaker optimality, α, were then fit by minimizing the RMSE of mean model predictions to mean human judgements in both experiments. The resulting fit was again good: RMSE = 8.36, $r = .95$. Looking at individual participants, we find median correlation of $r = .75$ between a participant's judgements and the model predictions based on his or her prior and knowledge scores. These results suggest that the rational speech-act model is able to capture the quantitative pattern of people's judgements both across the population and within individual participants.

Model   Experiment
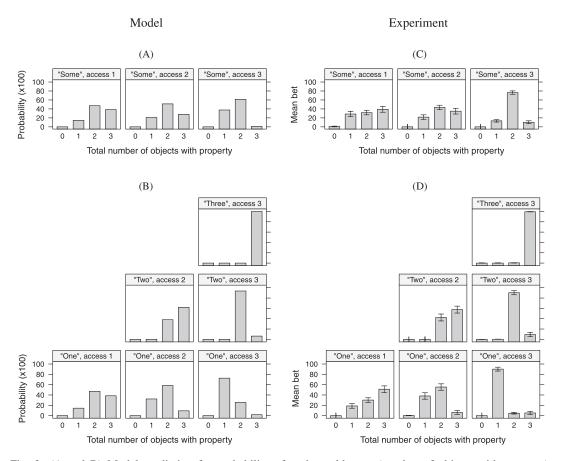

Fig. 2. (A and B) Model prediction for probability of each world state (number of objects with property), varying the word the speaker used and the speaker's perceptual access. The prior is assumed to be binomial with base rate 0.62, and the speaker optimality parameter is set to $\alpha = 3.4$. (C and D) Mean participant bet on each world state, varying the word the speaker used and the speaker's perceptual access. Data have been filtered to include only trials where the participant's bet that the speaker had complete knowledge was greater than 70 in the expected direction. Error bars are standard error of the mean.

5. Conclusion

We have described a rational speech-act model of scalar implicatures and their interaction with speaker knowledge. This model formalizes language understanding as social cognition, with language-specific goals and actions, using the tools of Bayesian statistics. In addition to predicting the standard implicatures ("some but not all") as an inference that depends on the alternative utterances, this model predicted that these implicatures could be canceled, completely or in part, when it was common knowledge that the speaker had incomplete knowledge. Experiments 1 and 2 verified these qualitative predictions and showed tight quantitative fits with the model.

The predicted interaction between interpretation and speaker's knowledge was not a peculiarity of this set of words or of scalar lexical items; instead, it follows from the general fact that rational decision makers must choose actions based on their expected utility. In contrast, a strongly modular model of implicatures would predict no such interactions. One could amend these theories (Chierchia et al., 2011) to allow a speaker's ignorance to affect the implicature mechanism. To predict the fine-grained interactions demonstrated for number words, additional machinery would be required, describing how the details of a speaker's knowledge state influence a listener's interpretation. In contrast, we have shown that the rational speech-act framework parsimoniously predicts these effects from high-level assumptions about the goals of listener and speaker.

The rational speech-acts framework we have used here is closely related to that of game-theoretic pragmatics (Jäger, in press), and particularly to the use of lifted games to capture epistemic effects (Franke, 2009). There are two principal differences of game-theoretic approaches from ours: In those approaches, it is assumed that speakers fully optimize ($\alpha \rightarrow \infty$ in our framework), and that they carry out deeply recursive social reasoning—"I think, that you think, that I think, that..."—while we assume only one such level of reasoning. The quantitative fits we have shown suggest that limited recursion and optimization are psychologically realistic assumptions. Future work will be needed to explore all the possible models in this space.

Our results support the rational speech-act framework for modeling pragmatics. More generally, they further boost the momentum building for quantitative models of language as a branch of rational social cognition. In the words of Grice (1975): "One of [our] avowed aims is to see talking as a special case or variety of purposive, indeed rational, behavior..."

Acknowledgments

Notes

1. Other communicative motivations could be added to this utility, such as a complexity term influencing the manner of expression.
2. Experiment 1 may be viewed at http://goo.gl/3S5zz, Expt. 2 at http://goo.gl/iSc6o.

References

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*, 329–349.

Barner, D., & Bachrach, A. (2010). Inference and exact numerical representation in early language development. *Cognitive Psychology*, *60*(1), 40–62.

Berger, J. (1985). *Statistical decision theory and bayesian analysis*. New York: Springer.

Chierchia, G., Fox, D., & Spector, B. (2011). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In C. Maienborn, K. von Heusinger, and P. Portner (Eds.), *Semantics: An international handbook of natural language meaning*. Berlin: Mouton de Gruyter.

Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.

Cover, T., & Thomas, J. (1991). *Elements of information theory*. New York: Wiley Online Library.

Fox, D., & Katzir, R. (2011). On the characterization of alternatives. *Natural Language Semantics*, *19*(1), 87–107.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336* (6084), 998.

Frank, M. C., Goodman, N. D., Lai, P., & Tenenbaum, J. B. (2009). Informative communication in word production and word learning. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society*. Austin, TX: Cognitive Science Society.

Franke, M. (2009). Signal to act: Game theory in pragmatics, Unpublished doctoral dissertation, Institute for Logic, Language and Computation.

Geisler, W. S. (2003). Ideal observer analysis. In L. Chalupa & J. Werner (Eds.), *The visual neurosciences* (pp. 825–837). Cambridge, MA: MIT Press.

Goodman, N. D., Baker, C. L., & Tenenbaum, J. B. (2009). Cause and intent: Social reasoning in causal learning. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society*.

Grice, H. (1975). Logic and grammar. In D. Davidson & G. H. Harman (Eds.), *Readings in language and mind*. (pp. 64–75). Encino, CA: Dickenson.

Horn, L. (1972). On the semantic properties of logical operators in english, Unpublished doctoral dissertation, UCLA.

Horn, L. (2004). Implicature. *The Handbook of Pragmatics*, 2–28.

Huang, Y., Snedeker, J., & Spelke, E. (2004). What exactly do numbers mean. In K. Forbus, D. Gentner, and T. Regier (Eds.), *26th annual meeting of the cognitive science society,* p. 1570. Chicago: Cognitive Science Society.

Jäger, G. (in press). Game theory in semantics and pragmatics. In C. Maienborn, P. Portner, & K. von Heusinger (Eds.), *Semantics: An international handbook of natural language meaning* Berlin: De Gruyter Mouton.

Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge, MA: The MIT Press.

Ullman, T., Baker, C. L., Macindoe, O., Evans, O., Goodman, N. D., & Tenenbaum, J. B. (2009). Help or hinder: Bayesian models of social goal inference. In Y. Bengio et al. (Eds.), *Advances in neural information processing systems 22*. Vancouver: NIPS Foundation.