Rational speech act models of pragmatic reasoning in reference games

Michael C. Frank

Department of Psychology, Stanford University

Andrés Gómez Emilsson

Department of Psychology, Stanford University

Benjamin Peloquin

Symbolic Systems Program, Stanford University

Noah D. Goodman

Department of Psychology, Stanford University

Christopher Potts

Department of Linguistics, Stanford University

Author Note

Thanks to Alex Stiller for work on a previous version of the “base rate” experiment, to Roger Levy for valuable discussion and the design of the “odd man out” stimulus, and to Ally Kraus for designing the stimuli. Data from Experiment 8 were presented in Vogel, Emilsson, Frank, Jurafsky, and Potts (2014). We gratefully acknowledge ONR N00014-13-1-0287, NSF #1456077, and the Merck Family Foundation. Please address correspondence to Michael C. Frank, Department of Psychology, Stanford University, 450 Serra Mall (Jordan Hall), Stanford, CA, 94305, tel: (650) 724-4003, email:mcfrank@stanford.edu.

Abstract

Human communication is almost always ambiguous, but it typically takes place in a context where this ambiguity can be resolved. A key part of this process of disambiguation comes from pragmatic reasoning about alternative messages that a speaker could have said in that context. Following previous work, we describe pragmatic inference as recursive reasoning – in which listeners reason about speakers and vice versa – using a “rational speech act” (RSA) model. We then systematically test the parameters and design decisions of this model through a series of ten experiments using one-shot reference-resolution games ($N = 7,569$). Such games present a valuable and tractable microcosm for studying broader questions about communication in context, and human behavior within them can be described using the RSA framework.

Rational speech act models of pragmatic reasoning in reference games

Human communication is astonishingly flexible and efficient. Between two people who know each other well, a single word, look, or gesture can convey volumes (Sperber & Wilson, 1986; Clark, 1996). Even two people who barely know one another can rapidly leverage their shared vocabulary to create new, economical ways of communicating about novel situations (Brennan & Clark, 1996; Clark & Brennan, 1991). These successes are all the more impressive considering the systematic issues of vagueness and ambiguity in natural language (Keefe & Smith, 1997; Wasow, Perfors, & Beaver, 2005). How do we use such slippery means to produce such concrete results?

Grice (1975) provided a critical insight into this problem, describing how listeners reason about speakers' language production goals. He posited that, by assuming that speakers were attempting to be cooperative in their communication, listeners could then infer the most likely intended goal that would have given rise to the speaker's particular utterance, given the current circumstances. In addition to introducing the idea of nested reasoning (listeners about speakers and speakers about listeners), this proposal allowed the separation of semantic content – those aspects of meaning that are invariant across contexts – from pragmatic content – those aspects of meaning that rely on contextual inferences about speakers' intentions.

Although many subsequent analysts have attempted to refine Grice's original proposal or even to reject substantial elements, the semantics/pragmatics distinction has been critical for progress in understanding language comprehension. Pragmatic phenomena in general have become increasingly central to research in language and cognition, following Bar-Hillel's (1971) call to rescue phenomena from the "pragmatic wastebasket." In some substantial sense, what linguists call "pragmatics" is what language users experience as the role of language in their everyday interpersonal communication – the reality of how context alters interpretation.¹ Thus, an important goal for research in linguistics and the

¹Other writers have explored pragmatic phenomena beyond Grice's initial focus on inferences that fol-

psychology of language is the development of formal tools for understanding pragmatic inferences and bringing them under experimental control.

In this article, we provide an extended presentation and evaluation of a formal framework for understanding pragmatics, known as the “rational speech act” (RSA) framework (originally introduced in Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013). This framework builds on Grice’s initial insight and combines it with tools from Bayesian cognitive modeling (Tenenbaum, Kemp, Griffiths, & Goodman, 2011), with the aim of building a quantitative understanding of social goal inference in communication. Evaluating this framework requires developing tools to make quantitative measurements of pragmatic inferences, so we present a set of experiments that uses simple signaling games as a framework for making these measurements. In the remainder of this introduction, we provide a brief survey of previous frameworks for pragmatic reasoning (for further review and discussion, see Goodman & Frank, in press) and give an informal introduction to RSA models, and then summarize the plan for the rest of the article.

Modeling pragmatic reasoning as rational action

Our conception of pragmatic enrichment – how the literal semantics of an utterance are modulated in context by a consideration of the speaker’s communicative goals – was initially established by Grice (1975). Grice’s original formulation of pragmatic inference is typically summarized with reference to the specific pragmatic *maxims* he proposed. These maxims, a set of normative guidelines that speakers are assumed to follow (be informative, truthful, relevant, and perspicacious), have been immensely influential. Despite this popularity, the maxims themselves do not figure heavily in contemporary empirical or theoretical work in linguistics and psycholinguistics (consider e.g., Chemla & Spector, 2011; Breheny, Ferguson, & Katsos, 2013); they are considered by some theorists to be a dead

low from a single speaker’s utterance (e.g., Clark, 1996). While noting the importance of discourse- and conversation-level pragmatics, here we pursue Grice’s original vision and focus on contextual inferences from individual utterances. Our hope is that this understanding will contribute to a broader theory of discourse phenomena.

end as a method for deriving pragmatic phenomena (e.g., Frederking, 1996). Indeed, though the general Gricean program has been advanced by a number of influential and important theories, none of these have incorporated the initial set of maxims as a central part (Levinson, 2000; Sperber & Wilson, 1986; Clark, 1996). For example, Sperber and Wilson (1986) construct their theory around a single maxim, that of relevance.

But Grice may not have intended the specific maxims he proposed to be the last word in the analysis of pragmatic phenomena. Instead, the maxims are a way to instantiate his more general proposal: a construal of language use as a subclass of rational action (“a special case or variety of purposive, indeed rational, behavior”; p. 47). This insight has arguably made a more lasting contribution, since grouping language together with other rational actions allows an analysis of linguistic behavior using a set of tools that have proven immensely powerful in other domains. These tools stem from the theory of rational economic action (Von Neumann & Morgenstern, 2007), and allow an analysis of the costs and benefits of individual linguistic actions relative to one another. The rational action perspective on linguistic pragmatics has been explored extensively in several different, and largely non-intersecting literatures.

First, researchers interested in the automatic generation of natural language referring expressions explored a utility-theoretic view of the Gricean framework. Dale and Reiter (1996) write “Ultimately... Grice’s maxims taken collectively mean ‘Don’t include elements that don’t do anything.’ Our position is that, under a goal-oriented view of language generation, there is no need to explicitly follow such a directive at all; the desired behavior just falls out of the mechanism.” They suggested that being informative, truthful, relevant, and perspicacious would emerge from the process of trying to communicate a particular meaning with maximal efficiency. This perspective inspired and motivated a tradition of work on systems for creating pragmatically-felicitous and practically-effective referring expressions (e.g., Dale & Reiter, 1995; Reiter, Dale, & Feng, 2000; Viethen & Dale, 2006). More recently, work in this tradition has begun to connect with the empirical literature on

language production (Van Deemter, Gatt, Sluis, & Power, 2012; Van Deemter, Gatt, Van Gompel, & Krahmer, 2012).²

Second, the general connection to game theory and rational action has been expanded productively in the theoretical linguistic literature. Following early work by Parikh (1991), the connection to game theory has been developed in analyses of a wide range of pragmatic phenomena (Van Rooij, 2004; Benz, Jäger, & Van Rooij, 2005; Jäger, 2008a). In particular, a game-theoretic approach to pragmatic phenomena can be used to derive models of back-and-forth reasoning (in which listeners reason about speakers and vice versa) (Franke, 2009; Jäger, 2008b). These models are deeply related to RSA, and in many ways can be viewed as cousins, often differing from RSA primarily in their use of more deterministic representations. We return to the connections between game-theoretic models and RSA in the formal presentation of the model, below.

Finally, in a more recent strand of work that has tied together some of this previous research, computational linguists interested in parsing human language production in grounded contexts have explored and expanded on ideas about connections between language production and rational action. Golland, Liang, and Klein (2010) introduced a simple one-step reasoning model in which speakers reason about listeners to produce good referring expressions. And Vogel, Potts, and Jurafsky (2013) attempted to derive the Gricean maxims directly from multi-agent decision theory. These pieces of work share deep formal similarities – in both their use of probabilistic models and nested reasoning processes – with the framework we next turn to, the Rational Speech Act framework.

²This viewpoint has been taken to be in conflict with the predictions of RSA models in some cases (e.g., Gatt, van Gompel, van Deemter, & Krahmer, 2013). Here we focus here on RSA as a model of listeners only, deferring discussion of the complex issue of how well RSA fits speakers' behavior. We refer listeners to Goodman and Frank (in press) for further discussion of the theoretical issues, and Monroe and Potts (2015) and Graf, Hawkins, and Goodman (2016) for two recent treatments of speaker behavior using RSA.

Rational Speech Act (RSA) models

RSA models have their antecedents in work on Bayesian cognitive modeling (for review, see Tenenbaum et al., 2011). These models are predicated on the idea that the probability calculus can be used to represent uncertainty in an agent’s internal representations of the world within structured hypothesis spaces. RSA models particularly build on Bayesian models of social reasoning (e.g., Baker, Saxe, & Tenenbaum, 2009), which use the tools of expected utility theory to capture agents’ reasoning about others’ actions.

Frank and Goodman (2012) and Goodman and Stuhlmüller (2013) introduced RSA models as a class of probabilistic models that capture the nested, back-and-forth of Gricean reasoning.³ In particular, listeners are modeled as reasoning about the decision-making – in terms of linguistic costs and benefits – of speakers. These speakers themselves reason about listeners. This recursion (listener - speaker - listener - speaker - ...) grounds out in the literal semantics of the utterance itself – described in the model as a “literal listener.” In the simplest form of the model, a listener hears an utterance and considers the goals that would have led a speaker to have uttered it. This (model-internal) speaker in turn considers the literal interpretation of that message compared with other alternatives that would be congruent with the speaker’s goal.

Both initial reports on RSA instantiated a model of this type and then tested the congruence of the predictions of the model with the judgments of human observers. Frank and Goodman (2012) used one-shot reference games in which participants were asked to reason about what speakers and listeners might say or understand given single-word utterances produced in response to displays of geometric objects. And Goodman and Stuhlmüller (2013) used vignettes to ask about number and quantifier implicatures (specific pragmatic inferences within lexical scales) in the presence of varying knowledge on

³Two other related formalisms, developed concurrently, provide alternative frameworks for modeling this kind of reasoning (Franke, 2009; Russell, 2012). In particular, the iterated-best response framework of Franke (2009) provides a complementary way to look at this kind of reasoning, building on the game-theoretic work described above.

the part of speakers (e.g., what does “some of the envelopes are empty” mean, given that the speaker hasn’t opened all of them?). Both reports found a tight congruence between participants’ aggregate judgments and model predictions.

Since these initial reports, a wide range of work has made use of RSA models and their variants to address a range of topics from individual differences in language comprehension (Franke & Degen, 2016) and the pragmatics of spatial language (Carstensen, Kon, & Regier, 2014) to non-literal language (Kao, Wu, Bergen, & Goodman, 2014), ambiguity (Lassiter & Goodman, 2015), and compositionality (Potts, Lassiter, Levy, & Frank, 2015; Bergen, Levy, & Goodman, 2016). While these extensions are important, there has been less work that provides a systematic assessment of the basic RSA model and the reference-game paradigms that have often been used to test them (but cf. Qing & Franke, 2015, for a replication of Frank & Goodman, 2012, and an exploration of some model assumptions). Here we attempt to address this gap, using “ad-hoc” implicatures (Stiller, Goodman, & Frank, 2015), pragmatic inferences that do not rely on specific lexical scales such as quantifiers.

The current work

The current work has two interlocking goals. The first goal is empirical. Via a sequence of large-scale, web-based experiments, we provide strong evidence that one-shot reference games provide a flexible and precise tool for studying quantitative patterns of human behavior. The second goal is formal. We assess the fit of basic RSA models to our data, exploring a range of model variants and parameter choices. This framework captures many of the patterns of performance we observe in our experiments, although some ambiguities – relating to parameters of the model – remain. Taken together, this body of work suggests that the RSA model provides a powerful set of tools for studying human pragmatic reasoning in quantitative detail.

Our experiments are listed in Table 1. To aid comprehension, they are broken into

three sequences. The first of these (“Preliminaries”) explores our experimental paradigm – one-shot reference games – providing experimental evidence that minor design choices do not account for the pattern of data we observe. The second section (“Priors”) tests the role of prior expectations in reference game behavior. The third (“Levels”) tests the level of recursion that speakers reason to and provides further data on the relationship between prior and posterior measurements.

Formal Framework

The RSA model has been described previously (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013; Goodman & Frank, in press). Our goal here is to provide a somewhat more systematic formal presentation, allowing us to highlight choice-points and discuss parallels and differences with other formal models in this space.⁴

Definitions

A reference game under our definition is a game in which a speaker S and listener L collaborate in a context C to identify a particular object in the context, known as the speaker’s intended referent $r_S \in C$. The game has two parts. First the speaker chooses a message w based on vocabulary V ; messages can be selected from a simple list or an open-ended grammar. Next the listener guesses a referent $r_L \in C$ after hearing w . The game is won if $r_S = r_L$. Although these game-theoretic foundations of RSA allow for definition of more complex payoff schemes (as in e.g., Franke, 2009), here we assume the simplest possible payoff, which is equal for S and L , positive if the game is won, and zero otherwise.

We assume that the world consists of a set of objects, some subset of which are present in C . There is a distribution σ over the objects $o_1 \dots o_n \in C$. This distribution picks out those objects in the context which are more or less likely to be talked about, either because of their intrinsic perceptual or conceptual noteworthiness, or because of some prior

⁴For another quite comprehensive presentation and comparison with extensions, see Bergen et al. (2016).

history between speaker and listener (e.g., one object having been talked about previously). Although in previous work we described it as a “contextual salience” distribution, we now refer to this distribution as the “prior.” “Salience” carries an implication that differences are perceptual rather than arising from other higher-level reasoning processes; we do not intend that interpretation and adopt a more neutral label here. The precise interpretation of this distribution and the extent to which it is assumed to be mutually-known by speaker and listener are both complex issues that we will return to throughout the paper.

We also assume a semantics for the set of possible messages w that a speaker can send. These messages can be linguistically complex phrases or utterances (see Potts et al., 2015 for an example of this approach), but here we examine only single words. Similarly, though these words could be learned pragmatically (Frank & Goodman, 2014; Smith, Goodman, & Frank, 2013), here we assume the semantics is known to both interlocutors. Here we treat messages as a Kronecker δ function that applies to objects and returns TRUE (1) if the message is true of that object, FALSE (0) otherwise. When applied to a context, a message specifies a uniform distribution over the objects for which it is true.⁵

The model

We define the RSA model in terms of sub-agents, listeners L and speakers S . Their goals are to transmit information efficiently. L s reason about what word S s would have said to describe a particular referent; S s reason about what interpretation L s would give to a particular message. Each agent uses Bayesian inference to reason about the other’s likely actions.

So, beginning our recursion, we can define:

⁵In other work, we have used continuous message semantics to capture how typical a particular exemplar is as the referent of a message (Graf et al., 2016).

$$P_{L_n}(r_S | w, C) \propto P_{S_n}(w | r_S, C) \quad (1)$$

$$P_{S_n}(w | r, C) \propto P_{L_{n-1}}(r_S | w, C) \quad (2)$$

where $n > 0$. The two agents are defined recursively and so individual probabilities are undefined unless the recursion ends. Agent L_0 , the “literal listener,” grounds the recursion by choosing an interpretation uniformly between available referents that are consistent with w :

$$P_{L_0}(r_S | w, C) \propto \delta_w(r_S). \quad (3)$$

In the case of a single cycle of recursion, L_1 reasons about S_1 , who in turn reasons about L_0 . But L_1 and S_1 are not special. L_n is simply a listener whose inferences reflect n full cycles of listener-speaker recursion. Clearly the level of recursion n is an important parameter, and it is one we discuss below in considerable depth.

We then define a top-level agent for the purposes of decision-making, L :

$$P_L(r_S | w, C) \propto P_{L_n}(r_S | w, C)P(r_S) \quad (4)$$

where $P(r_S)$ is the prior distribution over referents σ .⁶ In principle, we could also add a prior on words $P(w)$, but for simplicity here we assume that $P(w) \propto 1$ and do not discuss it further. In this statement of RSA we introduce the prior only at the top level of recursion but we test this assumption below.

⁶Note that L could be an agent at L_0 , in which case choices between alternatives would be a uniform multiplied by the prior.

Utility-theoretic definition

An alternative definition of our model can be posed in terms of rational action based on speakers’ utilities, following the game-theoretic definition above. Consider a speaker who weighs a message w in terms of its informational benefit $I(w)$ and cost $D(w)$:

$$U(w; r_S; D) = I(w; C) - D(w) \quad (5)$$

The speaker should then choose rationally between messages as a function of a choice rule, e.g.:

$$P_S(w \mid r_S, C) \propto e^{-\alpha U(w; r_S; C)} \quad (6)$$

where α is a “greed” parameter. If $\alpha = 0$, choices are random; as $\alpha \rightarrow \infty$, she always chooses the option with greater utility.

We define informational benefit with respect to the recursively lower listener’s surprisal, such that the speaker’s utility decreases as the listener (at the level they are assuming, namely $n - 1$) has more uncertainty: $I(w; C) = -(-\log(L_{n-1}))$. If we assume that cost is constant, $\alpha = 1$, and $n = 1$, then the log in surprisal and the exponent in the choice rule cancel (as in the derivation in Frank & Goodman, 2012), yielding $P_{S_1} \propto |w|^{-1}$.

Matrix definition

This probabilistic notation emphasizes individual interpretation probabilities; here we briefly describe how these models can be applied to entire signaling game matrices. We omit the prior here for simplicity of presentation but it can easily be added. We write a signaling game as a context matrix M with indices w and o corresponding to words and objects. If $M_{w,o} = 1$, then that word can be named with that object ($w(o) = \text{TRUE}$).

Now define our listener and speaker agents as two functions over signaling game matrices. (As above, for simplicity we consider the case where there are no costs and

$\alpha = 1$). The first function, $L(x)$, is the listener function; for a particular object-word pair, it returns the probability of that object, given the word:

$$L(x_{w,o}) = \frac{x_{w,o}}{\sum_{o' \in C} x_{w,o'}}. \quad (7)$$

Applying L to each element of matrix M results in the literal listener L_0 .

$S(x)$ is then the speaker function, which captures the intuition that the speaker chooses the best word to describe a particular object (from those that are available in the vocabulary). $S(x)$ takes an object-word pair from the matrix and returns the probability of speaking this word, normalized over the other possible words:

$$S(x_{w,o}) = \frac{x_{w,o}}{\sum_{w' \in V(C)} x_{w',o}}. \quad (8)$$

So now, if we successively apply these functions to matrix M , we can define agents with the same recursive depth as in the description above. For example, $L(S(L(M)))$ is equivalent to the L_1 agent above. In essence, repeated application of the L and S functions is equivalent to iterated renormalization alternating within rows and columns.

Worked example

We work through an example computation on the stimulus shown in Figure 1, which we refer to as our “simple” implicature stimulus because it is the simplest feature arrangement that results in an implicature (defined here as an inference that is not predicted by L_0). The stimulus is described by a matrix,

$$M_{simple} = \begin{matrix} & \begin{matrix} \emptyset & G & HG \end{matrix} \\ \begin{matrix} \text{hat} \\ \text{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \end{matrix}. \quad (9)$$

So then, normalizing within rows,

$$L(M) = \begin{matrix} & \emptyset & G & HG \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & .5 & .5 \end{bmatrix} \end{matrix}. \quad (10)$$

Note that $L(M)$ is equivalent to L_0 . Then, normalizing within columns,

$$S(L(M)) = \begin{matrix} & \emptyset & G & HG \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & .66 \\ 0 & 1 & .33 \end{bmatrix} \end{matrix} \quad (11)$$

and then, normalizing within rows again,

$$L(S(L(M))) = \begin{matrix} & \emptyset & G & HG \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & .75 & .25 \end{bmatrix} \end{matrix}. \quad (12)$$

Thus, at L_1 , the message “glasses” refers to G, the face with only glasses.

This example brings up one complexity not treated above, namely that the $\emptyset$ column has no non-zero entries, so technically would be undefined under Equation 8. Translated into practical terms, the question is how a speaker should refer to $\emptyset$, the foil entry, when there are no negative terms in the language. Should she simply speak at random and hope for the best? In the current paper, we largely duck this complexity, focusing on listener behavior and treating these entries as having value 0 rather than being undefined. Experiments 9–10 take a first step towards assessing listeners’ empirical behavior in these kinds of cases, however.

Relationship to other models

As reviewed above, a wide range of models in many different traditions have explored similar ideas to those in RSA. First, in a seminal early paper, Rosenberg and Cohen (1964) studied simple communication games much like the ones we study here. They model

speakers and listeners as non-recursive agents who nevertheless make symmetric choices using a Luce (1963) choice rule. Second, in a study of referring-expression production, Golland et al. (2010) describe a speaker-centric model in which an agent like $S(M)$ is the “reflex speaker” (producing true messages with equal probability) and $S(L(M))$ is called the “reasoned speaker.”

Finally, Jaeger (2010) describes an iterated best response (IBR) model in which each agent takes the highest probability alternative, e.g.,

$$L(x_{w,o}) \propto \begin{cases} P(r_S) & \text{if } x_{w,o} \max_{o' \in C} x_{w,o'} \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

This max operation can increase the frequency of a number of undesirable results such as ties or rows/columns with no non-zero elements (a problem that also occurs for our model, as discussed above). The IBR model thus includes some procedures for dealing with these cases. Otherwise, RSA and IBR are quite similar in spirit.

General Experimental Methods

This section describes the methods used in the experiments reported below. Our goal was to create a general method for measuring pragmatic inferences in simple reference games. Our taking-off point is two previous studies in which simple feature-based displays allowed measurement of pragmatic reasoning in grounded contexts (Frank & Goodman, 2012; Stiller et al., 2015). In what follows we describe some of the general features of these displays and the experiments that use them.⁷

⁷A general template for our experiments is available at <https://github.com/langcog/pragmods-expts> and all experimental data, analyses, and code for simulations are available at <https://github.com/langcog/pragmods>.

Participants

One challenge of pragmatic communication experiments is that repeated communication within a signaling game is very likely to influence responding (e.g., as in Brennan & Clark, 1996, where participants’ descriptions decreased in length as the experiment progressed). For this reason, we choose here to adopt a between-subjects approach, where each participant answers exactly one relevant question.

We adopted a general standard of 50 independent participants per cell, based on the tradeoff between cost and the desired precision of measurements. With samples of $N=50$ participants making binary decisions, we could assume 95% confidence intervals (CIs) with width .24 at their widest; doubling the sample size to $N=100$ per cell would only reduce max CI width to .18. While not every experiment reflects this precise standard due to idiosyncrasies of recruitment and exclusion, we have tried to maintain approximately this standard throughout.

All of the experiments described here were run as web experiments. We used Amazon’s Mechanical Turk crowdsourcing service to recruit participants between Fall 2013 and Spring of 2015. In general, each experiment was run as an independent human intelligence task (HIT); a few were posted as multiple HITs for convenience or due to experimenter error. Within experiments, we remove duplicated workers, so that data in each experiment represent a set of unique judgments by distinct participants.⁸

In addition to excluding duplicated participants, we also excluded participants who failed manipulation checks (see below). Table 1 gives details of participants for each experiment.

⁸We have not excluded workers who participated in multiple experiments. Although the MTurk population is not huge, it changes relatively quickly, with approximately half of all workers turning over six months (Stewart et al., 2015). Because of the wide temporal intervals between experiments and the randomization of experimental items, it is unlikely that participants would have seen similar items.

Stimuli

We created a set of six base “domains” that had features that could easily be composed independently: faces (pictured in Figure 1), boats, pizzas, sundaes, snowmen, and christmas trees. We created slightly different versions of each base so that they would appear to be unique (e.g., by varying the proportions and color tone of the face). We then were able to add features to each of these programmatically (with most experiments using two features, but some using three). Faces were supplemented with hats, glasses, and mustaches; boats had motors, sails, and cabins; pizzas had olives, peppers, and mushrooms; sundaes had whipped cream, chocolate, and cherries on top; snowmen had mittens, hats, and scarves; and christmas trees had lights, ornaments, and stars on top. All stimuli were judged to be appropriate for the U.S. participants we planned to recruit using Mechanical Turk.

Procedure

Participants viewed the experiment within a browser window. The first screen of the experiment presented a basic description of the paradigm and asked for informed consent. The second screen of the experiment presented the interlocutor, Bob (a cartoon picture of a man), and noted that he liked to do activities with the base item (e.g., visit friends or sail boats). In experiments with familiarization stimuli (e.g., Experiment 5), familiarization images were presented on this screen.

The third screen was the key screen: It presented the experimental stimulus (a set of base items augmented with features, representing the signaling game of interest) at the top of the screen. Below each stimulus a letter was displayed. In all experiments except for Experiment 2, two manipulation check questions were asked directly below the display using text entry boxes (e.g., “how many boats have cabins?”). Below this was the prompt, e.g.: “Bob can only say one word to communicate with you and he says: *glasses*. Click below on the option that represents the friend that you think Bob is talking about.” Below

this was a set of buttons labeled A, B and C for the participant to indicate which of the three options they thought the speaker was referring to. A final screen asked participants the following: “Who did you meet in this survey?”, “What was this survey about?” and “Any other comments for us?” Optionally, participants could also provide their age and gender, but finishing the experiment was not contingent on doing this.⁹

Methodological Preliminaries

Our first set of experiments, Experiments 1–3, establish the simple reference games in Figure 1 as a viable method, testing the replicability of judgments in this paradigm and examining particular experimental decisions.

Experiment 1 explored the use of different dependent variables. We considered a betting measure, in which participants were instructed to allocate \$100 across the three targets, as in Frank and Goodman (2012); a 7-point Likert scale, as in Goodman and Stuhlmüller (2013); and a simple three-alternative forced choice measure.¹⁰ Figure 2 shows results from each. With all three dependent variables, the pragmatic target was chosen more than the logical target, but this result was strongest for the forced choice, with 79% of participants choosing the pragmatic target.

In contrast, both the Likert and betting measures afforded the possibility of “hedging” by allocating equal bets or ratings to the pragmatic and logical targets. Conservative participants took advantage of this possibility very frequently, betting equal amounts on the logical and pragmatic targets 43% of the time and using equal Likert ratings 46% of the time. The forced-choice dependent variable did not allow this kind of hedging; thus, we adopted it in our further experiments. In practice, these hedges might signal opportunities for clarification questions, a common response to ambiguity in more

⁹In the first set of experiments, we manipulated a number of these experimental choices, including the dependent variable, the use of the manipulation check, and the linguistic framing of Bob’s utterance. Our description here reflects the defaults used in the majority of the experiments we report.

¹⁰As this experiment was one of our first and we were still actively developing items, we used only four of our six stimulus items: faces, boats, snowmen, and sundaes.

naturalistic communication (e.g., Clark & Schaefer, 1989).

Experiment 2 was designed to ensure that the presence of the manipulation check – which asked participants to count how many of each feature was present – did not induce a task demand that changed the magnitude of pragmatic responding. We saw 83% pragmatic responding in the manipulation check condition and 77% pragmatic responding in a matched condition with no manipulation check. Because of the high power of this experiment ($N_{include} = 513$), the difference between conditions was statistically significant ($\chi^2(2) = 6.97$, $p = .03$). Nevertheless, the modest (6%) difference between conditions suggests that the manipulation check does not *create* the pragmatic effect, though it may heighten participants’ attentions to the different distributions of the two features, leading to slightly more pragmatic responding. We continue to use the manipulation check as an exclusion criterion in further experiments.

Experiment 3 was designed to test the particular linguistic framing we used. In particular, we contrasted the presentation of the target word (e.g., “glasses”) as part of a phrase, “Bob says: ‘My friend has glasses.’” and a presentation of the target word with the framing that the speaker can only say a single word (described above in the General Methods). The “one word” framing produced 80% pragmatic responses, while the “My friend has glasses” framing produced 82% pragmatic responses. These frames did not differ significantly from one another in the distribution of responses they produced ($\chi^2(2) = 2.54$, $p = .28$). Except where noted below, we adopt the one-word framing.

In sum, these findings suggest that aspects of the experimental presentation and dependent variable have some minor effects on the magnitude of the pragmatic effect we observed. Nevertheless, the existence of the effect and its magnitude were quite robust to these variations.

Measuring and manipulating prior expectations

Experiments 4–7 examine the role of prior expectations – broadly construed – in reference resolution through the kinds of reference games introduced above. These experiments serve two goals. First, they create a set of quantitative measurements of prior expectations combined with listener inferences; these measurements can be used for model comparison and assessment, a topic that we return to in the section on model fits below. Second, they allow causal inferences about the role of the prior in participants’ pragmatic inferences. While Frank and Goodman (2012) *measured* prior expectations and showed that they improved model fit, that work did not *manipulate* prior expectations, so it was in principle possible that improvements in model fit did not result from a causal role played by prior expectations. We address this issue in the studies below. We begin by comparing methods for eliciting prior expectations (Experiment 4). We then investigate three ways of manipulating prior expectations: via base-rates (Experiment 5), via linguistic valence (Experiment 6), and via perceptual salience (Experiment 7).

It is not immediately clear what the best way is to measure prior expectations for purposes of communication. One particular conceptual issue seems important, though: does the term $P(r_S)$ refer to a prior over *meaning*, that is, over speakers’ particular communicative intentions? Or does it instead represent a prior over *world states*, for example, the things the speaker would do, think about, or request? Although the distinction is subtle in most of our studies, the question of priors over meanings vs. states is a fairly deep philosophical issue.¹¹ In addition, if these two kinds of priors differ substantially, then our estimates of $P(r_S)$ might lead to incorrect conclusions.

Frank and Goodman (2012) introduced a method in which priors were measured identically to listener judgements, with the exception that participants in the prior measurement condition were not given any linguistic information. They were told that the

¹¹See also Franke (2009), who notes that “we should not conflate the [listener]’s possible conjectures about actualities with his preferred interpretations of expressions” (p.129).

speaker could utter one word to signal their target referent, but that they (the listener) hadn't heard it. This method yielded reliable measurements, but presupposed that the prior was over meanings, rather than world states.

In Experiment 4, we compared prior elicitation methods. We presented participants with the same canonical display we had used in the earlier experiments, but asked them one of two prior questions. The meaning prior method was based on Frank and Goodman (2012), where the participant was told that "Bob can say one word" but the word was presented as "mumblemumble (you didn't hear what he said)." The world state prior method consisted of simply asking, e.g., "Which friend do you think Bob will visit next?" A simple inferential test failed to show a significant difference in the distributions between the two questions ($\chi^2(2) = 3.45$, $p = .18$, Figure 3). While this failure is of course not decisive, in both conditions, participants showed a tendency to choose the logical target (e.g., face with both hat and glasses) most, the foil at intermediate rates, and the pragmatic target relatively less. This experiment does not resolve the underlying philosophical issue of what kinds of expectations the prior actually tracks. Given the limited empirical differences, however, we put aside this issue for the time being. Since the choice does not appear critical and the focus of our model is communicative behavior, we adopt the meaning-based "mumble" prior elicitation method for subsequent experiments.

In Experiment 5, we attempted to manipulate participants' prior expectations about reference via a manipulation of the base rate of the interlocutor's actions.¹² The goal of this experiment was to provide evidence that prior expectations about actions influence inferences in the signaling games we were studying – licensing a causal inference about the role of priors.

Before presenting the inference question in this experiment, we exposed participants to evidence about the habits of the interlocutor (Bob). The familiarization screen in this

¹²This experiment is conceptually similar to a base-rate manipulation reported in Stiller, Goodman, and Frank (2011). However, the manipulation in that experiment used a different display, one more similar to those used in Experiment 8.

experiment stated (for example, in the case of the face stimulus) that Bob visited a friend every week; it then invited the participant to click on nine boxes to uncover the nine friends that Bob had visited most recently. Across four between-subjects base rate conditions, we manipulated how many of these boxes contained the pragmatic target (e.g., the face with glasses and no hat). One of the boxes always contained the foil (no glasses or hat) and another contained the logical target (hat and glasses). Then the other boxes contained either the pragmatic or the logical target. The four conditions were created such that 1/9, 3/9, 5/9, or 7/9 boxes showed the pragmatic target.

Figure 4 shows the results of this manipulation.¹³ The base rate had a strong effect on the prior elicitation, but it also boosted levels of pragmatic inferences. To quantify these effects, we used a logistic regression model (no hierarchical model was warranted because judgments were from independent participants). This model showed highly significant effects of base rate ($\beta = .27$, $p = .004$) and condition (with prior as the dummy-coded predictor, $\beta = -.57$, $p < .0001$), as well as a marginal interaction of the two ($\beta = .29$, $p = .06$). Critically, the main effect of base rate remained reliable in a similar logistic regression for just the inference trials ($\beta = .28$, $p = .0001$), indicating a highly reliable (and graded) effect of the base-rate manipulation on the strength of the pragmatic inference.

Our next experiment, Experiment 6, investigated the effects of linguistic manipulations on participants' priors and their resulting inferences. The goal of this experiment was to test whether, like base rate information, other, more subtle, manipulations of prior expectations would be integrated into the (putatively pragmatic) judgments we observed. In this experiment, we framed judgments as being about determining Bob's *favorite* or *least favorite* item. This experiment was conducted chronologically earlier in the sequence and had four rather than six experimental items:

¹³This experiment had unusually high rates of exclusion due to participants' confusion about whether they should report familiarization frequencies for the manipulation check (instead of counts in the display). We maintain our strict inclusion criterion here but note that the results are largely unchanged if we include these participants.

faces, boats, pizzas, and snowmen.

In the prior condition, the data overall supported the conclusion that Bob’s favorite item would likely have more features, while his least favorite would likely have fewer. The logical option, which had the most features, was chosen as “favorite” 70% of the time, while the distractor option, which had zero features, was chosen as “least favorite” 77% of the time.

In the inference condition, however, the linguistic manipulation had different results across items, as seen in Figure 5. Put simply, if your least favorite boat has a sail, it’s likely to have no other features; in contrast, if your favorite boat has a sail, it’s likely to have other features as well. The same was true of snowmen as well. But for pizzas, if your favorite *or* your least favorite has mushrooms, it is likely to have little else on it. Presumably participants reasoned that, if you like mushrooms, you want lots of them, and if you don’t like them, you would be most upset if they were alone, unmasked by any other topping. Confirming this impression of the data, a logistic regression predicting target responses as a function of item and condition showed a significant interaction between linguistic condition and the (dummy coded) pizza item ($\beta = -2.51$, $p = .008$). Thus, this experiment demonstrates that participants’ prior expectations about the importance of particular features for a particular domain interacted with their willingness to make a pragmatic implicature.

Our final experiment in this sequence, Experiment 7, attempted to manipulate prior expectations about reference by changing the salience of the different referents. The goal of this experiment was to use a more obvious, perceptual manipulation of expectations (compared with Experiment 6). In particular, we presented all but one of the images in grayscale, varying which of the images (pragmatic target, logical target, or foil) was shown in color.¹⁴ This manipulation had a strong effect on pragmatic inferences. In particular, when the logical target was shown in color, inferences to the pragmatic target dropped

¹⁴To ensure that participants still made the inference more generally, we included a condition where none of the referents was shown in color, notated by “none.”

significantly (Figure 6; in a logistic regression, $\beta = -1.42$, $p = .003$), presumably because the color either drew attention to the logical target or signaled to participants that their attention *should be* drawn to the logical target.

In sum, Experiments 5–7 together show strong evidence that many information sources can affect prior expectations about reference. In particular, these experiments license two inferences. First, manipulations that affect prior expectations about what a speaker will talk about (and, potentially, do) have a causal effect on inferences in signaling games. Second, these manipulations were reflected in the global rate of inference in a graded fashion that appeared quite sensitive to their type and magnitude. We return to these data below and use the RSA model to test the hypothesis that variation in inference levels can be accounted for by variation in prior expectations.

Levels of recursion and limits on human reasoning

Experiments 8–10 probe the limits on the kind of inferences participants can make in our reference game paradigm. The general goal of these experiments was to test whether the inference we saw so robustly above in the “simple” signaling game would generalize to others with more varied arrangements of features. In all of these experiments, we also conducted parallel experiments measuring prior expectations using the methods described above, but we largely omit discussion of these data until we use them for model fitting below.

In our first experiment, Experiment 8, we measured patterns of inference in more complex reference games. Following Stiller et al. (2011) and Degen and Franke (2012), we created a signaling game that required a somewhat deeper level of reasoning than our previous experiments. In particular, we compared our standard reference game (shown again for reference in Figure 7, bottom left) alongside a more complex game (Figure 7, bottom right):

$$M_{complex} = \begin{matrix} & \begin{matrix} M & GM & HG \end{matrix} \\ \begin{matrix} hat \\ glasses \\ mustache \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 0 \end{bmatrix} \end{matrix} \quad (14)$$

Different reference games require different depths of recursion to achieve a solution. We define a solution for a listener as a state where each message can be mapped to a particular referent (e.g., has a unique maximum value), and for a speaker as a state where each referent has a corresponding message (similarly, has a unique max). (Of course, some matrices do not have such solutions at all, but if a matrix is solved for a speaker, it is solved for all deeper listeners and vice versa.) For example, the “simple” matrix, which we have studied in the previous experiments, reaches a solution state with only half a level of recursion: As shown in the worked example above, a model that implements $S(L(M_{simple}))$ will show an implicature. In contrast, in the “complex” matrix, $L(S(L(M_{complex})))$ is required for one of the potential implicatures (that “glasses” refers to GM).

In Experiment 8, we used the stimulus in Figure 7, left bottom, to measure inferences about the previous pragmatic target (“glasses” interpreted as G, as in M_{simple}) and the target with the unique feature (“hat” means HG). We also used the stimulus in Figure 7, right bottom, to measure the more complex pragmatic inference (the target for this inference was M for “mustache”). The main contribution of this experiment, however, was our test of the $L(S(L(M)))$ inference (“glasses” implicating GM). This inference depends on having reasoned that, while both of the targets have another possible descriptor, for one of those targets, the second descriptor is unique, whereas for the other the second descriptor is also non-unique. As the preceding description shows, such reasoning is non-trivial.

Figure 7, top, shows the findings. Judgments in the simpler inference remained very close to 75% (the level observed across many of our other experiments), regardless of display. Literal meaning judgments were at ceiling. In contrast, more complex inferences

were at chance (53%). The prior for this target was lower than the prior for the logical distractor image, however (23% vs 39%). Thus, it may be the case that participants make this inference but it was partially or fully cancelled by an expectation against it (see Figure 1 of Frank & Goodman, 2012, for a potential example of this kind of cancellation, and see the next section for further discussion).

We considered conducting experiments with more complex signaling games that required deeper levels of recursion to achieve a solution for some messages. Brute-force searches of the space of possible matrices returns examples like Figure 8, with

$$M = \begin{matrix} & \text{GM} & \text{HG} & \text{H} & \text{GM} \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \\ \textit{mustache} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} \end{matrix} \quad (15)$$

and

$$L(S(L(S(L(M)))))) = \begin{matrix} & \text{GM} & \text{HG} & \text{H} & \text{GM} \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \\ \textit{mustache} \end{matrix} & \begin{bmatrix} 0.00 & 0.00 & 0.50 & 0.50 \\ 0.00 & 0.37 & 0.31 & 0.31 \\ 0.65 & 0.35 & 0.00 & 0.00 \end{bmatrix} \end{matrix}. \quad (16)$$

The key result for this matrix is that, at L_2 but not L_1 , “glasses” implicates HG. But the numerical difference is quite tiny (leading us to question whether we could measure such an effect), and the intuition felt even weaker than in the “complex” game (of which this matrix is a variant). Thus, we chose not to run these experiments.

Experiments 9–10 further considered the scope of the inferences that participants make. In particular, these experiments had as a goal asking whether participants would ever violate the literal semantics of the messages (especially in cases where it would otherwise be impossible to refer to a particular object). We began by examining displays

like the one shown in Figure 9, top:

$$M = \begin{matrix} & & \text{GM} & \text{HM} & \text{HM} \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \\ \textit{mustache} \end{matrix} & \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix} \end{matrix} \quad (17)$$

These displays had a uniform feature and another feature that was shared by two of the three objects. Intuitively, naming the unique feature (“glasses”) should reference the object with it (GM). But the predictions for the other features were perhaps less clear. Should the uniform feature also resolve to the unique object (e.g., “mustache” means GM)? And should “hat” resolve ambiguously to the duplicated objects (HM), or to the non-duplicated object (GM), violating the literal semantics of the message?

Referencing the duplicated feature (“hat”) led most participants to choose one of the duplicated referents (HG), suggesting that they did not violate the literal semantics (Figure 9). The message that applied to all three objects (“mustache”) led to a standard implicature (as in the “simple” matrix). Like Experiment 8, this experiment provided some evidence for the robustness of the basic implicature phenomenon.

Experiment 10 further considered whether participants would reach communicative equilibria that were at odds with those implied by the literal semantics of the words the speaker used, in cases where the matrix was systematically and frustratingly ambiguous. To address this question, we used a display like the one in Figure 10, top:

$$M = \begin{matrix} & & \text{HG} & \text{HM} & \text{GM} \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \\ \textit{mustache} \end{matrix} & \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \end{matrix} \quad (18)$$

Although every feature was ambiguous, a potential equilibrium would be the use of a

particular feature to describe the target *without* that feature. For example, here “hat” could refer to the face without a hat; “glasses” to the face without glasses; etc. But by and large, participants did not converge on this equilibrium and instead were content to choose between the two ambiguous referents (10, bottom, word condition), with only 17% of participants choosing to violate the literal semantics.

A second condition examined whether participants would be more or less likely to abandon the signal if it were an unconventional signal rather than a word with conventional semantics. In this condition, we recolored the features of each object (e.g., the hat and glasses) so that each had its own unique color. Bob then communicated by showing a color patch corresponding to the matching feature. Contra our initial prediction, participants stuck *more* closely to the match between color patches and particular features, choosing the “odd man out” at very low rates (8%). This condition suggests that it may be a feature of reasoning in these signaling games (rather than of words per se) that keeps participants from abandoning the literal semantics of the particular signals the communicator uses.

These results contrast with recent findings from Misyak, Noguchi, and Chater (in press), who studied whether a particular signal’s meaning could reverse in a signaling game where participants communicated about hidden objects. They found that signals could be used flexibly to indicate the locations of either rewards or dangers, depending on the particular constraints and affordances of the communicative situation. Messages in their study were “tokens” that were placed on top of locations in a virtual world, however, perhaps suggesting a more flexible semantics (e.g., “something [good or bad] is in here”) than the words and color patches we used in our studies. In addition, though they saw conventions appear even without repeated communication with a single partner, all participants in their studies built up some practice with the constraints of the virtual world of the game. In contrast, our participants got only a single opportunity to reason about a wholly-new set of constraints. More work will be needed to determine what aspects of the communicative task allow or prevent the emergence of reversible semantic conventions.

In sum, these experiments suggest the basic pragmatic inferences of the type studied throughout this paper were robust across a variety of displays (for example in the new matrix types shown in Experiments 8–9). But we saw very little evidence that, in a one-shot game of the sort studied here, participants would abandon the literal semantics of the words that were used (cf. Misyak et al., in press). And the evidence for inferences at deeper levels of recursion was limited at best (cf. Franke & Degen, 2016). Whether these are limitations of our paradigm or of human pragmatic reasoning more generally will require further study.

Model fits and model comparison

In the experiments reported above, we saw strong evidence that participants are engaging in reasoning that goes beyond the simple semantic interpretation of messages. Instead, they appear to be making inferences about reference in context that have the character of pragmatic reasoning – considering inferential alternatives that a speaker could have said. In this section, we test the ability of the RSA model to describe human behavior in these experiments.

Model comparison methods

For each experiment in the second and third sequences (Experiments 5–10), we compared predictions of the RSA model to participant data. As in Frank and Goodman (2012), for each experiment and experimental condition, we compared the proportion of participants choosing each different target to the model’s predictions about that target (e.g., proportion choosing target H, HG, etc.). This procedure yielded between 4 and 10 datapoints per experiment.

We used an implementation of the RSA model provided by a new package for the R statistical computing platform: `rrrsa`. The package is available at <http://github.com/benpeloquin7/rrrsa> and installation instructions and

comprehensive documentation are also available at that link.¹⁵

When we propagate empirical prior information into the model, we compute confidence intervals by simulating model fit using a non-parametric bootstrap resampling procedure in which the prior data are resampled with replacement and then the model is rerun 500 times using these resampled priors. Error bars then reflect percentile-based confidence intervals computed from these simulations.

Throughout, we assess model fit to data using a simple Pearson correlation statistic. Although a good fit alone cannot be probative (Roberts & Pashler, 2000), the RSA model as stated here and earlier contains a small number of parameters: depth of recursion and α (the greed parameter in the choice rule). We believe the use of correlations is relatively unproblematic since (as we show) changes in these parameters do not result in any radically different patterns of model predictions. Nevertheless, we do caution that our comparisons between models are exploratory (rather than being fully-preplanned confirmatory tests) and serve to elucidate differences in fit, rather than being a formal basis for quantitative model comparison. In particular, even if there are differences in fit between models, such a result doesn't guarantee better performance across stimulus conditions or parameters other than those evaluated here.

RSA models achieve globally good fit to experimental data

In our first simulation, we used the exact version of the RSA model given in Frank and Goodman (2012), with a recursive depth of 1 and $\alpha = 1$. Results of this simulation are given in Figure 11. The RSA model was correlated with human judgments ($r_{all} = .86$), and this correlation persisted even when removing those points predicted to be 0 or 1

($r_{no\ 0/1} = .68$).¹⁶

¹⁵The `rrrsa` package is optimized for analysis of experimental data like those presented here; other, more flexible variants of RSA models are available for exploration and extension at <http://forestdb.org/models/scalar-implicature.html> and <https://github.com/cgpotts/pypragmods>.

¹⁶For Experiment 7 – and to a lesser extent, for Experiment 6 – we saw a number of conditions in which model predictions had very high variance. This variance arises from variability in the prior on particular

Nevertheless, there were a number of experimental conditions for which this parameter setting did not yield good predictions. In Experiments 5–8, there appears to be a trend in which the model overpredicts for options that are judged to be lower probability, and underpredicts for those options judged more likely by participants. (This trend manifests itself as a cluster of predictions that are too close to the vertical midline of $p_{model} = .5$). In order to explore this trend further, we next varied the recursive depth in the model. Intuitively, further recursion should strengthen pragmatic inferences.

Deeper recursion improves model performance

Table 2 shows the results of simulations at depths 0 - 5. These results show two main findings. First, and critically, a single level of recursion dramatically improves performance over and above a flat (literal listener) model. This trend is especially pronounced for non-binary model predictions. Indeed, our experiments were designed to probe cases where a literal listener would make equal predictions for different objects but a pragmatic agent would differentiate between them. Thus, this result provides a (perhaps obvious, but still important) confirmation of our basic hypothesis, namely, that participants reason pragmatically in our experiments.

Second, however, the simulations showed an unexpected finding: deeper levels of recursion correlated more highly with human judgments. In particular, even four levels of recursion improved fit numerically compared with three levels (especially for non-0/1 predicted values). Figure 12 shows simulation results at depths 0 – 2 (the middle row is identical to Figure 11); we do not plot depths 3 – 5 because they are visually almost indistinguishable from depth 2 even though there are minor numerical differences in their correlation with participant data.

Recall, however, that we set the α parameter to 1 in these simulations. α controls the “greediness” of the model’s choice rule, with higher values indicating a greater willingness

objects (since they are chosen very infrequently in the prior condition, they show high numerical instability).

to choose higher-probability options. In the next set of simulations, we let α vary.

The α parameter trades off with recursive depth, limiting model-depth identifiability

Figure 13 shows the range of correlations with human data achieved at different levels of α , across recursive depths; Table 3 gives the highest correlation and corresponding α level. Two findings become apparent. First, at the highest level, although numerical correlations do vary across α levels, with $\alpha > 1$ the range of correlations is quite restricted across all depths. Thus, the performance of the model is relatively insensitive to this parameter.

Nevertheless, a close examination of Table 3 shows that values at depth 2 and especially depth 1 are noticeably higher than in Table 2 where $\alpha = 1$. In fact, by setting $\alpha = 3.1$, a model with depth 1 still achieves just about the same level of fit to data as a model with depth 4 or 5. This finding suggests that α and depth trade off against one another in most of the experiments that we conducted.

To understand this result, consider the standard matrix we used in most of our experiments (reproduced from above):

$$M = \begin{matrix} & & \emptyset & G & HG \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \end{matrix} \quad (19)$$

In this case, recursive depth 1 and $\alpha = 1$ produces:

$$L(S(L(M))) = \begin{matrix} & & \emptyset & G & HG \\ \begin{matrix} \textit{hat} \\ \textit{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & .75 & .25 \end{bmatrix} \end{matrix} \quad (20)$$

But note that recursive depth 2 produces:

$$L(S(L(S(L(M)))))) = \begin{matrix} & \emptyset & G & HG \\ \begin{matrix} \text{hat} \\ \text{glasses} \end{matrix} & \begin{bmatrix} 0 & 0 & 1 \\ 0 & .83 & .17 \end{bmatrix} \end{matrix} \quad (21)$$

which is identical to the same result at recursive depth 1 with $\alpha = 2$. But that is not always true. In principle, if a particular matrix is completely ambiguous at a particular recursive depth, then changing α will not affect performance. For example, the symmetry for the two identical elements in the matrix in Experiment 9 cannot be broken by changing α . (Another, more trivial, example of this observation would be changing levels of α for a literal listener; clearly this manipulation would have no effect since by definition the literal listener assigns probability evenly to all alternatives.)

Priors do not improve global fit and also trade off with α and recursive depth

The final issue we examine is the contribution of empirically-determined priors to overall model fit. We began with a model that used empirical prior measurements because Frank and Goodman (2012) proposed that they were necessary. Despite this previous work, we found that global fit of a model with $\alpha = 1$ and no prior was equivalent to the performance of the fitted- α model *with* priors. Correlations are shown in Table 4; numerically, results are almost identical with those shown in Table 3. Levels of recursion from 0–2 are shown visually in Figure 14.¹⁷

It is surprising that the use of the empirical prior fails to improve model fit. Yet despite this lack of global improvement, we note that there is unambiguous evidence from Experiments 5–7 that manipulations of participant expectations about the target object

¹⁷We also experimented with incorporating the prior at every level of recursion, rather than only at the top level of recursion. This model variant corresponds to a scenario in which the prior is mutually-known to both communicators (and they reason pragmatically about that mutual knowledge). In contrast the case we presented above corresponds to a scenario where the prior represents information that only the listener has access to (e.g., about perceptual salience or personal expectations about the situation). We found that the “mutual knowledge” variant performed worse than either the no-prior or the “private knowledge” model (the one presented in the main text) in all cases, with the fitted- α correlations lower by .02-.05 depending on the level of recursion.

lead to altered pragmatic judgments. In Experiment 5 in particular, we saw that manipulations of the base rate of a target object led to quantitative changes in both prior measurements and rates of pragmatic inference.

We discuss two observations about this unexpected finding: that there are tradeoffs between the prior, α , and depth of recursion and that there may be measurement issues that complicate model comparisons. These observations are based on an analysis of the individual correlations, separated by experiment (Table 5).

First, intuitively and as illustrated above, both increases in α and increases in recursive depth function to “pull apart” different targets, increasing the highest probability option and decreasing the lowest probability option. A comparison of Figures 12 and 14 suggests that including empirical prior measurements tends to exert the opposite pressure. In fact, this observation reveals an interesting generalization about our displays: our pragmatic targets tend to be *less* salient than our logical targets. For example, in the canonical display we use, the face with a hat is less salient than the display with the hat and glasses. Including information of this sort about salience via the empirical prior then works *against* the pragmatic inference, even as both increasing α and recursive depth strengthen the inference. Another way to think about this observation is that, in cases that the pragmatic target is more salient, it will reliably be picked by L_0 – such cases may not even appear pragmatic since no reasoning is necessary.

Second, adding empirical priors to the model complicates quantitative comparison with no-prior models substantially. Although adding the empirically-measured prior does not add any free variables to the model, it does add substantial measurement-related noise (consider the error bars shown in Figure 12 but not Figure 14). Thus, a failure to see increases in overall fit to data may be the result of a tradeoff between better fit in some cases due to the prior better accounting for the structure of the data *and* worse fit due to the prior measurements introducing random noise. In addition, errors due to random noise are systematic: They bias extreme values towards the center. This systematic bias would

almost certainly reduce global fit to data, although quantifying its magnitude might be difficult.

General Discussion

The flexibility of comprehension in context is one of the most important and striking features of human language. Yet in contrast to theories of syntax and semantics (e.g., Chomsky, 1965; Jackendoff, 2002), the most influential pragmatic theories have typically been quite informal in nature (Grice, 1975; Sperber & Wilson, 1986; Clark, 1996; Levinson, 2000). This gap has been addressed by recent efforts to develop probabilistic models of pragmatic reasoning, inspired by developments in game theory (e.g., Benz et al., 2005) and social cognition (Baker et al., 2009).

In the current work, we gave an extended presentation of one particular class of probabilistic pragmatics model: the rational speech act, or RSA, model (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013). Although a variety of models in this family have been presented previously, we provided a more extended description and comparison to other related models (along with software for using basic RSA models for simulations). We then conducted a series of large-scale, single-judgment pragmatic inference experiments to validate and explore one-shot reference games as a method for studying pragmatic inference quantitatively.

These experiments support three main conclusions. First, language games of the type we studied here are a flexible and robust tool for probing pragmatic judgments (despite their schematic nature). Second, judgments in these games were systematically affected by a variety of manipulations of the linguistic, perceptual, and statistical structure of the game; these variables appeared to influence participants' pragmatic inferences in a graded way. Third, we found that participants typically began their reasoning with the literal semantics of the message that they were interpreting and converged to an interpretation consistent with a low (but definitely non-zero) number of recursive iterations.

Our simulations using the RSA model largely confirmed these conclusions. We found that the basic RSA model provided relatively good fit to our corpus of experimental data. Increasing either the α (greed) parameter *or* the level of recursive depth in the pragmatic computation each led to increases in fit, but these two parameters traded off against one another. For this reason, we were not able to make a conclusive statement about the level of recursion used by comprehenders in our study (for more evidence on this question, see Franke & Degen, 2016). Surprisingly, we did not find evidence for a global boost in fit by including empirical prior information, although we did see that some experiments were fit better by models that included a prior. We caution against too broad an interpretation of this last finding, however. Priors have played a very important role in applications of RSA models to other phenomena (e.g., Kao et al., 2014), therefore there may be complexities in the estimation and use of the prior either in our specific paradigm or in grounded object reference games generally.

Despite our attempts to make a comprehensive study of the RSA model, our work here has many limitations that should form a basis for future investigations. First of all, our experiments were focused on the judgments of single Mechanical Turk participants in one-off games. This methodology limits interpretation in a number of important ways. Although preliminary evidence suggests that such judgments are similar for other convenience populations (e.g., college undergraduates in the U.S.; Frank, Goodman, & Tenenbaum, 2009), it is an open question how similar rates of pragmatic inference would be across cultures and contexts. In addition, because our design elicited exactly one judgment per participant, we have no information about how rates of pragmatic inference vary across individuals. Individual differences in inference and their connections with other aspects of language processing are a rich topic for future study, and reference games like ours may provide a possible tool for executing well-controlled studies in this area.

More generally, however, our modeling framework and experiments probe only a small part of the immense space of contextual inferences in language comprehension. We

restricted our scope to single interactions, purposefully blocking collaborative inferences about reference across a discourse (Clark, 1996). Similarly, although details of the broader context, discourse history, and physical setting can deeply influence the details of the pragmatic inference (Sperber & Wilson, 1986), we focused on aspects of the situation that we could either manipulate systematically or hold constant. The strength of this empirical strategy is that it yields large amounts of quantitative data that can be compared with quantitative models. The weaknesses, however, are apparent as well. First and foremost among these is that our work eliminates much (but hopefully not all) of the contextual details that create pragmatic inferences in the first place.

Nevertheless, we are hopeful that the models we describe here can be extended to a broader range of phenomena. RSA-type models have been used to capture judgments across a wide range of stimuli (e.g., Franke & Degen, 2016; Carstensen et al., 2014) and phenomena (e.g., Potts et al., 2015; Kao et al., 2014; Bergen et al., 2016); for review, see Goodman and Frank (in press). By treating RSA as a base for the formal exploration of pragmatic phenomena, we hope to build towards a more comprehensive quantitative theory of language comprehension in context.

References

- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, 113, 329–349.
- Bar-Hillel, Y. (1971). Out of the pragmatic wastebasket. *Linguistic inquiry*, 2, 401–407.
- Benz, A., Jäger, G., & Van Rooij, R. (Eds.). (2005). *Game theory and pragmatics*. Palgrave Macmillan.
- Bergen, L., Levy, R., & Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics & pragmatics*.
- Breheny, R., Ferguson, H. J., & Katsos, N. (2013). Taking the epistemic step: Toward a model of on-line access to conversational implicatures. *Cognition*, 126, 423–440.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1482.
- Carstensen, A., Kon, E., & Regier, T. (2014). Testing a rational account of pragmatic reasoning: The case of spatial language. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 2009–2013). Austin, TX: Cognitive Science Society.
- Chemla, E., & Spector, B. (2011). Experimental evidence for embedded scalar implicatures. *Journal of semantics*, 28, 359–400.
- Chomsky, N. (1965). *Syntactic structures*. Walter de Gruyter.
- Clark, H. H. (1996). *Using language*. Cambridge, UK: Cambridge University Press.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. *Perspectives on socially shared cognition*, 13, 127–149.
- Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive science*, 13, 259–294.
- Dale, R., & Reiter, E. (1995). Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive science*, 19, 233–263.
- Dale, R., & Reiter, E. (1996). The role of the Gricean maxims in the generation of

- referring expressions. *arXiv preprint cmp-lg/9604006*.
- Degen, J., & Franke, M. (2012). Optimal reasoning about referential expressions. *Proceedings of SemDIAL*.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*, 998.
- Frank, M. C., & Goodman, N. D. (2014). Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology*, *75*, 80–96. Retrieved from <http://langcog.stanford.edu/papers/FG-cogpsych2014.pdf>
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, *50*, 578–585.
- Franke, M. (2009). *Signal to act: Game theory in pragmatics*. Amsterdam, NL: Institute for Logic, Language and Computation, University of Amsterdam.
- Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual-vs. population-level probabilistic modeling. *PloS one*, *11*, e0154854.
- Frederking, R. E. (1996). Grice's maxims: Do the right thing. In *Proceedings of the Computational Implicature Workshop at the AAAI-96 Spring Symposium Series* (pp. 1–6). Stanford, CA.
- Gatt, A., van Gompel, R. P., van Deemter, K., & Krahmer, E. (2013). Are we bayesian referring expression generators. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Golland, D., Liang, P., & Klein, D. (2010). A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 conference on empirical methods in natural language processing* (pp. 410–419).
- Goodman, N. D., & Frank, M. C. (in press). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling

- language understanding as social cognition. *Topics in Cognitive Science*, 5, 173–184.
- Graf, C., Hawkins, R. X. D., & Goodman, N. D. (2016). Animal, dog, or dalmatian? Level of abstraction in nominal referring expressions. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- Grice, H. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics* (Vol. 3, pp. 41–58). New York, NY: Academic Press.
- Jackendoff, R. S. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford, UK: Oxford University Press.
- Jaeger, G. (2010). Game-theoretical pragmatics. In *Handbook of logic and language* (p. 467). Elsevier.
- Jäger, G. (2008a). Applications of game theory in linguistics. *Language and linguistics compass*, 2, 406–421.
- Jäger, G. (2008b). *Game theory in semantics and pragmatics*. (Manuscript, University of Bielefeld)
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111, 12002–12007.
- Keefe, R., & Smith, P. (Eds.). (1997). *Vagueness: A reader*. Cambridge, MA: MIT Press.
- Lassiter, D., & Goodman, N. D. (2015). Adjectival vagueness in a Bayesian model of interpretation. *Synthese*, 1–36.
- Levinson, S. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.
- Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology*. New York: Wiley.
- Misyak, J. B., Noguchi, T., & Chater, N. (in press). Instantaneous conventions: The emergence of flexible communicative signals. *Psychological Science*.
- Monroe, W., & Potts, C. (2015). Learning in the rational speech acts model. *arXiv preprint*.

- Parikh, P. (1991). Communication and strategic inference. *Linguistics and Philosophy*, 14, 473–514.
- Potts, C., Lassiter, D., Levy, R., & Frank, M. C. (2015). Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. *Journal of Semantics*, 10.1093/jos/ffv012.
- Qing, C., & Franke, M. (2015). Variations on a bayesian theme: Comparing bayesian models of referential reasoning. In *Bayesian natural language semantics and pragmatics* (pp. 201–220). Springer.
- Reiter, E., Dale, R., & Feng, Z. (2000). *Building natural language generation systems* (Vol. 33). MIT Press.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological review*, 107, 358.
- Rosenberg, S., & Cohen, B. (1964). Speakers' and listeners' processes in a word-communication task. *Science*, 145, 1201.
- Russell, B. (2012). *Probabilistic reasoning and the computation of scalar implicatures* Unpublished doctoral dissertation. Citeseer.
- Smith, N. J., Goodman, N., & Frank, M. C. (2013). Learning and using language via recursive pragmatic reasoning about other agents. In *Advances in neural information processing systems*. Retrieved from <http://langcog.stanford.edu/papers/SGF-NIPS2013.pdf>
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford, UK: Blackwell Publishers.
- Stewart, N., Ungemach, C., Harris, A. J., Bartels, D. M., Newell, B. R., Paolacci, G., & Chandler, J. (2015). The average laboratory samples a population of 7,300 amazon mechanical turk workers. *Judgment and Decision Making*, 10, 479.
- Stiller, A. J., Goodman, N. D., & Frank, M. C. (2011). Ad-hoc scalar implicature in adults and children. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science*

- Society*. Austin, TX: Cognitive Science Society.
- Stiller, A. J., Goodman, N. D., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development*, 11, 176–190.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331, 1279–1285.
- Van Deemter, K., Gatt, A., Sluis, I. v. d., & Power, R. (2012). Generation of referring expressions: Assessing the incremental algorithm. *Cognitive science*, 36, 799–836.
- Van Deemter, K., Gatt, A., Van Gompel, R. P., & Krahmer, E. (2012). Toward a computational psycholinguistics of reference production. *Topics in cognitive science*, 4, 166–183.
- Van Rooy, R. (2004). Signalling games select Horn strategies. *Linguistics and philosophy*, 27, 493–527.
- Viethen, J., & Dale, R. (2006). Algorithms for generating referring expressions: Do they do what people do? In *Proceedings of the fourth international natural language generation conference* (pp. 63–70).
- Vogel, A., Emilsson, A. G., Frank, M. C., Jurafsky, D., & Potts, C. (2014). Learning to reason pragmatically with cognitive limitations. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society* (pp. 3055–3060). Austin, TX: Cognitive Science Society.
- Vogel, A., Potts, C., & Jurafsky, D. (2013). Implicatures and nested beliefs in approximate decentralized-POMDPs. In *Proceedings of the 2013 Annual Conference of the Association for Computational Linguistics* (pp. 74–80). Stroudsburg, PA: Association for Computational Linguistics.
- Von Neumann, J., & Morgenstern, O. (2007). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Wasow, T., Perfors, A., & Beaver, D. (2005). The puzzle of ambiguity. In O. Orgun & P. Sells (Eds.), *Morphology and the web of grammar: Essays in memory of Steven G.*

Lapointe (pp. 265–282). Stanford, CA: CSLI Publications.

Table 1

Outline of the experiments reported here.

Expt. #	Sequence	Experiment	N_{total}	$N_{include}$	Summary
1	Prelims	Dependent variable	689	554	Forced choice yields best DV
2		Manipulation check	580	513	Modest increase in inference for MC
3		Linguistic frame	100	89	No difference
4	Priors	Elicitation frame	200	175	No difference
5		Base rate	800	488	Base rate affects inference
6		Linguistic framing	550	502	Linguistic valence affects inference
7	Levels	Perceptual salience	300	267	Perceptual salience affects inference
8		Level of inference	416	362	Deeper recursion difficult
9		Twins	220	194	Few non-literal readings
10		Oddman	300	270	Few non-literal readings

Depth	r_{all}	$r_{no\ 0/1}$
0	0.74	0.14
1	0.86	0.68
2	0.93	0.86
3	0.95	0.90
4	0.95	0.91
5	0.95	0.91

Table 2

Correlations of model and data on both the full dataset and the subset of conditions for which the model did not predict a probability of either 0 or 1.

Depth	α_{best}	r_{all}	$r_{no\ 0/1}$
0	0.0	0.74	0.14
1	3.1	0.95	0.89
2	1.4	0.95	0.91
3	1.1	0.95	0.91
4	1.0	0.95	0.91
5	1.0	0.95	0.91

Table 3

Correlations of model and data (as above) with best-fitting α value at each depth.

Depth	r_{all}	$r_{no\ 0/1}$
0	0.74	0.14
1	0.94	0.90
2	0.95	0.90
3	0.95	0.90
4	0.94	0.90
5	0.94	0.90

Table 4

Correlations of no-prior model and data at varying recursive depths.

Depth	Model	E5	E6	E7	E8	E9	E10
1	No prior, $\alpha = 1$	0.96	0.89	0.95	0.96	0.88	0.95
1	Prior, α fit	0.99	0.93	0.89	0.97	0.75	0.98
1	Prior, $\alpha = 1$	0.87	0.87	0.67	0.95	1.00	0.98
2	No prior, $\alpha = 1$	0.99	0.87	0.95	0.97	0.80	0.95
2	Prior, α fit	0.99	0.93	0.89	0.98	0.84	0.98
2	Prior, $\alpha = 1$	0.95	0.96	0.81	0.98	0.99	0.98
3	No prior, $\alpha = 1$	0.99	0.86	0.95	0.97	0.74	0.95
3	Prior, α fit	0.99	0.93	0.89	0.98	0.89	0.98
3	Prior, $\alpha = 1$	0.98	0.95	0.86	0.98	0.95	0.98

Table 5

Correlations of model and data for three model variants and varying recursive depths, broken down by experiment. Results for depths 4 and 5 are similar to depth 3. α values are fit globally rather than individually for each experiment.

Figure 1. An example stimulus item showing our canonical inferential context. The prompt in this case would be “glasses.” We refer to the face with glasses but not a hat (G) as the “pragmatic target.” The “logical target” is the face with a hat and glasses (HG), and the “foil” is the face with nothing ($\emptyset$). In all experiments reported here, position and identity of all features is randomized across displays; the pragmatic target was not always the face with glasses nor was it always positioned in the middle.

Figure 2. Data from Experiment 1. Each panel shows results from a different dependent variable (normalized for simplicity into the interval $[0,1]$). The three columns show results for the foil (e.g., $\emptyset$), logical (e.g., HG), and pragmatic (G) targets. Error bars show 95% confidence intervals, computed by non-parametric bootstrap using the percentile method.

Figure 3. Data from Experiment 4. Colors indicate different methods for eliciting prior expectations.

Figure 4. Data from Experiment 5 from prior elicitation and inference conditions. Horizontal axis shows the proportion of familiarization trials that showed the pragmatic target.

Figure 5. Data from Experiment 6 from the inference conditions, broken down by item.

Figure 6. Data from Experiment 7 in the inference conditions.

Figure 7. Data from Experiment 8, plotted by display (left: simple, right: complex) and inference level (see text).

Figure 8. An example signaling game for which a solution is present for L_2 but not L_1 .

Figure 9. Data from Experiment 9, along with the display used for this experiment.

Figure 10. Data from Experiment 10, along with the display used for this experiment.

Figure 11. Human data plotted against RSA model predictions for a recursive depth of 1, and $\alpha = 1$. Each point shows proportion of choices for a single target object. Colors denote the pragmatic target (e.g., glasses), logically true object (e.g., hat and glasses), foil (e.g., no hat, no glasses). Vertical error bars are binomial confidence intervals; horizontal error bars are computed by resampling prior data and recomputing model predictions (see text). Black lines show the line of best fit.

Figure 12. Human data plotted against RSA model predictions for recursive depths 0 – 2 (shown as different rows of panels), and $\alpha = 1$. Plotting conventions are as above.

Figure 13. Correlation between full human dataset and RSA models at different levels of α . Colors show levels of recursive depth from 0–5, dashed lines mark the maximum correlation for each depth.

Figure 14. Human data plotted against RSA model predictions for recursive depths 0 – 2, and $\alpha = 1$. RSA model does not include estimates of the prior, hence there are no error bars on the horizontal axis. Plotting conventions are as above.