

# Variations on a Bayesian Theme: Comparing Bayesian Models of Referential Reasoning

Ciyang Qing and Michael Franke\*

Institute for Logic, Language and Computation  
University of Amsterdam, The Netherlands

**Abstract.** Recent developments in Bayesian experimental pragmatics have received much attention. The Rational Speech Act (RSA) model formalizes core concepts of traditional pragmatic theories quantitatively and makes predictions that fit empirical data nicely. In this paper, we analyze the RSA model and its relation to closely related game theoretic approaches, by spelling out its *belief*, *goal* and *action* components. We introduce some alternatives motivated from the game theoretic tradition and compare models incorporating these alternatives systematically to the original RSA model, using Bayesian model comparison, in terms of their ability to predict relevant empirical data. The result suggests that the RSA model could be adapted and extended to improve its predictive power, in particular by taking speaker preferences into account.

## 1 Introduction

Human language communication is efficient, in that people need not always say every detail in order to be understood. Rather, speakers often only say what is most relevant, and listeners can often seemingly effortlessly grasp the intended meaning beyond what is literally said. This ability to do pragmatic inference has been long studied and the *conversational implicature* theory by Grice [1] is one of the most prominent theories in the field. However, since it is hard to formalize the *Cooperative Principle* and the *Conversational Maxims*, Grice's theory is not precise enough to give empirically testable predictions, especially quantitative predictions.

The Bayesian *Rational Speech Act* (RSA) model attempts to address this issue by using information-theoretic notions together with the Bayesian inference framework [2,3,4,5]. It has been shown that the model could yield quantitative predictions that are highly correlated to actual human judgments [3,5].

Despite the RSA model's theoretical promise and empirical success, if we analyze the design of the model, we will find several choices that are not *prima facie* obvious and might thus need further consideration and justification. These choices have their alternatives in closely related game-theoretic approaches to pragmatics [6,7,8,9,10], which have similar but slightly different conceptual motivations. Hence it is important to systematically compare the original RSA model

---


\* Thanks to ...

with all these alternatives and test their predictions on empirical data, so as to gain a better understanding of the relation between these models. Doing so will help illuminate the theoretical commitments and empirical implications of each design choice, which will enhance our understanding of the nature of human pragmatic inference.

The outline of the paper is as follows. Section 2 introduces the referential communication games investigated by [3] and the original RSA model proposed there. Section 3 analyzes a few design choices in this model and introduces their alternatives in game theoretic pragmatics, emphasizing the distinction between *belief*, *goal* and *action* and unifying various models as different interpretations of these three components. Section 4 reports on the results of an experiment similar to the one of [3]. In Section 5, we compare the predictions of different models of our new data. Finally, we discuss the implications of the results of the model comparison in Section 6, concluding that the original RSA model can be improved on, in particular by taking speaker preferences into account, but we also point out that a more complex picture of probabilistic pragmatic reasoning about referential expressions awaits further exploration.

## 2 Referential Games and the RSA model

Referential games are confined interactive reasoning tasks used to study pragmatic inference in a controlled environment [3]. A referential game consists of a set of objects called a *context*. A context contains different shapes of different colors in various arrangements, such as shown in Figure 1(a). The speaker in the game is asked to refer to one of the objects (the *target*) by uttering an expression (typically a single word denoting a feature, i.e., color or shape, of the target) to the listener. The listener, who does not know which object is the target, needs to recover the target based on the speaker’s choice of expression.



**Fig. 1.** A simple referential game

Let us first try to play this game ourselves, to get a better sense of what it is about and why it is relevant to human pragmatic reasoning. For simplicity let us assume that the utterances that the speaker is allowed to use are from a *vocabulary* which is commonly known between the speaker and the listener. The

content of the vocabulary and the (truth-conditional/literal) meaning of each word are shown in Figure 1(b). Suppose we play as the speaker and the target is  $t_1$ , the green square. We can either use “Square” or “Green” to truthfully refer to it, but it seems more prudent to use “Square”, as there is only one square in the context and thus the listener can easily identify it, whereas there are two green objects, which makes “Green” ambiguous. In terms of the Gricean *Maxim of Quantity*, we should use “Square” because it is more informative than “Green” in the given context; the surplus informativity moreover seems relevant. Similarly, we should use “Blue” to refer to  $t_3$ , the blue circle. However, using “Blue” to refer to the blue circle might intuitively sound a little unnatural, as color terms are usually used as adjectives<sup>1</sup>, while we usually use nouns to refer to concrete objects. This inclination becomes more evident when we want to refer to  $t_2$ , the green circle. While “Green” and “Circle” are equally ambiguous in this case, we might nevertheless prefer the latter.

Now let us turn to play as the listener. If we hear “Square” or “Blue”, we easily know the intended referent as there is no ambiguity, but what if we hear “Circle” (or “Green”)? There are two circles in the context that we need to choose from. On the one hand, the blue circle, having the unique color in the context, seems to be perceptually dominant and thus easily captures our attention. On the other hand, from the previous analysis we know that if the blue circle were the intended referent, the speaker could have chosen “Blue” which is not ambiguous and thus more informative. Hence the listener needs to balance two sources of information, i.e., the (presumably subconscious) perceptual salience of different objects and the rational expectation of the likelihood of the speaker making the utterance for each object. The latter line of reasoning is crucial to the classic Gricean account of scalar implicatures, and the major challenges in terms of formal modeling are how to quantify notions such as informativeness and how different pieces of information should be integrated.

The RSA model addresses the problems by using information-theoretic concepts to measure informativeness and Bayesian inference to integrate different sources of information [2,3,4,5]. In order to measure the informativeness of an utterance, the RSA model starts with the literal listener who upon receiving an utterance  $m$  does a conditioning on its literal meaning:

$$\rho_0(t \mid m) = \mathcal{U}(t \mid \llbracket m \rrbracket), \quad (1)$$

where  $\mathcal{U}$  is a uniform distribution over all the possible referents. The informativeness of utterance  $m$  for the intended referent  $t$  can be measured as the negative Kullback-Leibler divergence of the induced literal listener’s belief  $\rho_0$  from the speaker’s own belief  $\delta_t$ :

$$\text{Info}(m, t) = -\text{KL}(\delta_t \parallel \rho_0) = - \sum_{t'} \delta_t(t') \log \left( \frac{\delta_t(t')}{\rho_0(t')} \right), \quad (2)$$

---

<sup>1</sup> They are used as nouns mostly to refer to the colors themselves.

where  $\delta_t$  is a delta distribution with all probability mass on target object  $t$ , as the speaker knows her intended referent:

$$\delta_t(t') = \begin{cases} 1 & \text{if } t' = t \\ 0 & \text{otherwise} \end{cases} . \quad (3)$$

The speaker acts sub-optimally by choosing the utterance that soft-maximizes her expected utility, which is defined as the informativeness of the utterance subtracted by its cost:

$$\sigma(m | t) \propto \exp(\lambda_S \cdot U(m, t)) = \exp(\lambda_S \cdot (\text{Info}(m, t) - \text{Cost}(m))), \quad (4)$$

where  $\lambda_S$  is a parameter measuring the speaker’s degree of rationality, i.e., to what extent the speaker sticks to the strict optimum. The cost term is used to encode preference in different utterances, be it about the utterances’ lengths or syntactic categories.

From (1)-(4) we obtain the speaker’s production rule:

$$\sigma(m | t) \propto \exp(\lambda_S \cdot (\log \mathcal{U}(t | \llbracket m \rrbracket) - \text{Cost}(m))) . \quad (5)$$

The pragmatic listener in the RSA model, upon receiving the utterance  $m$ , performs a Bayesian update on his prior belief  $\mathcal{S}(t)$  by using an estimate of the speaker’s behavior (5):

$$\rho(t | m) \propto \mathcal{S}(t) \cdot \sigma(m | t) . \quad (6)$$

Bayes’ rule naturally integrates the perceptual salience of each object, which is treated as the prior  $\mathcal{S}(t)$  and can be empirically measured, with listener’s expectation of the speaker being informative, which is incorporated as the likelihood, thus addressing the previously mentioned challenge of balancing different sources of information. Setting  $\lambda = 1$  and  $\text{Cost}(m) = 0$  for all  $m$ , [3] obtained a highly significant correlation between the model prediction and their experiment data on actual speaker and listener behavior gathered from referential games of varying complexity.

### 3 Alternatives in Design Choices

The RSA model has a theoretical advantage over traditional formal pragmatic theories in that it provides quantitative predictions, and it has been shown to fit relevant empirical data well. However, if we want to learn more about pragmatic reasoning about referential expressions, it will be worthwhile to examine RSA carefully to pin down its major components, to spell out its main design choices and their underlying assumptions, and to test their contribution to the predictive power of the model through statistical comparison with natural alternatives. In this section, we will therefore compare the RSA model to closely related game theoretic approaches, that likewise assume that speakers and listeners form cascading beliefs about mutual behavior and seek to optimize their behavior

based on these beliefs [6,7,8,9,10].<sup>2</sup> Doing so let's us describe a class of potential alternatives to RSA. We will then compare the predictive power of these given relevant empirical data in later sections.

There are three closely related notions, i.e., *belief*, *goal* and *action*, that play a crucial role in RSA as well as game theoretic models. The relation and distinction between them can be best illustrated by looking at (5), repeated here:

$$\sigma(m | t) \propto \exp(\lambda_S \cdot (\log \mathcal{U}(t | \llbracket m \rrbracket) - \text{Cost}(m))) .$$

The term  $\mathcal{U}(t | \llbracket m \rrbracket)$  is the speaker's *belief* about the listener's interpretation of the utterance. The expected utility  $U(m, t) = \log \mathcal{U}(t | \llbracket m \rrbracket) - \text{Cost}(m)$  describes the speaker's *goal* of communication, i.e., inducing a belief that is as close as possible to her own with minimal effort. The production rule  $\sigma(m | t)$  on the left-hand side specifies the speaker's *action*, as it defines the probability with which, according to RSA, speakers would choose a particular expression. The soft-max rule connects these notions, as a (sub-)rational agent's action depends on his belief and goal. While the RSA and game theoretic models all share this general architecture in their designs, they vary in the specific interpretations of the three ingredients, which reflect different views and emphasis on language and communication and give rise to a series of interesting and important empirical questions. Again, let us continue with the speaker model (5) to illustrate some points of divergence and formally define the corresponding alternatives.

First of all, although it is *prima facie* reasonable to hypothesize, as [3] do, that the empirically measured perceptual salience  $\mathcal{S}(\cdot)$  is common knowledge between speaker and listener, it does not actually affect RSA's production rule (5) at all. This makes it unclear in what sense the speaker can be said to know the listener's perceptual salience, since neither his beliefs nor actions depend on it. A natural variant, which is used in some game theoretic models, would replace the literal listener's uniform distribution  $\mathcal{U}(t)$  in (1) with the salience prior  $\mathcal{S}(t)$ . This leads to the alternative production rule:

$$\sigma_{\mathcal{S}}(m | t) \propto \exp(\lambda_S \cdot (\log \mathcal{S}(t | \llbracket m \rrbracket) - \text{Cost}(m))) . \quad (7)$$

Secondly, RSA measures the informativeness of an utterance  $m$ , which is a crucial part of the communicative goal, in terms of how close the induced belief of the literal listener is from the speaker's own belief. This means that the RSA model sees the goal of communication as conveying belief. While it is normally true that language does convey the belief of the speaker, it is questionable at least in this referential scenario whether letting the listener form a proper belief is the ultimate goal of the communication. After all, if the speaker wants to refer to something, it seems that in the end what matters is whether the listener actually picks out the intended referent successfully (e.g., when the speaker wants the listener to pass something). This view is inherent in game theoretic approaches where agents' beliefs are backed up by explicit formulations of their utilities. We might call this latter view *action-oriented*, in contrast to the *belief-oriented*

---

<sup>2</sup> See [11] for overview and comparison of game theoretic and Bayesian models.

view of communication which the RSA model adopts, as it interprets the goal of communication as invoking the intended action rather than forming an accurate belief.<sup>3</sup> Thus, according to this view, the informativeness of a message would be measured as the probability of the listener choosing the intended referent. Formally speaking, we have the action-oriented speaker model

$$\sigma_a(m | t) \propto \exp(\lambda_S \cdot (\rho_0(t | m) - \text{Cost}(m))) . \quad (8)$$

Hence we have four types of speaker models that differ in either the speaker’s belief about the literal listener, or the speaker’s goal of communication. We now introduce a uniform notation  $\sigma_{xy}$ ,  $x \in \{a, b\}$ ,  $y \in \{\mathcal{U}, \mathcal{S}\}$  for them:

$$\sigma_{ay}(m | t) \propto \exp(\lambda_S \cdot (y(t | m) - \text{Cost}(m))), \quad (9)$$

$$\sigma_{by}(m | t) \propto \exp(\lambda_S \cdot (\log y(t | m) - \text{Cost}(m))), \quad (10)$$

where  $\mathcal{U}$  is the uniform prior and  $\mathcal{S}$  is the salience prior. For example, in the original RSA model, the speaker does not take listener’s salience prior into account and he has a belief-oriented goal of communication. Thus it will be denoted as  $\sigma_{b\mathcal{U}}$ .

Finally, the original RSA speaker model of [3] has  $\text{Cost}(m) = 0$  for all  $m$ , which means that the potential speaker preference in different utterances is not taken into account (but not so in, e.g., [4]). As we previously pointed out, our intuition seems to suggest a preference for nouns over adjectives as referential expressions.<sup>4</sup> Since it is an empirical question whether, to what direction or to what extent such a preference exists, we leave open all the possibilities. Technically speaking, since only the difference in costs matters, we define the cost function using a constant  $c \in \mathbb{R}$

$$\text{Cost}(m) = c \text{ if } m \text{ is an adjective and } 0 \text{ otherwise} . \quad (11)$$

If  $c > 0$  it means there is a preference for nouns and if  $c < 0$  then the preference is for adjectives. No preference exists if  $c = 0$ .

Now we turn to the listener model (6), where the boundary between belief, goal and action becomes less clear:

$$\rho(t | m) \propto \mathcal{S}(t) \cdot \sigma(m | t) .$$

---

<sup>3</sup> Note that the two views are clearly intimately related, despite our emphasis on their differences here. An accurate belief is usually helpful for choosing a desired action. In fact, forming a belief can itself be seen as an action that the listener performs, which might well be all that the speaker cares about in many everyday situations, e.g., casual talks or responses to inquiries of information. Also, the relevant actions might vary greatly across situations which makes belief formation a practical alternative to serve as the goal. Nevertheless, they are still essentially different views in spite of the fact that they often coincide in practice. The referential communication scenario we investigate in this paper reflects the distinction.

<sup>4</sup> The rich tradition on referential expression generation [12] has also identified the need to take speaker preferences into account, and [13] criticize RSA on these grounds as well.

Let us start from the relatively clear part. As pointed out previously, the likelihood term  $\sigma(m | t)$  is the listener’s belief about how the speaker behaves. One natural question is to what extent the listener’s belief correlates with the actual production. We thus treat the speaker’s production term  $\sigma(m | t)$  in the listener’s model as a parameter, making the association with a belief about production of the listener model explicit:

$$\rho(\sigma_{xy})(t | m) \propto \mathcal{S}(t) \cdot \sigma_{xy}(m | t) . \quad (12)$$

Note that the speaker’s production rule has two parameters  $\lambda_S, c$  which are also included in the above specification of the listener model.

Next, even though our intuition suggests different objects have different perceptual salience and thus might affect our judgment, it is after all an empirical question whether it is relevant in the interpretation of referential expressions. It is in principle possible that the listener does not take into account the perceptual salience in his reasoning, which means he has a uniform prior over the referents:

$$\rho_{\mathcal{U}}(\sigma_{xy})(t | m) \propto \mathcal{U}(t) \cdot \sigma_{xy}(m | t) . \quad (13)$$

Finally, since the RSA model adopts a belief-oriented view on the goal of communication, its listener model only consists of the belief about the intended referent after hearing the utterance, obtained by a Bayes update. However, according to the action-oriented view, this is not quite the end of the story. The listener often needs to decide the exact referent of a referential expression. (Again, consider the case in which the listener is asked to pass something.) Hence in such cases the listener will choose an object by, essentially, (soft-)maximizing over his beliefs. Formally, the action-oriented listener model becomes:

$$\rho_{av}(\sigma_{xy})(t | m) \propto \exp(\lambda_L \cdot \rho_{bv}(\sigma_{xy})(t | m)), \quad (14)$$

where  $v \in \{\mathcal{U}, \mathcal{S}\}$ ,  $\lambda_L$  is the parameter measuring the listener’s degree of rationality, and  $\rho_{bv}$  is the belief-oriented model that does the Bayesian update:

$$\rho_{bv}(\sigma_{xy})(t | m) \propto v(t) \cdot \sigma_{xy}(m | t) . \quad (15)$$

For instance, the original RSA listener model is a belief-oriented one with the perceptual salience as prior, whose belief about the speaker is  $\sigma_{b\mathcal{U}}$ , hence it is denoted as  $\rho_{b\mathcal{S}}(\sigma_{b\mathcal{U}})$ .

## 4 The Experiment

In the previous section we analyzed the design of the original RSA model, spelled out the relation and distinction between belief, goal and action, and proposed a family of alternative models based on different interpretation of these notions. This gives rise to a series of interesting empirical questions regarding the underlying assumptions of the models. In order to gain insight into these questions by comparing the predictive power of different models, we conducted the following experiment to collect empirical data, which we will use to test the predictions of different models in the next section.

## Participants

We recruited 1032 US participants via Amazon’s Mechanical Turk, an online crowd-sourcing tool. Each participant played a one-shot referential game either as a speaker or as a listener and received a payment of 5 cents. If a participant contributed multiple trials, only the first was included. All participants passed the attention check described below.

## Materials and Procedure

Each participant saw a context consisting of three objects (Figure 2 as an example) and reported on the number of objects in the context as an attention check. Participants in the listener conditions were told to imagine someone was talking to them who used some utterance (depending on the specific condition tested) to refer to exactly one of the objects in the context. Then they were asked to choose what they thought to be the intended referent. Participants in the speaker conditions were told to imagine they were talking to someone and they wanted to refer to one of the objects (indicated by a small red arrow<sup>5</sup>). Then they were asked to choose between two words to refer to it.



**Fig. 2.** A sample context

Each context had a square and a circle sharing the same color, and another circle having a different color. Hence each object had two features: shape (square/circle) and color (green/blue) and each context had two ambiguous (shared by two objects) features. In the listener salience condition, participants were told that the person talking to them was using a foreign language they could not understand, and in the other two listener conditions, the utterances used were the ambiguous feature of shape and color, respectively. In the speaker conditions, the target could be any of the three objects: the one with the unique color, the one with the unique shape and the one with both features shared. The two words for the participants to choose between were the features of color and shape of the target object.

In order to minimize the effect of the confounding factors, we counterbalanced the positions and colors of the objects, as well as the orders of the candidate words in the speaker conditions. Thus the unique color in a context could be

---

<sup>5</sup> Participants were told that the arrow was only for illustration and thus the person they were talking to could not see it.



either blue (as in Fig. 2) or green and the order of the objects could be any of the permutations.<sup>6</sup>

## Design

There are four major design choices in our own experiment that differ from the original experiment in [3]. First of all, the original experiment used a *betting paradigm*, i.e., participants were asked to bet over a range of options and they were instructed that the amount of money that is bet on an option should reflect their confidence that the option is correct. Even though the betting paradigm has the merit of providing us with graded responses from each individual participant, the caveat is that it is unclear whether it measures beliefs or actions. This can lead to confusion when we are to fit the model predictions to the empirical data without knowing whether they are directly comparable. In addition, since the betting paradigm is more or less introspective in nature, it tends to be not very accurate. Thus we used a *forced choice design* instead, which clearly measures the action, does not rely on introspection and, technically speaking, provides for a straightforward likelihood function to be used in subsequent statistical model comparison. Secondly, since we decided to investigate the influence of the speaker’s preference as well as the listener’s perceptual salience, we focused on contexts equivalent to Figure 1(a), which require reasoning that is highly reminiscent of Gricean accounts of scalar implicature calculation, and examined different features separately. Thirdly, we only included the features of color and shape to ensure the vocabulary to be common knowledge, excluding the feature of texture. Finally, we did not use dotted lines surrounding an object as the way to indicate the target to the speaker, as [3] did, because that might unduly emphasize the feature of shape and thus be a confound in the experiment. Instead we used a small arrow pointing to the target.

Despite these differences, the remaining aspects in our experiment were almost identical to those of the original one, e.g. the phrasing of the instructions.

The results of the speaker and listener conditions of our experiment are as follows.




### 4.1 Speaker Conditions

There were 432 participants in the speaker conditions, 144 in each condition. The numbers of participants choosing each word in each condition are shown in Table 1. Note that when the target was the object with the unique shape (as in the first row of the table), the feature of shape (“Square”) should be the optimal utterance because the listener could uniquely identify it. Similarly when the target was the object with the unique color (the second row), the optimal utterance would be the feature of color (“Blue” in this case). When the target was the object with both features shared, both features should be equally ambiguous because of the context’s symmetric nature.

---

<sup>6</sup> Due to the scale of our preliminary experiment we did not counterbalance the shape of the objects, i.e., a context always had one square and two circles.

**Table 1.** Speaker Conditions




Target	“Green”	“Square”	“Blue”	“Circle”	Total
	9	135	–	–	144
	–	–	119	25	144
	63	–	–	81	144

From the data in Table 1, we can see that the speakers tend to choose the optimal feature more often when the target has the unique shape than when it has the unique color ( $\chi^2 = 8.5, p < .01$ ). Even though they seem to prefer the feature of shape when both of the target’s features are shared, the difference is not statistically significant from uniform random choice ( $\chi^2 = 2.25, p = .13$ ).

## 4.2 Listener Conditions

There were 600 participants in the listener conditions, 240 in the salience condition and 180 in each of the remaining two conditions. The numbers of participants choosing each object in each condition are shown in Table 2.

**Table 2.** Listener Conditions

				Total
Saliency	71	30	139	240
“Green”	65	115	0	180
“Circle”	1	62	117	180

The result of the salience condition will be used as the empirical estimation of contextual salience. For the other two conditions, the object with both features shared (the green circle in the above table, but note that we counterbalanced colors) is what the Gricean pragmatic account predicts to be the target. The result in Table 2 shows that when the message is “Green”, listeners prefer the green circle which is predicted by the Gricean pragmatics, while they tend to stick to the blue circle which is more perceptually salient when the message is “Circle”. The behavioral patterns in both conditions are significantly different from uniform random choice, and they significantly differ from each other in whether they conform to the predictions by Gricean pragmatic theory.

## 5 Model Comparison

We use a Bayesian approach to model comparison [14,15,16,17] to find out which alternative best explains the observed data. The models under investigation have unspecified parameters: the speaker’s degree of rationality  $\lambda_S$ , the cost of adjectives  $c$ , and, for those listener models that have an action-oriented communication goal, the listener’s degree of rationality  $\lambda_L$ . Since there is no principled

theory to determine the value of the parameters, we will rely mostly on relatively uninformed hyperpriors (so-called to distinguish them from the salience priors). Based on a specification of hyperpriors, we calculate the models' *evidences* and compare them by their *Bayes factors*. The evidence of a model  $M$  is the weighted average of the likelihood of observing the data under all parameter values:

$$\text{Ev}(M) = \int \text{Pr}(\theta) \cdot \text{Pr}(D \mid M, \theta) d\theta, \quad (16)$$

where  $\text{Pr}(\theta)$  is the hyperprior over parameter(-tuple)  $\theta$  associated with model  $M$  and  $\text{Pr}(D \mid M, \theta)$  is the likelihood of the observed data  $D$  given  $M$  and each concrete instantiation of the parameter(-tuple)  $\theta$ . The Bayes factor  $K_{M_2}^{M_1}$  is a comparative measure for the plausibility of model  $M_1$  over  $M_2$ , given their respective hyperpriors and the data in question:

$$K_{M_2}^{M_1} = \frac{\text{Ev}(M_1)}{\text{Ev}(M_2)}. \quad (17)$$

Model  $M_1$  makes the data more likely whenever  $K_{M_2}^{M_1} > 1$ , but normally only a Bayes factor  $K_{M_2}^{M_1} > 3$  (or sometimes  $K_{M_2}^{M_1} > 5$ ) is considered substantial. Values  $K_{M_2}^{M_1} > 10$  are considered strong evidence.

In a sense, comparison by Bayes factors is comparing models in a wider sense of the term: we actually compare pairs consisting of a model and its associated hyperprior. For clarity, we refer henceforth to a model-hyperprior pair as a Model.

*Speaker data.* First we look at the speaker models  $\sigma_{xy}$ ,  $x \in \{a, b\}$ ,  $y \in \{\mathcal{U}, \mathcal{S}\}$ . Each model has two parameters  $\lambda_S$  and  $c$ . We assume that they are independent of each other:

$$\text{Pr}(\lambda_S, c) = \text{Pr}(\lambda_S) \cdot \text{Pr}(c). \quad (18)$$

We are uncertain about the rationality of the speaker:

$$\text{Pr}(\lambda_S) = \mathcal{U}_{(0,11)}(\lambda_S), \quad (19)$$

which is a uniform distribution over  $(0, 11)$ . Excluding  $\lambda$ -values  $\geq 11$  serves practical purposes only, but is innocuous since the regions of non-negligible posterior likelihood of  $\lambda$  lie safely in the chosen interval for all models. Next, to allow for the possibility of speaker preferences (nouns over adjectives or vice versa), we consider two types of hyperpriors for costs  $c$ . The first hyperprior has  $\text{Pr}(c) = \delta(c)$ , the Dirac delta distribution, that assigns all probability mass to  $c = 0$ . This captures the assumption that there is no speaker preference. The second hyperprior is  $\mathcal{U}_{(-0.4, 0.4)}$ , which captures the notion that a preference exists, without commitment to either direction. We restrict our attention to the interval  $(-0.4, 0.4)$ , because we consider higher levels of cost implausible, given that utilities for successful communication live in  $[0, 1]$  and that we believe that strive for communicative success should outrank preference satisfaction in a rational model of communication. Taken together, there are four speaker models,


two hyperpriors for each, so that we compare eight Models with respect to their evidences.

Evidences of speaker Models were calculated by grid approximation. The results are shown as log-evidences in Table 3. (Notice that the Bayes factor between models is easily compared by taking the differences between log-evidences.) We can see that the data very strongly supports speaker models that do not take the listener’s perceptual salience into account. Also, it seems that action-oriented models are slightly better than their belief-oriented counterparts, even though the relevant Bayes factors are not substantial by common standards. Finally, our data makes each speaker Model that does allow for a speaker preference strongly more plausible than its counterpart that does not. A look at the posterior likelihood of  $c$  for each speaker model informs us that our data supports the belief in a speaker preference for nouns, see Figure 3.<sup>7</sup> In sum, our data supports the

**Table 3.** Log-evidences of speaker Models. The darker a cell’s background the higher the evidence.

support of $P(c)$	$\sigma_{bU}$	$\sigma_{aU}$	$\sigma_{bS}$	$\sigma_{aS}$
$[0, 0]$	-67.48	-67.81	-104.14	-157.75
$(-0.4, 0.4)$	-32.92	-33.15	-81.44	-139.27

**Fig. 3.** Posterior distributions over costs given the data for each speaker model.



view that the speaker does not take into account the perceptual salience of the listener, while having her own preference for shape terms over color terms.

<sup>7</sup> A technical aside: formally, action-based speaker models, as formulated in Section 3, assign a non-zero probability to the event that a speaker chooses a description which is not true of the target object (unlike belief-based models). But since subjects in our experiments could only select between true alternatives, the predictions of action-based speaker models were restricted to truthful choices in the reported model comparison.

*Listener data.* Each listener model has a speaker model nested inside as a belief of the listener about the speaker’s behavior. Section 3 introduced a total of 16 potentially relevant listener models, but we will focus on a selection only. For one, we restrict our attention to those listener models that are either entirely belief-based or entirely action-based. In other words, we exclude as notionally inconsistent models like  $\rho_a(\sigma_b)$  where the receiver part assumes an action-based goal structure and the speaker part a belief-based goal structure. For another, we assume that the listener’s model of the speaker is a reasonable one and therefore put to the side listener models that embed speaker models that are highly implausible, given the speaker data as discussed above. Effectively, this discards listener models that include speaker models that take the listener’s salience prior into account. These two principled restrictions leave us with four listener models to compare.<sup>8</sup>

Further variation comes from different relevant hyperpriors. Belief-based models of the listener have the same parameters as speaker models: the speaker’s rationality  $\lambda_S$  and the speaker’s preference costs  $c$ . Importantly though, hyperpriors over  $\lambda_S$  and  $c$ , although formally parallel to those for speaker models, have a different interpretation in listener models where they capture our prior beliefs about the listeners’ beliefs about the speakers’ likely rationality and preferences. This is true as well for action-based models, which also have an additional parameter  $\lambda_L$ . Hyperpriors for the latter encode our prior beliefs about the listener’s actual rationality.

We consider a variety of hyperpriors that differ in whether they take the speaker’s costs into account, whether the listener’s beliefs about the speaker’s rationality and preferences are uninformed (i.e., flat) or informed (i.e., given by the posterior likelihood of parameters given the actual speaker data) and, in action-based models, whether the listener’s level of rationality corresponds to his “informed” beliefs. The latter option effectively implements the assumption that there is a tight correlation between the speaker’s actual rationality, the listener’s actual rationality and the listener’s beliefs about the speaker’s rationality.

Concretely, we consider the following “flat” hyperpriors:

$$\begin{aligned}\Pr(\lambda_S, c) &= \mathcal{U}_{(0,11)}(\lambda_S) \cdot \Pr(c) \\ \Pr(\lambda_S, c, \lambda_L) &= \mathcal{U}_{(0,11)}(\lambda_S) \cdot \Pr(c) \cdot \mathcal{U}_{(0,11)}(\lambda_L),\end{aligned}$$

for belief-based and action-based models respectively, where  $\Pr(c)$  is the hyperprior for the speaker’s preference parameter  $c$ . When  $\Pr(c) = \delta(c)$  we have a hyperprior that does not take costs into account; otherwise we assume  $\Pr(c) = \mathcal{U}_{(-0.4,0.4)}(c)$ , as before.

---

<sup>8</sup> Similar to the remarks about speaker models in Footnote 7, we need to be mindful about the fact that action-based listener models assign non-zero probability to the choice of referents of which the given description is not true. Here, despite the fact that subjects in our experiments could in principle choose referents of which the given description is not true (this indeed happened once), we restricted the predictions of these models to only semantically compatible referent choices. We then discarded the one data point in which a subject choose a semantically non-compatible referent.

Hyperpriors that capture the idea that the listener’s beliefs are good guesses of speaker behavior are modeled as if informed by the data from the speaker experiments:

$$\begin{aligned} \Pr(\lambda_S, c) &= \Pr(\lambda_S, c \mid D_S, M_S) \\ \Pr(\lambda_S, c, \lambda_L) &= \Pr(\lambda_S, c \mid D_S, M_S) \cdot \mathcal{U}_{(0,11)}(\lambda_L) . \end{aligned}$$

Here,  $D_S$  is the data from the speaker experiments, and  $M_S$  the relevant speaker model. For a given listener model, we consider only the embedded speaker model as relevant. We call hyperpriors of the above form “informed” or, in the case of action-based models, “informed uncorrelated”. We distinguish the latter from “informed correlated” hyperpriors of the form:

$$\Pr(\lambda_S, c, \lambda_L) = \Pr(\lambda_S, c \mid D_S, M_S) \cdot \Pr(\lambda_L \mid D_S, M_S),$$

where the listener’s rationality parameter is distributed according to the relevant posterior. All of the three types of informed hyperpriors were tested in two varieties: whether the listener takes the sender’s preferences into account or not. If he does not, the posterior  $\Pr(\lambda_S, c \mid D_S, M_S)$  is derived from a speaker-hyperprior  $\Pr(c) = \delta(c)$ ; otherwise from  $\Pr(c) = \mathcal{U}_{(-0.4,0.4)}(c)$ .

Taken together, we consider two belief-based models, paired with four hyperpriors, and two action-based models, paired with six hyperpriors (see Table 4). Comparing these ten Models is meant to address the following general questions:

1. Is the goal structure assumed by participants in our task belief-based or action-based?
2. Does the listener take the estimated salience prior into account or not?
3. Is the listener’s belief about the speaker’s level of rationality “correct”, i.e., in line with the observed speaker data?
4. Is the listener’s level of rationality related to the speaker’s actual level of rationality?

Answers to these questions can be found by comparing the evidences of the Models listed in Table 4.

**Table 4.** Log-evidences for listener Models

prior type	costs	$\rho_{bU}(\sigma_{bU})$	$\rho_{bS}(\sigma_{bU})$	$\rho_{aU}(\sigma_{aU})$	$\rho_{aS}(\sigma_{aU})$
flat	no	-24.83	-11.66	-24.06	-10.62
flat	yes	-24.80	-9.84	-23.84	-10.81
informed	no	-59.87	-10.44	-25.42	-8.48
informed	yes	-68.43	-13.35	-24.21	-11.85
informed-correlated	no	NA	NA	-64.70	-6.80
informed-correlated	yes	NA	NA	-82.36	-29.01

The most striking contrast is that Models that do not take the salience prior into account fare much worse than those that do. The data makes a total rejection of Models that rely on uninformative salience priors strongly plausible.

Another eye-catching feature is that there is a clear winner in the list. The single best Model, which is substantially better than all the others, is action-oriented and takes perceptual salience into account. It has an informed-correlated hyperprior and assumes that the listener does not take the speaker’s preferences into account. This result is highly thought-provoking, but should be taken with a grain of salt. It might be that we merely got lucky by restricting the range of  $\lambda_S$  and  $\lambda_L$  just to a small region of reasonably high posterior likelihood for these parameters. Further experimentation should tell us whether informed-correlated hyperpriors are good predictors in general.

Most generally reliable are the results from flat hyperpriors. Here it is interesting to note that all of the models that take salience priors into account, with the exception of  $\rho_{bU}(\sigma_{bS})$  are equally plausible, by common standards of Bayes factor comparisons. This is despite the fact that action-based models have an additional parameter.

Action-based models are also better at accommodating the idea that the listener’s estimate of the speaker behavior is roughly correct. The belief-based models are substantially or strongly less plausible in this case. In other words, action-based models, possibly due to the added flexibility of another parameter, are better at explaining production and comprehension data in unison.

This latter point is of interest with respect to the findings of [3] that a single parameter choice  $\lambda = 1$  provided a good fit for both production and comprehension data. The RSA model  $\rho_{bU}(\sigma_{bS})$  with fixed parameter  $\lambda = 1$  is a very poor predictor of our data. (We focus here on comprehension, but the same applies to production data.) Assuming  $\lambda = 1$  as the null hypothesis of a nested-model comparison, we can use the Savage-Dickey method to compute a Bayes factor [18,19]. Let  $\mathcal{M}$  be the parameterized Model with model  $\rho_{bU}(\sigma_{bS})$  and a flat hyperprior, not taking costs into account. Given our data we should adjust our beliefs in parameter value  $\lambda = 1$  by the Bayes factor (computed approximately via MCMC sampling):

$$K = \frac{P(\lambda = 1 \mid D_L, \mathcal{M})}{P(\lambda = 1 \mid \mathcal{M})} \approx \frac{1.25e-9}{8.18e-2} = 1.52e-8$$

where  $D_L$  is the listener data. That means that our data provides very strong evidence that the null hypothesis  $\lambda = 1$  is incorrect.

## 6 Conclusion

Our experiment and model comparison suggest the following: (1) On the conceptual level, it is helpful to clarify the distinction between belief, goal and action, and the Bayesian framework provides us with a natural way to structure these components into a formal model. By examining each component systematically, we can explicitly spell out the underlying assumptions and formulate alternative hypotheses to be tested. (2) On the empirical level, we tested the RSA model with a family of variants motivated from a game theoretic perspective on communication, and in particular, we investigated the speaker’s preference as well

as the listener’s perceptual salience. Our data showed a more intricate picture of the roles that various factors play in pragmatic reasoning about referential expressions. We found evidence for a correlation between the speaker’s and the listener’s rationality. Either side appears to have his own biases, but appears to be negligent of the other’s. Also, an action-based view might better reflect the goal of communication, at least in the forced-choice experimental setting. Understanding the relation between the forced-choice and the betting paradigms would be an important next step to gain more insight into the differences between action- and belief-based goal structures.

Our work shows that careful conceptual analysis of the design choices for quantitative models can lead to a better understanding of the phenomenon and further improvement in the formal model’s predictive power. Of course, our restricted empirical data can only serve as a start and more data is needed to fuel further model comparison towards a more robust pragmatic theory of people’s probabilistic reasoning about referential expressions.

## References

1. Grice, H.: Logic and conversation. 1975 (1975) 41–58
2. Frank, M.C., Goodman, N.D., Lai, P., Tenenbaum, J.B.: Informative communication in word production and word learning. In: Proceedings of the 31st annual conference of the cognitive science society. (2009) 206–211
3. Frank, M.C., Goodman, N.D.: Predicting pragmatic reasoning in language games. *Science* **336** (2012) 998
4. Bergen, L., Goodman, N.D., Levy, R.: That’s what she (could have) said: How alternative utterances affect language use. In: Proceedings of the thirty-fourth annual conference of the cognitive science society. (2012)
5. Goodman, N.D., Stuhlmüller, A.: Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science* **5** (2013) 173–184
6. Rabin, M.: Communication between rational agents. *Journal of Economic Theory* **51** (1990) 144–170
7. Stalnaker, R.: Saying and meaning, cheap talk and credibility. In Benz, A., Jäger, G., van Rooij, R., eds.: *Game Theory and Pragmatics*. Palgrave MacMillan, Hampshire (2006) 83–100
8. Benz, A., van Rooij, R.: Optimal assertions and what they implicate. *Topoi* **26** (2007) 63–78
9. Franke, M.: Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics & Pragmatics* **4** (2011) 1–82
10. Jäger, G.: Rationalizable signaling. *Erkenntnis* (2013)
11. Franke, M., Jäger, G.: Pragmatic back-and-forth reasoning. manuscript, Amsterdam & Tübingen (2013)
12. Kramer, E., van Deemter, K.: Computational generation of referring expressions: A survey. *Computational Linguistics* **38** (2012) 173–218
13. Gatt, A., van Gompel, R.P.G., van Deemter, K., Kramer, E.: Are we Bayesian referring expression generators? In Knauff, M., Pauen, M., Sebanz, N., Wachsmuth, I., eds.: *Proceedings of the 35<sup>th</sup> Annual Meeting of the Cognitive Science Society*. (2013)



14. Jeffreys, H.: Theory of Probability. 3rd edn. Oxford University Press, Oxford (1961)
15. Kass, R.E., Raftery, A.E.: Bayes factors. Journal of the American Statistical Association **90** (1995) 773–795
16. Jaynes, E.T.: Probability Theory: The Logic of Science. Cambridge University Press, Cambridge, UK (2003)
17. Vandekerckhove, J., Matzke, D., Wagenmakers, E.J.: Model comparison and the principle of parsimony. In Busemeyer, J., Townsend, J., Wang, Z.J., Eidels, A., eds.: Oxford Handbook of Computational and Mathematical Psychology. Oxford University Press, Oxford (to appear) ???–???
18. Dickey, J.M., Lientz, B.P.: The weighted likelihood ratio, sharp hypotheses about chances, the order of a Markov chain. The Annals of Mathematical Statistics **41** (1970) 214–226
19. Wagenmakers, E.J., Lodewyckx, T., Kuriyal, H., Grasman, R.: Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. Cognitive Psychology **60** (2010) 158–189