# Super Semantics – Mini-term Paper Odd Utterances in the RSA Framework

Samuel Debray

January 28, 2022

#### Abstract

This mini-term paper tackles the RSA model's account for scalar implicature and its flaws in the context of odd utterances like *some of the marbles sank*. We review two solutions from the literature, namely Degen et al.'s (2015) and Spector's (2017), and propose an alternative solution which is inspired by Degen et al.'s.

### Contents

1	Introduction	1
2	The Rational Speech Act framework         2.1       The RSA model	<b>2</b> 2 2
3	The literature's solutions to the odd utterance problem         3.1       Wonky worlds (wRSA)         3.2       Semantic exhaustivity operator         3.3       Disscussion and criticism	<b>3</b> 3 4 5
4	Alternative solutions and follow-up	5
<b>5</b>	Conclusion	7

### **1** Introduction

The process of understanding language is much more complicated than the mere decoding of syntactic trees with fixed meanings. Depending on the context, the relation between the speaker and the listener, or even their state of mind, a single sentence can have a whole lot of different meanings. For example, *Mary has told Jane a billion times not to put her feet on the couch* does not genuinely mean that Mary told Jane a billion times, but conveys additional information about the way the speaker feels about Jane not listening to Mary when she repeatedly asks her not to put her feet on the couch. Grice [4] tried first to give an account for the context-sensitivity of language with his maxims but these are quite difficult to formalise, and fail to give a good account for such phenomena as sarcasm or irony.

Goodman and Frank [3] proposed, following Grice's line, that speaker and listener be modeled by Bayesian agents reasoning about each other and trying to maximise some utility function. This framework proves to be very general and easily generalisable to reason jointly about other properties (e.g. type of speaker, see [3]). The key assumption of this model is that both speaker and listener are rational and of goodwill.

However, in the case of odd utterances, such as *Max threw fifteen marbles in the water, some of the marbles sank*, RSA fails to provide a good account for the scalar implicature – that not all marbles did – which we sense that must triggered. In section 2 we present in more details the RSA framework, its assets and limitations in the context of odd utterances. Sections 3 and 4 are dedicated to proposing

and evaluating solutions, either from the literature (section 3) or from personal considerations (section 4), and comparing them with available data.

### 2 The Rational Speech Act framework

#### 2.1 The RSA model

The RSA framework represents the speaker and the listener as Bayesian agents. Speaker is a somewhat rational (rationality is captured by the parameter  $\lambda$  in equation 2.3) agent who wants to maximise the utility of her message, i.e. find a tradeoff between the quantity of information conveyed and the cost of utterance. Given a prior probability distribution on the possible states of the world **P**, Listener uses Bayes law to update his beliefs about the current state world. Thus, Speaker and Listener reason about each other to update their behaviour. So as to avoid infinite recursion, the model considers a naive literal listener L<sub>0</sub> who interprets messages without taking Speaker's intentions into account : if a given state s is compatible with the utterance u, then L<sub>0</sub> considers that the state of the world is s with probability  $\mathbf{P}(s \mid [\![u]\!])$ , where  $[\![u]\!]$  denotes the subset of states which are compatible with u. The model is described by the following equations.

$$\mathcal{L}_{0}(s \mid u) = \delta_{\llbracket u \rrbracket(s)} \frac{\mathbf{P}(s)}{\mathbf{P}(\llbracket u \rrbracket)}$$
(2.1)

$$\forall n \in \mathbf{N} \cdot U_{n+1}(u \mid s) = \log(\mathsf{L}_n(s \mid u)) - c(u)$$
(2.2)

$$\forall n \in \mathbf{N} \cdot \mathsf{S}_n(u \mid s) = \frac{\exp(\lambda U_n(u \mid s))}{\sum_{u'} \exp(\lambda U_n(u' \mid s))}$$
(2.3)

$$\forall n \in \mathbf{N}^* \cdot \mathsf{L}_n(s \mid u) = \frac{\mathbf{P}(s) \cdot \mathsf{S}_n(u \mid s)}{\sum_{s'} \mathbf{P}(s') \cdot \mathsf{S}_n(u \mid s')}$$
(2.4)

Here, we suppose we are given a set of states of the world  $\Omega$  along with a prior probability distribution for Listener  $\mathbf{P} \colon \Omega \to [0,1]$ , a set of utterances and, for each utterance u, a cost c(u) and a semantics  $\llbracket u \rrbracket \colon \Omega \to \text{Bool such that } \llbracket u \rrbracket(s) = \top \text{ iff } s \text{ is compatible with } u.$ 

As explained above, equation 2.1 describes the behaviour of a naive literal listener. Equation 2.2 describes how Speaker at level n + 1 updates her utility function given the behaviour of Listener at level n: it is the difference between the quantity of information conveyed by u from  $L_n$ 's perspective and the cost c(u) of said utterance. Equations 2.3 and 2.4 model the behaviour of Speaker and Listener at level n. They can be rewritten for first-level speakers and listeners as follows.

$$U_1(u \mid s) = \log(\mathsf{L}_0(s \mid u)) - c(u) \tag{2.5}$$

$$\mathsf{S}_1(u \mid s) \propto \exp(\lambda U_1(u \mid s)) \tag{2.6}$$

$$\mathsf{L}_1(s \mid u) \propto \mathbf{P}(s) \cdot \mathsf{S}_1(u \mid s) \tag{2.7}$$

In equations 2.3 and 2.6, we see that  $\lambda > 0$  captures the amount of rationality of Speaker : when  $\lambda$  grows to  $+\infty$ ,  $S_n$  is a Dirac distribution centered on the u which maximises  $U_n(u, s)$ . When  $\lambda$  is finite, Speaker is thus only approximately rational. Equations 2.4 and 2.7 are just applications of Bayes law as previously explained.

It is not clear how deep social recursion actually is, but it seems most authors consider only firstlevel speaker and listener  $S_1$  and  $L_1$  as defined in 2.6 and 2.7. With these settings, the RSA model has shown good fits with experiments using betting paradigms to assess confidence in signaling games [3], and provides a pretty good account for scalar implicature too.

#### 2.2 The odd utterance problem

Consider a situation where Speaker sees a pool with fifteen marbles and tells Listener some of the marbles sank. Denoting  $\Omega = \{s_i \mid 0 \le i \le 15\}$  the possible states ( $s_i$  is the world where exactly i marbles sank) and  $U = \{u_{\text{none}}, u_{\text{some}}, u_{\text{all}}\}$  the set of possible utterances, equation 2.7 gives

$$\mathsf{L}_{1}(s_{15} \mid u_{\text{some}}) = \frac{\mathbf{P}(s_{15})}{\mathbf{P}(s_{15}) + \left(1 + \exp\left(\lambda \left(1 - \log\frac{\mathbf{P}(s_{15})}{1 - \mathbf{P}(s_{0})}\right)\right)\right) \left(1 - \mathbf{P}(s_{0}) - \mathbf{P}(s_{15})\right)}$$
(2.8)

taking  $c(u_{\text{none}}) = c(u_{\text{some}}) = c(u_{\text{all}}) = 0$  and  $\llbracket u_{\text{some}} \rrbracket = \{s_i \mid i \neq 0\}.$ 



Figure 1: Evolution of  $L_1(s_{15} | u_{\text{some}})$  as a function of the prior  $\mathbf{P}(s_{15})$  under the assumption that for all  $i \neq 15$ ,  $\mathbf{P}(s_i) = \frac{1 - \mathbf{P}(s_{15})}{14}$  (all  $(s_i)_{i \neq 15}$  are equiprobable) and  $\lambda = 1$ .

Now assume that Listener has a strong prior that all marbles sank, i.e.  $\mathbf{P}(s_{15}) \sim 1$  and, for  $i \neq 15$ ,  $\mathbf{P}(s_i) \sim 0$ . Then equation 2.8 gives  $\mathsf{L}_1(s_{15} \mid u_{\text{some}}) \sim 1$ , that is to say that the scalar implicature that some but not all marbles sank disappears, as shown on figure 1. This prediction is highly unsatisfactory: intuitively, the subjective probability that all marbles sank given the utterance  $u_{\text{some}}$  is close to 0.

### 3 The literature's solutions to the odd utterance problem

In this section, we present two solutions to the odd utterance problem presented in 2.2 from the literature [2], [7]. Section 3.3 is a comparison and criticsm of these two solutions.

#### 3.1 Wonky worlds (wRSA)

Degen et al. [2] proposed to introduce a new variable w to characterise the world's wonkiness. The intuition is the following: when listening to Speaker, Listener makes joint inference on the state of the world and the type of world (wonky or not). In the case of a normal world, the model behaves just as the baseline RSA model. However, in the case of a wonky world, Listener forgets all his priors and assumes a uniform distribution over  $\Omega$ . In other words, the priors themselves are uncertain. The model is described by the following equations.

$$\mathbf{P}(s \mid w) = \begin{cases} \frac{1}{\#\Omega} & \text{if the world is wonky} \\ \mathbf{P}(s) & \text{otherwise} \end{cases}$$
(3.1)

$$\mathsf{L}_{0}(s \mid u, w) = \delta_{\llbracket u \rrbracket(s)} \frac{\mathbf{P}(s \mid w)}{\mathbf{P}(\llbracket u \rrbracket \mid w)}$$
(3.2)

$$U_1(u \mid s, w) = \log(\mathsf{L}_0(s \mid u, w)) - c(u)$$
(3.3)

$$S_1(u \mid s, w) = \frac{\exp(\lambda U_1(u \mid s, w))}{\sum_{u'} \exp(\lambda U_1(u' \mid s, w))}$$
(3.4)

$$\mathsf{L}_{1}(s,w \mid u) = \frac{\mathbf{P}(w) \cdot \mathbf{P}(s \mid w) \cdot \mathsf{S}_{1}(u \mid s, w)}{\sum_{(s',w')} \mathbf{P}(w') \cdot \mathbf{P}(s' \mid w') \cdot \mathsf{S}_{1}(u \mid s', w')}$$
(3.5)

The above equations describe that  $L_0$ , the naive literal listener, just assumes a world w, so does  $S_1$  who thinks the world under study is common knowledge. The key assumption of the model is that Listener trusts Speaker's background assumption, that is to say  $P(s \mid w)$  represents the probability of s given the world w in 3.2 but it represents the probability of s given that Speaker believes w is

common knowledge in 3.5, so these two quantities must be equal. Under the previous assumption, equation 3.5 describes how  $L_1$  uses Bayes law to reason jointly about s and w.

Degen et al. run an experiment to assess their wRSA model, which has shown good fit with the data for  $\lambda = 2$ . However, in their analysis, Listener's wonkiness prior was a parameter to be fitted and the best fit was obtained for  $\mathbf{P}(w) = \frac{1}{2}$ . This is questionable as we will discuss in section 3.3.

### 3.2 Semantic exhaustivity operator

Spector [6] proposed a connexion between the RSA framework and the observed ambiguity between literal and exhaustive meanings [5]. An utterance A is always considered ambiguous between its literal meaning lit(A) = A and its exhaustified meaning exh(A) =  $A \land \neg B$  if B is an alternative of A such that  $A \not\vDash B$ . Thus exh can be seen as a sort of *only* operator. Spector has proposed this model in the context of plural definites, but Cremers et al. [1] proposed that it be tested as a solution to the anti-exhaustivity problem<sup>1</sup>. Since the odd utterance problem can be seen as an anti-exhaustivity issue (since the priors on  $\mathbf{P}(s_{15})$  are high, Listener interprets *some* as *some and all*), we can examine this proposal too.

In Spector's model, Speaker is aware that Listener may interpret her utterance in different ways and wants to maximise the expected utility of her message. In addition, the model allows for Listener to make joint inference on the state of the world s and the question under discussion (QUD) Q, in case it is unclear what QUD Speaker wants to address.

Suppose we are given a set of utterances u, costs c(u), a set of QUDs along with a prior probability distribution, a set of states of the world along with a prior probability distribution too and two interpretation functions  $i_{\text{lit}}$  and  $i_{\text{exh}}$ . The latter make the semantics function  $\llbracket u \rrbracket$  also depend on the interpretation :  $\llbracket u \rrbracket^i(s) = \top$  iff s is compatible with u under interpretation i and  $\llbracket \cdot \rrbracket$ , that is a functional of type Utt  $\rightarrow$  Int  $\rightarrow \Omega \rightarrow$  Bool. Assuming all the previous variables are independent, the model is decribed by the following equations.

$$\mathsf{L}_{0}(s,\mathcal{Q} \mid u,i) \propto \delta_{\llbracket u \rrbracket^{i}(s)} \mathbf{P}(s) \cdot \mathbf{P}(\mathcal{Q})$$
(3.6)

$$U_1(u \mid s, \mathcal{Q}) = \sum_i \mathbf{P}(i) \cdot \log\left(\sum_{s' \in \mathcal{Q}(s)} \mathsf{L}_0(s', \mathcal{Q} \mid u, i)\right) - c(u)$$
(3.7)

$$\mathsf{S}_1(u \mid s, \mathcal{Q}) \propto \exp(\lambda U_1(u \mid s, \mathcal{Q})) \tag{3.8}$$

$$\mathsf{L}_{1}(s,\mathcal{Q} \mid u) \propto \mathbf{P}(s) \cdot \mathbf{P}(\mathcal{Q}) \cdot \mathsf{S}_{1}(u \mid s,\mathcal{Q})$$
(3.9)

Here, a question Q is modeled by an equivalence relation  $\sim_Q$  over states of the world and Q(s) denotes the equivalence class of s under  $\sim_Q$ .

Now let us define  $s_{\exists} = \bigcup_{1 \le i \le 14} s_i$  the event "some but not all marbles sank". The possible states of the world are  $\Omega = \{s_{\exists}, s_{15}\}$  and the utterances  $u_{\text{some}}, u_{\text{all}}$  and  $u_{\text{some but not all}}$ , the latter having positive cost. We consider two different QUDs :  $\mathcal{Q}_{\text{total}}$ , where Speaker wants to communicate the actual state of the world ( $\mathcal{Q}_{\text{total}}(s) = \{s\}$  for every  $s \in \Omega$ ), and  $\mathcal{Q}_{\exists}$ , which is only interesed in knowing whether any marble sank ( $\mathcal{Q}_{\exists}(s) = \Omega$ ). If Speaker believes  $s_{15}$  and wants to answer  $\mathcal{Q}_{\text{total}}$ , she cannot utter  $u_{\text{some}}$  since  $U_1(u_{\text{some}} \mid s_{15}, \mathcal{Q}_{\text{total}}) = -\infty$  because  $[\![u_{\text{some}}]\!]^{i_{\text{exh}}}(s_{15}) = \bot$ . Thus, if Speaker utters  $u_{\text{some}}$  and believes  $s_{15}$ , she must be answering  $\mathcal{Q}_{\exists}$ , in which case  $u_{\text{some}}$  is a reasonable choice (though less reasonable than  $u_{\text{all}}$ ). Besides,

$$\mathsf{L}_{1}(s_{15} \mid u_{\text{some}}) = \mathsf{L}_{1}(s_{15}, \mathcal{Q}_{\text{total}} \mid u_{\text{some}}) + \mathsf{L}_{1}(s_{15}, \mathcal{Q}_{\exists} \mid u_{\text{some}}).$$
(3.10)

We have seen that  $L_1(s_{15}, Q_{\text{total}} | u_{\text{some}}) = 0$ , therefore  $L_1(s_{15} | u_{\text{some}})$  depends on the prior  $\mathbf{P}(Q_{\exists}) \sim 0$ : if Listener really believes that Speaker wants to answer  $Q_{\text{total}}$ , then  $\mathbf{P}(Q_{\exists}) \sim 0$  and  $L_1(s_{15} | u_{\text{some}})$ ; and if Listener believes that Speaker might be addressing  $Q_{\exists}$ , then  $L_1(s_{15} | u_{\text{some}}) \neq 0$  but, intuitively, the scalar implicature should not be triggered in this case anyway.

<sup>&</sup>lt;sup>1</sup>Anti-exhaustivity is the (unsatisfactory) prediction made by RSA that, if the prior on  $s_A$  is much lower than that on  $s_{A \wedge B}$ , then  $L_1$  will understand the utterance A as describing  $A \wedge B$ .

Lastly, this model predicts that, if Speaker believes  $s_{\exists}$ , she will have the choice between  $u_{\text{some}}$ and  $u_{\text{some but not all}}$ .  $u_{\text{some but not all}}$  will definitely be a less good option if she wants to answer  $Q_{\exists}$ , but depending on the value of  $c(u_{\text{some but not all}})$ , it maybe a better option if she wants to answer  $Q_{\text{total}}$ . By reasoning on this,  $L_1$  can also infer that, if  $c(u_{\text{some but not all}})$  is low enough, then Speaker is probably willing to address  $Q_{\exists}$  if she believes  $s_{\exists}$  and utters  $u_{\text{some}}$ .

#### 3.3 Disscussion and criticism

Cremers et al. [1] ran an experiment to compare the two models from sections 3.1 and 3.2 in the more general context of anti-exhaustivity. They found that the two models fit the data equally well (AIC of 492 for both models). For wRSA, they fitted the prior  $\mathbf{P}(w)$  and found the best fit for  $\mathbf{P}(w) = .86$  whereas, for the semantic exhaustivity one, they fitted the prior  $\mathbf{P}(\mathcal{Q}_{\text{total}})$  and found the best fit for  $\mathbf{P}(w) = .89$  (the prior on interpretation was assumed to be uniform  $\mathbf{P}(i_{\text{lit}}) = \mathbf{P}(i_{\text{exh}}) = \frac{1}{2}$ ). Since the two model make as good a prediction as each other, we need to compare them on other grounds.

First, there are some general considerations that should be mentioned. Spector's proposal is just an adaptation of a model he proposed to account for a very different phenomenon, hence it is more general than wRSA, which was designed ad hoc. On the other hand, wRSA is pretty much in the spirit of the RSA framework, variables can be added to reason jointly about different features (e.g. the type of speaker for uRSA [3]), whereas Spectors's is more complex since it adds two variables which do not have to do with contextual features. Indeed, while w measures the wonkiness of the world (an environmental parameter), i and Q capture a different kind of feature, which is not just the behaviour of Speaker since they are taken into account by L<sub>0</sub> (who is supposed to be naive). Spector's model seems thus a little less in the spirit of RSA.

However, to obtain so good a fit with experimental data, both Degen et al. and Cremers et al. have decided to fit the prior  $\mathbf{P}(w)$  as a parameter and they both found the best fit for very high values: .5 and .86 respectively. This is highly problematic. This would mean that, before hearing anything at all, Listener assumes with a probability greater than  $\frac{1}{2}$  that he lives in a wonky world (or that the situation itself is wonky, e.g. the marbles do not have the same density). It could have been wiser to empirically measure the prior  $\mathbf{P}(w)$ , but we can expect that this is not an easy parameter to measure. Perhaps this suggests that we do not manipulate  $\mathbf{P}(w)$  directly, as discussed in section 4.

Finally, Cremers et al. also fitted the costs  $c(u_{all})$  and  $c(u_{some but not all})$  (they took  $c(u_{some}) = 0$ ) and, while wRSA set  $c(u_{all}) = 0$  and  $c(u_{some but not all}) = .37$ , Spector's model assigned  $c(u_{all}) = 44$ and  $c(u_{some but not all}) = .2$ . This is because the experiment is conducted in a more general framework where  $u_{all}$  is seen as a more complex  $u_{some and all}$  (as opposed to  $u_{some but not all}$ ). The reason why the model needed to set a high cost for  $u_{some}$  is probably to ensure that it is never used to answer  $Q_{\exists}$ , as discussed in section 3.2. However, assigning so high a cost to  $u_{all}$  does not seem appropriate in the context we consider ( $u_{all}$  and  $u_{some}$  intuitively have the same cost) and thus Spector's model is questionable for this matter.

Nonetheless, given its generality and wRSA's unreasonable prior w, Spector's proposal seems to be slightly more adequate to account for the odd utterance problem.

### 4 Alternative solutions and follow-up

**Revising Degen et al.'s wRSA.** Cremers et al. [1] highlighted that equation 3.1 does not describe a Bayesian process given the way Degen et al. collected data. Indeed, they measured the prior  $\mathbf{P}(s)$ before the subject heard anything, and used this value as  $\mathbf{P}(s \mid \neg w)$ . This would only make sense if the measured prior is itself conditional, that is if subjects only communicated their prior conditional to  $\neg w$ . This remark and the fact that Degen et al. did not address this issue show that  $\mathbf{P}(s \mid w)$  is not the most intuitive feature to consider. Besides, it seems quite unrealistic to consider that Listener has a prior on the wonkiness of the world: he should instead have a prior on the state of the world and, given any state s, derive a probability of wonkiness  $\mathbf{P}(w \mid s)$ . This leads to rewrite Bayes law in equation 3.5 as

$$\mathsf{L}_1(s, w \mid u) \propto \mathbf{P}(s) \cdot \mathbf{P}(w \mid s) \cdot \mathsf{S}_1(u \mid s, w).$$
(4.1)

The prior  $\mathbf{P}(w)$  would thus no longer be a parameter to be fitted and  $\mathbf{P}(w \mid s)$  could be measured empirically by presenting subjects a state of the world and asking them to assess how bizarre it seems to them. Since we do not have access to Degen et al.'s experimental data, we cannot assess quantitatively the accuracy of this alternative model, but our intuition is that it could fix all the issues raised in 3.3 and provide a suitable solution to the odd utterance problem.

**Challenging the semantics of** *some.* One possible and naive way of solving the problem raised in 2.2 is to change the semantics of *some*. It is clear in some contexts that *some* may well mean *all*, as in *if you have eaten some of the cookies you found on the floor, you should go to the hospital* : you should definitely go to the hospital if you ate all of the cookies you found. However, in many other contexts, *some* triggers the implicature *some but not all*, as in *some students failed the class*. Maybe we should apply to *some* the same treatment that is usually applied to other ambiguous operators between inclusive and exclusive readings like *or* [5]. Given that most people, when facing a *or* situation, assume that the *all* situation is false<sup>2</sup>, we could just let  $[[u_{some}]] = \{s_i \mid 1 \le i \le 14\}$  in section 2.2, which would entail that  $L_1(s_{15} \mid u_{some}) = 0$ . This definition would be supported by Degen et al.'s experimental data, since they found that  $L_1(s_{15} \mid u_{some}) \simeq \frac{1}{8}$  even when  $P(s_{15}) = 1$ . However, this solution is not really satisfactory since it does not provide a unified vision of *some*, neither does it explain how we are to choose between the two possible semantics we proposed in other contexts.

**Greater social recursion.** Lastly, it is worth mentioning that the problem described in 2.2 persists even in the context of deeper social recursion. Indeed, while Cremers et al. and Degen et al. only consider first-level speakers and listeners, it could be tempting to have a look at what happens for deeper levels of social recursion. Figure 2 plots the evolution of  $L_k(s_{15} \mid u_{\text{some}})$  for  $k \in \{1, 2, 3, 4\}$  in the baseline RSA model and shows that, even if  $L_k(s_{15} \mid u_{\text{some}})$  seems to depend less on  $\mathbf{P}(s_{15})$  for small values of  $\mathbf{P}(s_{15})$  as k grows, the scalar implicature still disapears when  $\mathbf{P}(s_{15}) \sim 1$ .



Figure 2: Evolution of  $L_k(s_{15} \mid u_{\text{some}})$  for  $k \in \{1, 2, 3, 4\}$  under the assumption that for all  $i \neq 15$ ,  $\mathbf{P}(s_i) = \frac{1 - \mathbf{P}(s_{15})}{14}$  (all  $(s_i)_{i \neq 15}$  are equiprobable) and  $\lambda = 1$ .

Thus, although deepening social recursion was a tempting solution to consider, it actually does not solve the problem.

<sup>&</sup>lt;sup>2</sup>For example, when asked to judge if the statement *some of the marbles sank* is true when all marbles sank, most people will answer "no".

## 5 Conclusion

Although the RSA framework has proved very efficient in predicting the behaviour of pragmatic locutors in many contexts, it fails to provide an account for scalar implicature when priors are too high in favour of the richer interpretation, as it is the case with the odd utterance problem. A few authors have tried to adapt the RSA framework to provide a solution to this issue, and Spector's answer seems to be the most satisfactory so far, both because it gives a good account for the phenomenon under study and because it is not ad hoc – in that it was designed to solve a very different issue and happens to also handle the one at hand. However, the revision of Degen et al.'s model we proposed in section 4 could also be tested and compared to Spector's model, since it intuitively seems to be a reasonable alternative which follows the line of other RSA refinements like uRSA [3].

# References

- [1] Alexandre Cremers, Ethan G. Wilcox, and Benjamin Spector. *Exhaustivity and anti-exhaustivity in the RSA framework: Testing the effect of prior beliefs.* arXiv:2202.07023 [cs]. Feb. 2022. DOI: 10.48550/arXiv.2202.07023. URL: http://arxiv.org/abs/2202.07023.
- [2] Judith Degen, Michael Henry Tessler, and Noah D Goodman. "Wonky worlds: Listeners revise world knowledge when utterances are odd". In: *CogSci.* 2015.
- [3] Noah D. Goodman and Michael C. Frank. "Pragmatic Language Interpretation as Probabilistic Inference". In: *Trends in Cognitive Sciences* 20.11 (Nov. 2016), pp. 818-829. ISSN: 13646613. DOI: 10.1016/j.tics.2016.08.005. URL: https://linkinghub.elsevier.com/retrieve/pii/ S136466131630122X.
- [4] H. P. Grice. "Logic and Conversation". In: Speech Acts (Dec. 1975). Publisher: Brill, pp. 41–58. DOI: 10.1163/9789004368811\_003. URL: https://brill.com/view/book/edcoll/9789004368811/BP000003.xml.
- Philippe Schlenker. "The semantics-pragmatics interface". In: The Cambridge Handbook of Formal Semantics. Ed. by Maria Aloni and Paul Dekker. Cambridge: Cambridge University Press, 2016, pp. 664-727. ISBN: 978-1-139-23615-7. DOI: 10.1017/CB09781139236157.023. URL: https: //www.cambridge.org/core/product/identifier/CB09781139236157A036/type/book\_part.
- [6] Benjamin Spector. "The pragmatics of plural predication: Homogeneity and non-maximality within the rational speech act model". In: Proceedings of the 21st Amsterdam colloquium. 2017, pp. 435–444.
- [7] Ethan Wilcox and Benjamin Spector. "The Role of Prior Beliefs in The Rational Speech Act Model of Pragmatics: Exhaustivity as a Case Study". In: CogSci. 2019, pp. 3999–3105.