

Pragmatic coordination on context via definite reference

Larry Muhlstein (UC San Diego), Christopher Potts (Stanford), Michael C. Frank (Stanford), and Roger Levy (UC San Diego)

Background Significant progress has been made in understanding how context influences language use. However, our understanding of what context *is*, and how it is implicitly coordinated upon (Clark and Marshall, 1981; Hawkins, 1991), remains limited. In utterance interpretation, what counts as context is not the totality of information mutually known between participants, but rather some circumscribed subset of this information. For example, definite reference in *Let's go to the restaurant* may be easy to understand even when uttered on a street full of restaurants if the speaker and listener have previously agreed upon a plan of action regarding dinner. Here we explore the central role of language itself in solving this coordination problem: just as context influences language, language influences context, giving rise to complex patterns of joint inference.

Summary We study the power of definite descriptions like *the fork* to discriminate among contexts and thereby help language users solve the context coordination problem. We first present a simple experiment providing evidence for the idea that definites can drive contextual inferences even in highly concrete and delimited visual scenarios, and that such inferences interact with more general pragmatic reasoning. We then extend the Rational Speech Act model (RSA; e.g., Frank and Goodman 2012) to one that makes joint probabilistic inferences about language and context and show that the predictions align with our experimental data.

Motivation Because the referents of definite descriptions vary by context, a speaker's use of such phrases, together with pragmatic reasoning, could guide the listener toward a particular kind of context. Fig. 1a depicts an example. The complete visual scenario includes several indistinguishable forks, so *the fork* should fail to refer if the context is taken to be the complete scenario. In contrast, narrowing the context to include only the plate on the left (but not the right) would permit successful definite reference. Thus, if the listener assumes (i) a cooperative speaker and (ii) mutual knowledge that context can be flexibly narrowed (Frazier, 2008), then she should be able to jointly infer from *the fork* that context is circumscribed to the left plate and that the definite's intended referent is the fork on that plate. We also expect such reasoning to interact with implicature calculation. In Fig. 1b, for instance, definite description is successful if context is narrowed to either plate, but *the fork* should bias in favor of the left-hand plate due to ad-hoc scalar implicature, since *the spoon* could have been used to identify the right-hand context unambiguously.

Experiment We tested the above predictions with the scenarios in Fig. 1. Each participant was presented with one of the two-plate scenarios in the left-hand column of the figure and asked to make a forced-choice decision between them based on the utterance prompt. In the DEFINITE condition, the prompt read "The plate with the fork is mine". In the ONE-WORD condition, the prompt read "The speaker can only use a single word to refer to a plate, and says 'fork'." (These one-word prompts have been shown to work well for studying ad-hoc scalar implicature – e.g., Frank and Goodman, 2012 – and allow us to largely avoid complex issues surrounding bare plurals like *forks* and indefinite phrases like *a fork*, which are likely to generate their own specificity inferences.) The middle panel of Fig. 1 summarizes responses for each condition. (A separate PRIOR-SALIENCE condition, not depicted in the figure, used the prompt "Imagine that someone is about to talk to you about one of the plates. Which one do you think it will be?" and yielded at-chance preferences between plates.) The bars depicts the proportion of choices for the left context. The results strongly support our guiding hypotheses. First, using *the fork* in the Fig. 1a scenario led subjects to infer the one-fork context. The

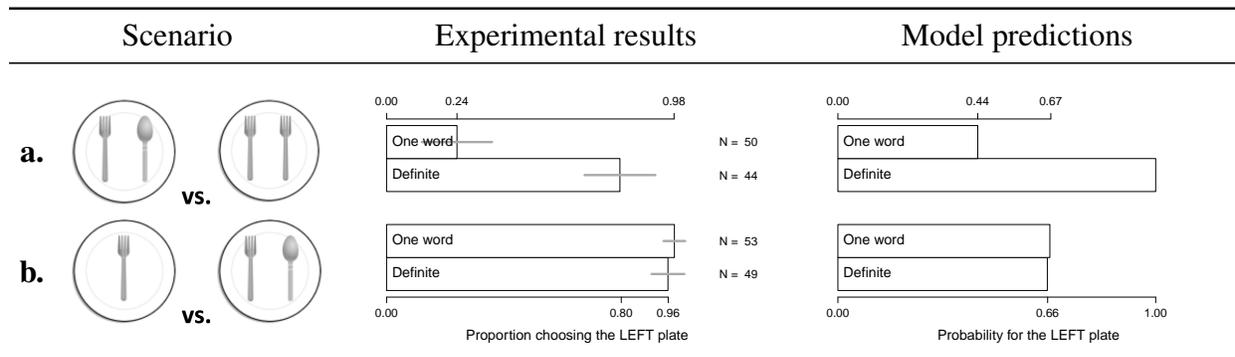


Figure 1: Summary of pilot experiment and model predictions.

definite led to the same pattern in the Fig. 1b scenario, but here the primary contextual disambiguation seems to trace to use or avoidance of *spoon*.

Model We began with the RSA model in which speakers and listeners are assumed to reason in terms of a fixed context (see also iterated best response models; Jäger 2007; Franke 2009). This model captures, in formal terms, the the back-and-forth pattern characteristic of Gricean reasoning, and it has in turn been shown to achieve many kinds of conversational implicatures. To capture contextual uncertainty and contextual inference, we propose three crucial modifications. First, building on formal properties of the lexical uncertainty model of Bergen et al. (2014), we treat each context as an independent inference problem with its own set of priors over potential referents. Second, each context additionally has a special ‘Fail’ state indicating a failure of reference. This state is given very low prior probability, reflecting our assumption that speakers generally tend to refer and that (charitable) listeners will conclude that reference failed only as a last resort. The result is a probabilistic version of a presuppositional definite determiner (for discussion and qualification, see Beaver 2001; Coppock and Beaver 2012). Third, to characterize our pragmatic listener, the inferences from each context are combined into a single set of joint context–referent probabilities, and a further marginalization step over referents determines the probability for each context given the speaker’s utterance. The predictions of this model are given in the right panel of Fig. 1. They achieve the same qualitative preferences as human data, whereas standard RSA has no account of how successful reference is achieved in the scenario of Fig. 1a.

Conclusion The results suggest that humans can leverage the requirements of definite reference to help coordinate on their presumed of utterance, and that this process is interwoven with general pragmatic reasoning. We have also shown how the lexical-uncertainty variant of the RSA model can be readily adapted to the problem of context coordination, allowing listeners to flexibly infer both context and reference from utterances using general principles of cooperative communication.

References

- Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI.
- Bergen, L., Levy, R., and Goodman, N. (2014). Pragmatic reasoning through semantic inference. <http://web.mit.edu/bergen/www/papers/BergenLevyGoodman2014.pdf>.
- Clark, H. H. and Marshall, C. R. (1981). Definite reference and mutual knowledge. In Joshi, A., Weber, B. H., and Sag, I. A., editors, *Elements of discourse understanding*, pages 10–63. Cambridge University Press.
- Coppock, E. and Beaver, D. I. (2012). Weak uniqueness: The only difference between definites and indefinites. In Chereches, A., editor, *Proceedings of Semantics and Linguistic Theory 22*, pages 527–544. Ithaca, NY. CLC Publications.
- Frank, M. C. and Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084):998.
- Franke, M. (2009). *Signal to Act: Game Theory in Pragmatics*. Institute for Logic, Language and Computation, University of Amsterdam.
- Frazier, L. (2008). The big fish in a small pond: Accommodation and the processing of novel definites. Ms., UMass Amherst.
- Hawkins, J. A. (1991). On (in)denite articles: implicatures and (un) grammaticality prediction. *Journal of Linguistics*, 27(2):405–442.
- Jäger, G. (2007). Game dynamics connects semantics and pragmatics. In Pietarinen, A.-V., editor, *Game Theory and Linguistic Meaning*, pages 89–102. Elsevier.