EasyChair Preprint

№ 9409

Presupposition Triggering Reflects Pragmatic Reasoning About Utterance Utility

Alex Warstadt

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

December 3, 2022

Presupposition Triggering Reflects Pragmatic Reasoning About Utility*

Alex Warstadt

ETH Zürich
Zürich, Switzerland
awarstadt@ethz.ch

Abstract

I present an RSA account of presupposition triggers building off prior work by Qing, Goodman, and Lassiter (2016). Intuitively, presupposition accommodation can occur through a listener's reasoning about the speaker's model of the common ground. Listeners assume that speakers try produce utterances with high contextual utility, and they will modify their own understanding of the context to rationalize low utility utterances. The account derives phenomena including accommodation of genus-species presuppositions due to Abusch (2002), QUD sensitivity of presuppositions, and gradience in the strength of some presuppositions.

1 Introduction

Theories of presupposition *triggering* (Abusch 2002; Simons, Tonhauser, et al. 2010; Abrusán 2011; Schlenker 2021) must contend with defeasibility and context sensitivity of projective inferences arising from the diverse class of *soft triggers* (Abusch 2002; Abusch 2010). This work presents a pragmatic account of triggering, building off notions of utterance utility from the Rational Speech Acts (RSA) framework (Frank and Goodman 2012; Qing, Goodman, and Lassiter 2016): I propose that a listener is more likely to accommodate a proposition p in circumstances where the speaker's utterance has higher utility if p is in the common ground. More specifically, listeners and speakers necessarily have separate models of the common ground (Stalnaker 1978), which they try to coordinate as much as possible. However, if their models of common ground diverge, sometimes the speaker may produce an utterance which they incorrectly judge will be informative to the listener. In such cases, a pragmatic listener can reason about what the speaker's model of common ground might have looked like and update their beliefs accordingly.

This basic story has been suggested previously in various forms (Stalnaker 1974; Simons 2001; Simons, Beaver, et al. 2017). Most relatedly, Qing, Goodman, and Lassiter (2016) propose an extension to the RSA model in which the pragmatic listener tries to infer not just the true state of the world, but also the speaker's model of the common ground. However, this earlier work fails to recognize some of the more interesting predictions of this kind of model. Exploring these predictions is the focus of this paper.

The Genus-Species Presupposition For the sake of illustration, I focus my discussion on genus-species presuppositions such as (1-a), first discussed by Abusch (2002). In these examples, a "species" predicate *have a green card* triggers a defeasible inference that a "genus" predicate

*I am grateful to Craige Roberts and Mandy Simons for extensive discussions on this topic. The analysis in this paper also benefited from input from Philippe Schlenker, Will Merrill, and Ethan Wilcox.

is a non-US citizen applies to Tom.[1]. As expected with presuppositions, this inference seems to project (defeasibly) out of negation (1-a) given that the Question Under Discussion (QUD) is as in (2-a). It also projects universally out of quantifiers (1-b). However, Abusch 2002) also considers another context where the local bar is giving free drinks to anyone with a green card, and the QUD is whether Tom qualifies for a free drink (2-b). In this context, (1-a) is reported to lack the genus presupposition.

(1) a. Tom doesn't have a green card.
b. None of the new hires has a green card.

(2) a. Does Tom need a visa?
b. Can Tom get a free drink?

The lexical semantic account of presupposition triggering assumed by many (e.g., Heim 1983; Schlenker 2009) struggles with soft triggers for two reasons: (a) While local accommodation provides a mechanism for blocking projection optionally, it does not explain why projection is only blocked in some special contexts; and (b) whereas the plausibility of lexical triggering is debatable for triggers like factives and change-of-state predicates, genus presuppositions are certainly not lexically encoded on species predicates (Abusch 2002; Abusch 2010; Schlenker 2021). Pragmatic approaches to these and other soft triggers have suggested that projection occurs when doing so increases the utility of the response (Simons 2001; Simons, Beaver, et al. 2017). However, the notion of utility in these works is left vague.

2 The RSA Approach to Accommodation

I adopt an earlier approach to presuppositions in RSA by Qing, Goodman, and Lassiter (2016) with a few modifications, as shown in (3).

(3) a. $L_0(Q(w)|u,C) \propto \sum_{w \in Q(w) \cap [[u]] \cap C} P(w)$
b. $S_1(u|Q,w,C) \propto exp(\alpha * Util(u||Q,w,C))$
c. $Util(u|Q,w,C) = \begin{cases} log(L_0(Q(w)|u,C)) - Cost(u) & if [[u]](w) = 1 \\ -\infty & if [[u]](w) = 0 \end{cases}$
d. $L_1(w,C|u,Q) \propto S_1(u|Q,w,C) * P(Q|w,C) * P(w|C) * P(C)$

Definitions The agents reason about a fixed set of worlds $W = \{w_1, ..., w_n\}$ and utterances $U = \{u_1, ..., u_m\}$, provided in Table 1 in the Appendix. Additionally, the model includes two contextual components: a context C and a QUD Q. Following Stalnaker's definition of the Context Set (Stalnaker 1978), a context in this model is the set of possible worlds consistent with publicly shared beliefs of the conversational agents. The QUD serves to delimits the topic of conversation (Van Kuppevelt 1995; Ginzburg 1996; Roberts 1996/2012). A QUD in this model is a partition over the set of possible worlds, where each cell represents a possible answer to the question (Groenendijk and Stokhof 1984). I write $Q(w)$ to denote the cell of the partition to which w belongs. The list of QUDs in the model is also provided in table 1.

Agents The literal listener in the model (3-a) is provided with a QUD, an utterance, and a context. They infer a probability distribution over the set of possible answers to the QUD by considering the prior probabilities of worlds consistent with the utterance and the context.

[1] A green card grants permanent residence to a non-US citizen. A US citizen cannot be a green card holder, and a green card holder does not need a visa to live or work in the US.

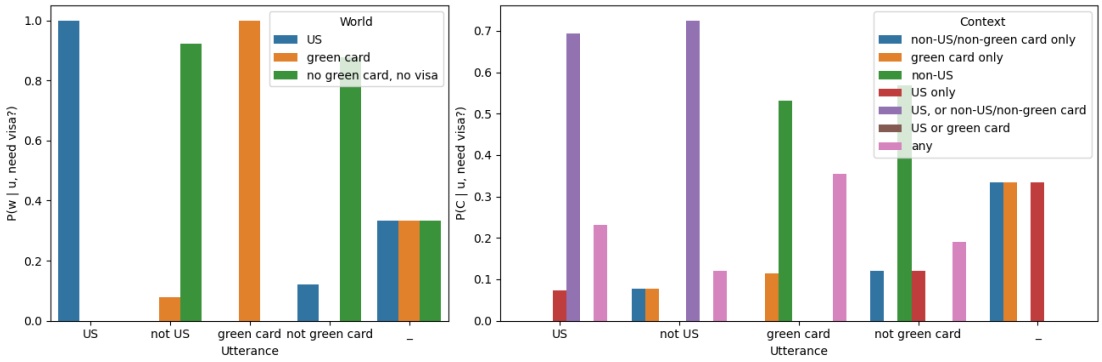

Figure 1: Pragmatic listener's posteriors over worlds (left) and contexts (right) given the QUD *Does Tom need a visa?*.

The speaker in the model (3-b) produces an utterance u with probability in proportion to the exponential of u's utility, times a rationality parameter α. The utility function is of key importance to the model. In prototypical RSA models (Frank and Goodman 2012), the speaker's goal is to reduce the listener's surprisal upon learning the actual state of the world. In this model, however, the speaker does not care to reduce the surprisal of the actual state of the world *per se*, but rather the correct answer to QUD. Furthermore, the speaker considers the reduction of surprisal with respect to C. These two factors will affect a speaker's choices in several ways. For instance, two utterances may have the same utility even if they differ in informativity overall, as long as they have the same informativity with respect to the QUD. Additionally, an utterance that has low utility in the weakest context $C = W$ may have high utility in a context C' that excludes some worlds.[2]

The pragmatic listener L_1 (3-d) differs from the prototypical RSA model in that they do joint inference over a context C in addition to reasoning over the set of possible worlds W. This is the main insight of Qing, Goodman, and Lassiter's model that I make use of. L_1 reasons about this distribution using Bayes' Rule. By assumption, $P(u|Q, w, C)$ is modeled as $S_1(u|Q, w, C)$, the speaker's probability distribution over utterances. The listener also considers three priors. First, there is a prior over QUDs $P(Q|w, C)$, which I assume to be uniform, for simplicity. Second, there is a prior over worlds given a context $P(w|C)$, which I take to be proportional to the prior over worlds $P(w)$ when $w \in C$, and 0 otherwise. Third, there is a prior over contexts C. The choice of this prior has significant implications for the predictions of the model, as I discuss in greater detail in Section 4. In the simplest case, this distribution is uniform over the set of all contexts, i.e. $\mathcal{P}(W)$.

3 Predictions

Genus-Species Presuppositions This model generates the basic genus-species presupposition. Recall that when the QUD is *Does Tom need a visa?* (2-a), the speaker's utterance

[2]The utility function also includes a condition that it returns $-\infty$ when the utterance is strictly false in the actual world. Without such an "honesty correction", the speaker will assign high utility to false utterances that happen to get the listener to believe the correct answer to the QUD. This is an interesting result which may be realistic for some speakers in some settings, and it has been leveraged to explain speakers' use of non-literal language (Kao et al. 2014). However, it is not crucial to account for projective inferences, and I find that adding this correction leads to more intuitive results.

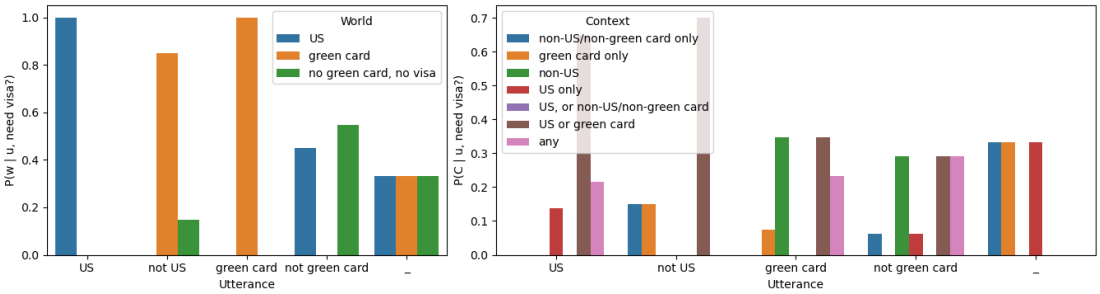

Figure 2: Pragmatic listener's posteriors over worlds (left) and contexts (right) given the QUD *Can Tom get a free drink on green card night?*.

Tom doesn't have a green card triggers an inference that Tom is not a US citizen (and therefore does need a visa). This inference is in essence a kind of relevance implicature, but the way in which the implicature is computed is not simply by reasoning about alternative utterances that the speaker could have produced, but also about alternative common grounds that the speaker could have been assuming.

The results for this model are shown in Figure 1.[3] The most interesting case is the one where the speaker's utterance is *not green card* (shorthand for (1-a)). In the listener's posterior over worlds, we can see that they assign significantly higher probability to the world in which Tom is a non-US citizen without a green card over the world in which Tom is a US citizen. This is notable, as without this kind of pragmatic reasoning, the literal listener assigns equal probability to these worlds given the same utterance.

The listener's posterior over contexts shows a related picture. When the utterance is *not green card*, the listener assigns over 50% probability mass to the context that entails that Tom is a non-US citizen. This is exactly the predicted accommodation. Interestingly, the listener still assigns a good deal of probability mass to other contexts, which may be a reasonable prediction, given the defeasibility of the genus-species presupposition.

QUD Sensitivity The model also predicts the QUD sensitivity of this inference. In particular, when the QUD is *Can Tom get a free drink on green card night?* and the utterance is *Tom doesn't have a green card*, there is little temptation to accommodate that Tom is a non-US citizen. Intuitively, this is because the speaker's utterance is an exhaustive answer in all contexts, meaning it is equally informative with respect to the QUD regardless of the common ground. In other words, there is no reasoning about the speaker's model of the common ground that will allow the listener to ascribe greater utility to the speaker's utterance.

This intuitive story is borne out in the results in Figure 2. In the listener's posterior over worlds given the utterance *no green card*, they assign nearly equal probability to the world in which Tom is a US citizen and the world in which he is a non-US citizen without a green card. Similarly, in the posterior over contexts, we see no special preference for the context that entails that Tom is a non-US citizen.

Family-Genus-Species Presuppositions I observe that the strength of Abusch's genus-species presupposition is related to the relative probability of the genus predicate and the

[3] In these and subsequent simulations, the value of α is set to 100. I found that a relatively high rationality parameter was important for generating the desired inferences.

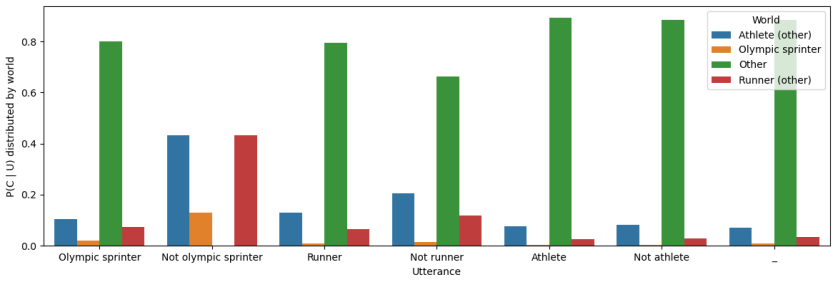

Figure 3: Pragmatic listener's posterior over contexts. Rather than plotting the probability for each of the 15 possible contexts, I plot the probability of each world according to the listener's posterior over contexts. This is not the same as the listener's posterior over worlds. Instead, given a context, I assume the probability of worlds is proportional to the world's prior probability. Thus the total probability of a world according to the listener's context posterior is the sum over all contexts of the world's probabilities in the context.

species predicate. Equivalently, rather than a pair of predicates, we can consider a triple of predicates, which I refer to as family, genus, and species. For example, if the species predicate is *Olympic sprinter*, a reasonable genus predicate is *runner*, and a reasonable family predicate is *athlete*. My observation is the projective inference that the family predicate is true (4-c) from negating the species predicate (4-a) is subjectively stronger than the projective inference from negating the genus predicate (4-b).

(4) a. Tom is not an Olympic sprinter.
b. Tom is not a runner.
c. Tom is an athlete.

This intuition fits well with the utility based story advocated above. In particular, Olympic sprinters are so rare, that there is almost no utility in asserting that Tom is not an Olympic sprinter in a typical context. Therefore, if a speaker utters (4-a), the listener may infer that the speaker is assuming a common ground in which it is significantly more likely that Tom could be an Olympic sprinter, for instance a context in which Tom is an athlete of some kind. By contrast, runners are common enough that it is typically informative to assert that Tom is not a runner. True, this assertion has even higher utility given the assumption that Tom is an athlete, but this inference is not so necessary to rationalize the speaker's utterance.

To test the model against this intuition, I use a different set of worlds and utterances from the previous examples, as illustrated in Table 2 in the Appendix. I also set a prior for each of the four worlds which reflects their approximate relative values. Note that while I do not vary the QUD in this example, the QUD can be implicitly understood to be *What is Tom's hobby*, or something similar.

Figure 3 shows the listener's probability distribution over contexts aggregated by world, with $\alpha = 1000$. For nearly all utterances, the listener infers that the speaker's model of the common ground makes it highly likely that Tom is a non-athlete. This reflects the high prior probability of this world. Crucially, this is not the case when the utterance is *Tom is not an Olympic sprinter*. In this case, the probability of the non-athlete world shifts to the worlds where Tom is a runner or another kind of athlete. Also as predicted, we see a much weaker version of this effect for the genus-family case, when the utterance is *Tom is not a runner*.

4 Discussion and Outstanding Issues

Extension to Other Triggers The success of this model for one kind of presupposition trigger (genus-species presuppositions) raises the question, "Can a similar story be told for all triggers?" For hard triggers such as additive particles like *too*, this is almost certainly not the case, as these triggers show almost no gradience or contextual sensitivity. However, there is a wide variety of other soft triggers for which this analysis is appealing, including change-of-state triggers and factives. Indeed, if we consider entities who have stopped smoking to be the "species", then entities who have smoked at some point may be considered the "genus". A classic example due to Geurts (1994) shows the contextual sensitivity of *stop* (5): Given an adequate QUD, there is less need to accommodate that Tom was ever a smoker.[4]

(5) Why is Tom chewing on his pencil? Did he recently stop smoking?

However, I am cautious to jump to the conclusion that this kind of on-the-fly pragmatic reasoning is solely responsible for the presuppositions of change-of-state predicates and factives. Intuitively, the inferences from these triggers are more indefeasible—and more conventionalized—than genus-species presuppositions. In fact, change-of-state predicates were the focus of Qing, Goodman, and Lassiter (2016)'s analysis, and they make some questionable predictions. Namely, they predict that given the QUD *Did Tom smoke in the past?*, the assertion *Tom didn't stop smoking* should give rise to a projective inference that Tom does not currently smoke, and that this inference should be just as likely as the typical initial state presupposition given the QUD *Does Tom smoke now?*.

Still, there is one parameter in the model that could potentially derive this asymmetry: the prior over contexts. If we assume that facts about Tom's past are more likely to be common ground than facts about Tom's present, this might help to explain why we more easily accommodate facts about the initial state of change-of-state predicates. A similar idea is proposed by Schlenker (2021).

Choosing What to Accommodate The analysis gives the listener great freedom to accommodate anything that might increase the utility of the speaker's utterance. In reality, listener's accommodations may be much more constrained. For instance, given the utterance *Tom is not an Olympic sprinter*, a listener could in theory accommodate that Tom has an Olympic gold medal, or that Tom has an appointment with a world famous orthopedist. Either accommodation would greatly increase the utility of the utterance, but they seem far less likely than the accommodation that Tom is a runner.

Part of the explanation must come down to the listener's prior over worlds. It is far more likely that Tom is a runner than that he is an Olympic gold medalist. If a listener is trying to find a *likely* explanation for a low utility utterance, they should accommodate facts which are more likely *a priori*. Indeed, Degen and Tonhauser (2022) do find that the likelihood of a projective inference is related to its prior probability in this way.

Future Work Future work should extend this analysis to projection of presupposition out of environments other than negation, including quantifiers, conditionals, and questions. It will also be important to test quantitative predictions of the RSA analysis against human judgments.

[4]The account could be extended to projection from interrogatives by operationalizing the utility of a question as its entropy, following van Rooy (2003) and van Rooy and Schulz (2004).

References

Abrusán, Márta (2011). "Predicting the presuppositions of soft triggers". In: *Linguistics and philosophy* 34.6, pp. 491–535.

Abusch, Dorit (2002). "Lexical alternatives as a source of pragmatic presuppositions". In: *Semantics and linguistic theory*. Vol. 12, pp. 1–19.

— (2010). "Presupposition triggering from alternatives". In: *Journal of Semantics* 27.1, pp. 37–80.

Degen, Judith and Judith Tonhauser (2022). "Prior beliefs modulate projection". In: *Open Mind* 5, pp. 59–70.

Frank, Michael C and Noah D Goodman (2012). "Predicting pragmatic reasoning in language games". In: *Science* 336.6084, pp. 998–998.

Geurts, Bart (1994). "Presupposing". PhD thesis. Universität Osnabrück.

Ginzburg, Jonathan (1996). "Dynamics and the semantics of dialogue". In: *Seligman, Jerry, & Westerst ahl, Dag (eds), Logic, language and computation* 1.

Groenendijk, Jeroen and Martin Stokhof (1984). "Studies on the semantics of questions and the pragmatics of answers". PhD thesis. Institute for Logic, Language & Computation (ILLC).

Heim, Irene (1983). "On the projection problem for presuppositions". In: *Formal semantics–the essential readings*, pp. 249–260.

Kao, Justine T et al. (2014). "Nonliteral understanding of number words". In: *Proceedings of the National Academy of Sciences* 111.33, pp. 12002–12007.

Qing, Ciyang, Noah D Goodman, and Daniel Lassiter (2016). "A rational speech-act model of projective content." In: *CogSci*.

Roberts, Craige (1996/2012). "Information structure: Towards an integrated formal theory of pragmatics". In: *Semantics and Pragmatics* 5, pp. 6–1.

Schlenker, Philippe (2009). "Local contexts". In: *Semantics and pragmatics* 2, pp. 3–1.

— (2021). "Triggering presuppositions". In: *Glossa: a journal of general linguistics* 6.1.

Simons, Mandy (2001). "On the Conversational Basis of Some Presuppositions". In: *Semantics and Linguistic Theory*. Vol. 11, pp. 431–448.

Simons, Mandy, David Beaver, et al. (2017). "The best question: Explaining the projection behavior of factives". In: *Discourse processes* 54.3, pp. 187–206.

Simons, Mandy, Judith Tonhauser, et al. (2010). "What projects and why". In: *Semantics and linguistic theory*. Vol. 20, pp. 309–327.

Stalnaker, Robert (1974). "Pragmatic presuppositions". In: *Semantics and Philosophy*. Ed. by Milton K. Munitz and Peter K. Unger. New York University Press, pp. 135–148.

— (1978). "Assertion". In: *Formal Semantics: The Essential Readings*, pp. 147–161.

Van Kuppevelt, Jan (1995). "Discourse structure, topicality and questioning". In: *Journal of linguistics* 31.1, pp. 109–147.

van Rooy, Robert (2003). "Questioning to resolve decision problems". In: *Linguistics and Philosophy* 26.6, pp. 727–763.

van Rooy, Robert and Katrin Schulz (2004). "Exhaustive interpretation of complex sentences". In: *Journal of logic, language and information* 13.4, pp. 491–519.

Appendix

	QUD_1 Need visa?	QUD_2 Free drink?	u_1 US	u_2 not US	u_3 GC	u_4 not GC	u_5 ϵ
w_1 (US citizen)	no	no	1	0	0	1	1
w_2 (Green card holder)	no	yes	0	1	1	0	1
w_3 (None of the above)	yes	no	0	1	0	0	1

Table 1: Worlds, utterances, and QUDs in the RSA model of example (1-a).

	prior	u_1 O.S.	u_2 Not O.S.	u_3 Runner	u_4 not runner	u_5 Athlete	u_6 Not athlete	u_7 ϵ
w_1 (Ol. Spr.)	1%	1	0	1	0	1	0	1
w_2(Runner (other))	5%	0	1	1	0	1	0	1
w_3(Athlete (other))	10%	0	1	0	1	1	0	1
w_4 (Other)	84%	0	1	0	1	0	1	1

Table 2: Worlds, utterances, and priors in the RSA model for family-genus-species inferences.