

MIT Open Access Articles

Strengthening, exhaustification, and rational inference

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Asherov, D., Fox, D. & Katzir, R. Strengthening, exhaustification, and rational inference. *Linguist and Philos* (2024).

As Published: 10.1007/s10988-023-09406-0

Publisher: Springer Netherlands

Persistent URL: <https://hdl.handle.net/1721.1/155947>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution





Strengthening, exhaustification, and rational inference

Daniel Asherov¹ · Danny Fox¹ · Roni Katzir²

Accepted: 14 December 2023
© The Author(s) 2024

Abstract The literature in semantics and pragmatics provides extensive evidence for the strengthening of linguistic expressions, both in matrix positions and when embedded under various operators. We study the properties of such strengthening using a very simple setting. Specifically, we look at when the expression “crate with a banana” can be understood as a unique crate even though two different crates have a banana in them. By varying the scenarios in which an expression such as “Pick the crate with a banana” is evaluated, we show that the strengthening of “crate with a banana” within the scope of the definite article parallels the entailments of “crate with only a banana” (with an overt exhaustivity operator, ‘only’). We use this observation to argue that strengthening in embedded positions follows the logic of an exhaustivity operator rather than that of rational inference. We then note that a similar pattern obtains in matrix positions.

Keywords Strengthening · Exhaustification · Rational inference

1 Introduction: strengthening in the scope of a definite article

Consider Scenario A in Fig. 1 (modeled after Stiller et al., 2011, 2015; Vogel et al., 2014): crate I is empty, crate II has a banana and nothing else in it, and crate III has both a banana and an apple in it. Given Scenario A, an utterance of (1) (= “Pick the

✉ Daniel Asherov
asherov@mit.edu

¹ Department of Linguistics and Philosophy, Massachusetts Institute of Technology, Cambridge, MA, USA

² Department of Linguistics and Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel



Fig. 1 Scenario A

crate with a banana”) is acceptable and can be used as a request to pick crate II.¹ The acceptability of (1) is surprising in light of the presuppositions of the definite article. In particular, the uniqueness presupposition of the definite article requires that there be no more than one salient individual in the extension of the NP ‘crate with a banana’, but in Scenario A there are two such individuals: crate II and III. The acceptability of (1) can be explained if the NP can be strengthened to be a predicate that has just crate II in its extension.

1. Pick the crate *with a banana*.

The strengthening of ‘crate with a banana’ to crate with only a banana in (1) requires an explanation, and in this note we will compare two potential mechanisms from the literature that could in principle account for it: (a) exhaustification, and (b) rational inference. We will describe each of them below. While the two accounts make different architectural assumptions (primarily on the semantics pragmatics divide), we find it convenient to focus first on the logic of strengthening without asking if the relevant computation is semantic or pragmatic. For much of this note we will continue to use the embedding under the definite article in (1), which serves as a helpful environment for studying the properties of strengthening, as it allows us to test predictions about strengthening through acceptability judgments. When there are two crates with a banana, as in Scenario A above and as in other scenarios that we will consider below, ‘the crate with a banana’ will be acceptable only if strengthening succeeds. More specifically, strengthening is needed to narrow down the NP to denote just a single crate, as required by the uniqueness presupposition of the definite article. This, in turn, will help us identify scenarios in which strengthening succeeds in yielding this narrowing (as in Scenario A above) and others where it fails (as in the variants considered below). This will also allow us to ask what characterizes the successful cases and whether this is predicted by either exhaustivity or rationality. Specifically, by varying the scenarios in minimal ways in Sects. 2–5, we will provide evidence that the embedded strengthening in (1) follows the logic of exhaustification and not that of rational inference.² In Sect. 6 we turn to matrix

¹ Some speakers judge (1) as less than perfect. What matters here is not the absolute acceptability of (1) but rather its relative acceptability in Scenario A and in the scenarios below, where (1) becomes clearly unacceptable.

² As we noted, the definite article makes the kind of investigation we are interested in relatively straightforward. Other environments are often less helpful. Consider the indefinite article, for example: ‘Pick a crate with a banana’ would in general be felicitous regardless of whether strengthening narrows down ‘crate with a banana’ to denote just a single crate, which makes it harder to tell whether strengthening has taken place and, if so, in what way. In principle one can set aside acceptability and simply use introspection to determine whether the relevant inference is possible. Studying strengthening in this way, however, is made difficult by the fact that strengthening is in general optional and might fail

positions and ask whether there, too, we can find evidence that teases apart the predictions of exhaustivity and rationality. While we cannot use the exact same methodology as in embedded positions, we will provide a different tool that will translate strengthening into felicity. By looking again at the same scenarios just mentioned we will conclude that strengthening follows the logic of exhaustivity rather than that of rationality also in matrix positions. Finally, in Sect. 7 we turn to the division of labor between grammatical computations and pragmatic inferences.

In the remainder of this section we briefly present the two potential approaches just mentioned to the embedded strengthening in (1). To facilitate the focus on the logical properties of the two approaches we will present them in terms of two silent strengthening operators: O-Exhaust and O-Rationality.³ Each can be thought of as attaching to the NP ‘crate with a banana’ in (1) and, in Scenario A at least, leading to the strengthened meaning—a predicate that picks out crate II (as opposed to crate III).⁴ Both operators are alternative-sensitive, and for the most part we will assume that the alternatives when attached to ‘crate with a banana’ are of the form ‘crate with an *x*’, where *x* is a single fruit name. We will revisit this assumption in Sect. 5.

O-Exhaust is a silent counterpart of the overt exhaustification operator ‘only’ (e.g., *Exh* of Fox, 2007). It affirms the prejacent (in (1), ‘crate with a banana’) and negates all of the alternatives to the prejacent that are safe to negate (e.g., ‘crate with an apple’, ‘crate with a pear’).⁵ The result, in the case of ‘crate with a banana’, is very similar to the meaning of ‘crate with *only* a banana’, which uses the overt exhaustivity operator ‘only’. Both with the silent O-Exhaust and with the overt ‘only’ the strengthened meaning—the result of affirming banana and negating other fruits—has just crate II in its denotation in Scenario A. As a singleton set this satisfies the presuppositions of the definite article, which in turn accounts for the felicity of sentence (1) in this scenario.

O-Rationality derives a similar strengthening of ‘crate with a banana’ so as to pick out crate II, but it does so through a different route: instead of negating alternatives, it attempts to capture how discourse participants might reason about the intended referents of various messages in a given context. There are of course different ways in which such reasoning might proceed, but in the case of Scenario A the following seems like a natural sequence of steps:

Footnote 2 continued

to occur for a variety of reasons and is subject to further complications when an indefinite article is involved, as indefinites often lead to anti-uniqueness inferences incompatible with the simple visual display in Scenario A. The definite article allows us to bypass these complications.

³ The choice is one of presentational convenience and does not imply the presence of operators in the syntax. An alternative possibility is that the relevant strengthening arises through local, meta-linguistic use of global pragmatics, along the lines of Horn (1985, 1989). See also Geurts and van Tiel (2013) on pragmatic truth-conditional narrowing in embedded positions.

⁴ There are other conceivable positions in which the operators might attach within (1). These extra possibilities will not affect our conclusions.

⁵ Informally, an alternative *m* is safe to negate given an assertion if doing so does not lead to contradiction (or to arbitrary choices in order to avoid contradiction). See Fox (2007) for a precise notion of when an alternative is safe to negate. See Spector (2016) for a comparison with other methods of dealing with potential contradictions.

2. Rational inference about Scenario A (schematic)

Step 1: Although the message ‘crate with a banana’ is not sufficiently narrow so as to pick out a specific crate, the message ‘crate with an apple’ is. Specifically, ‘crate with an apple’ identifies crate III since it is true of this crate and is false of every other crate.

Step 2: Given step 1, the message ‘crate with an apple’ and crate III are paired together and peeled off. At this point the message ‘crate with a banana’ is only true of crate II and identifies it.

Implicit in (2) is the notion of a noun phrase identifying a crate, or (more generally) a message identifying a state of the world. The following criterion, which we take from Fox and Katzir (2021) and revise below, will serve as a starting point:

3. **Semantic State Identification:** Given a set of messages M and a set of states T , a message identifies a state if it is true in that state and there is no other state in which it is true.

As with O-Exhaust, O-Rationality correctly predicts that the uniqueness presupposition of the definite article is satisfied and that (1) should be acceptable in Scenario A.

2 A problem for O-Rationality

Consider Scenario B in Fig. 2, which is based on Frank and Goodman (2012), Stiller et al. (2011), and Vogel et al. (2014). It differs from Scenario A in the following way: a pear is added to both crate I and crate II. Unlike Scenario A, here there is no state with *only* a banana.

Differently from Scenario A, in which (1) was judged as acceptable, in Scenario B the same sentence is judged as unacceptable, presumably because ‘crate with a banana’ is no longer an acceptable way to single out crate II (or any other individual crate for that matter).

O-Exhaust makes the correct prediction for Scenario B. If (1) is parsed without O-Exhaust, then, as in all the scenarios that we consider in this paper, both crate II and crate III are in the extension of ‘crate with a banana’, so the uniqueness presupposition of the definite article is not satisfied. Differently from Scenario A, however, parsing the sentence with O-Exhaust is now of little help. Specifically, the



Fig. 2 Scenario B

NP [*O-Exhaust* [crate with a banana]], just like [crate with only a banana], means crate with a banana and not with an apple and not with a pear. There is no such crate in this scenario, so the extension of the NP is empty, which leaves the existence presupposition of the definite article unsatisfied. Regardless of whether (1) is parsed with *O-Exhaust*, then, the sentence is correctly predicted to be unacceptable.

Consider now how *O-Rationality* would strengthen ‘crate with a banana’ in Scenario B:

4. *O-Rationality* in Scenario B

Step 1: Although the message ‘crate with a banana’ is not sufficiently narrow so as to pick out a specific crate, the message ‘crate with an apple’ is. Specifically, ‘crate with an apple’ would lead the listener directly to crate III, since it is false of the other two states. Thus, if a rational speaker had intended crate III, they would use the message ‘crate with an apple’.

Step 2: Given step 1, the speaker who used the message ‘crate with a banana’ did not intend crate III. Once crate III is eliminated as the intended state, the only remaining option is crate II.

It is easy to see that the reasoning in (4) is identical to that in (2), which was used for Scenario A. *O-Rationality* thus incorrectly predicts that listeners will take ‘crate with a banana’ to identify crate II in Scenario B as well, which satisfies the uniqueness and existence presuppositions of the definite article. This, in turn, predicts that (1) should be acceptable in Scenario B, contrary to fact.

3 Attempting to rescue *O-Rationality* through probabilistic identification and blocking iterations

Is there a difference between Scenarios A and B that might be relevant for the workings of *O-Rationality* so as to overcome this problem of overgeneration? A potentially relevant difference pertains to the perspective of a speaker that needs to select from among the alternative messages that are true in a given state. In Scenario A, the message ‘crate with a banana’ is the only message (among the assumed set of alternatives) that such a speaker can select for crate II but one out of two possible messages for crate III. In Scenario B, on the other hand, ‘crate with a banana’ is one out of two possible messages both for crate II and for crate III. This could be stated in probabilistic terms as follows, on the assumption of a *naïve speaker*, who chooses randomly between the messages that are possible given the state they wish to refer to: in Scenario A, ‘crate with a banana’ will be uttered with probability 1 for crate II and with probability 0.5 for crate III, while in Scenario B, ‘crate with a banana’ will be uttered with probability 0.5 both for crate II and for crate III. Suppose, then, that we changed the identification criterion in a way that can exploit this distinction:⁶

⁶ This probabilistic notion of identification can be motivated in terms of Bayesian reasoning. Such a move brings *O-Rationality* closer to iterated models of rationality in the literature such as Iterated Best Response (Franke 2009, 2011) and the Rational Speech Act framework (Frank & Goodman 2012). The motivation, however, leads to the problematic prediction that strengthening should be sensitive to the

5. **Probabilistic State Identification (PSI):** Given a set of messages M and a set of states T , a message identifies a state if the likelihood that a naïve speaker would use the message to describe that state is higher than for any other state.

In Scenario A, O-Rationality with PSI correctly finds that the message ‘crate with a banana’ identifies crate II. Here, unlike the non-probabilistic State Identification, O-Rationality discovers this message-state pairing *within one step*.

In Scenario B, PSI still arrives at the incorrect result that the message ‘crate with a banana’ identifies crate II, just like non-probabilistic State Identification (see note 6):

6. O-Rationality in Scenario B (Probabilistic)

Step 1: Although the message ‘crate with a banana’ does not identify a state, the message ‘crate with an apple’ does (as does the message ‘crate with a pear’). Specifically, ‘crate with an apple’ would lead the listener directly to crate III, since $P(\text{‘crate with an apple’}|\text{crate III})=0.5 > 0=P(\text{‘crate with an apple’}|\text{crate II})=P(\text{‘crate with an apple’}|\text{crate I})$. Thus, if a rational speaker who is trying to identify a state with PSI had intended crate III, they would use the message ‘crate with an apple’. (Similarly, ‘crate with a pear’ identifies crate I.)

Step 2: Given step 1, the speaker who used the message ‘crate with a banana’ did not intend crate III. Once crate III is eliminated as the intended state, the only remaining option is crate II.

While PSI still makes the wrong prediction for Scenario B, there is now a distinction between Scenario A and Scenario B: while in Scenario A identification is achieved *within one iteration (in Step 1)*, in Scenario B this is only achieved *in the second iteration (only in Step 2)*.

Suppose, then, that this difference is responsible for the difference in the acceptability of (1) in Scenarios A and B. In particular, suppose that O-Rationality cannot proceed past the step of iteration needed for identification in Scenario A (past step 1 in our simplified system). If that were the case, ‘crate with a banana’ would still have crate II (and nothing else) within its extension in Scenario A, which would keep the presuppositions of the definite article satisfied as before. In Scenario B, on the other hand, the restriction to a single step will prevent identification, and both crate II and crate III will be in the extension of ‘crate with a banana’, which will correctly lead to unacceptability. The goal of the next two sections is to test this idea by constructing scenarios that, on the one hand, have no crate with a banana and no other fruit (so that O-Exhaust fails), but, on the other hand, make it possible for O-Rationality to identify crate II even under the idea that O-Rationality is

Footnote 6 continued

prior probabilities of states. See Degen et al. (2015), Fox & Katzir (2021), Schreiber & Onea (2021), and Cremers et al. (2023) for discussion. Note that whenever message m identifies state t by (3) it will also identify t by (5) (at least when the set of messages is finite). This means that the move to (5) does not block identification in Scenario B as we discuss below. The only effect of this move on scenarios A and B pertains to the number of steps needed for identification (fewer in A than in B).



Fig. 3 Scenario C

limited to no more than one iteration of PSI.⁷ We will find that, even with this limitation, O-Rationality still suffers from an overgeneration problem.

4 The problem arises even within a single iteration

Consider Scenario C (Fig. 3), where (1) is judged unacceptable. Like Scenario B, there are two crates with a banana but no crate with *only* a banana, so O-Exhaust predicts that the sentence will be unacceptable regardless of whether ‘crate with a banana’ is strengthened.

While O-Exhaust makes the correct prediction for Scenario C, PSI does not, even with the restriction to a single step of identification. This is so since $P(\text{‘crate with a banana’} | \text{crate II}) = 0.5 > 0.33 = P(\text{‘crate with a banana’} | \text{crate III})$. This should make ‘crate with a banana’ identify crate II within the first iteration, which in turn would satisfy the uniqueness presupposition of the definite article and make (1) acceptable, contrary to fact.

5 Further attempts to avoid the overgeneration problem for O-Rationality

We wish to discuss two further potential attempts to salvage O-Rationality in the face of the overgeneration problem identified above: one response based on the idea that a message that identifies a state must be no worse than other available messages that can serve the same purpose and another response based on the idea that sensitivity to probabilities might be more limited than assumed in PSI. We present minimal variations on Scenario C which will demonstrate that neither response can solve the overgeneration problem for O-Rationality.

5.1 No sub-optimal identifiers

The first response is based on the intuition that a speaker who wants to identify crate II in Scenario C above has a more obvious strategy than using the message ‘crate with a banana’, namely using the message ‘crate with a pear’ (after all, ‘crate with a

⁷ The idea that humans cannot proceed past a first step of PSI clashes directly with the idea that something like O-Rationality might account for other inferences such as conjunctive readings of disjunction (see Franke 2009, 2011, van Rooij 2010, and Fox & Katzir 2021). We set aside this concern for the purposes of the present discussion.




Fig. 4 Scenario D

pear’ is true exclusively in state II, while ‘crate with a banana’ is true in state II and also in another state). Assuming that this intuition can be grounded in a principle, then we can say that even though ‘crate with a banana’ is an identifier for crate II according to PSI, it is a suboptimal message. One can use this observation as the basis for a modification of O-Rationality so that the problem of overgeneration in Scenario C will not arise. (We thank Anton Benz, p.c., for bringing up this point.) This, together with restricting identification to only one iteration, could give us an O-Rationality story for Scenarios A–C.

But we think that this move will not be helpful. To see why, consider Scenario D (Fig. 4), which is a minimal variation on Scenario C such that a pear replaces the orange in crate III. In Scenario D, ‘crate with a banana’ still identifies crate II using PSI, just like it did in Scenario C. But differently from Scenario C, in Scenario D ‘crate with a pear’ is not a better message than ‘crate with a banana’ in any conceivable sense. This means that ‘crate with a banana’ would identify crate II in Scenario D by any modification of PSI that avoids identification by suboptimal messages. And yet, the utterance “Pick the crate with a banana” remains as bad in Scenario D as it was in Scenario C.

Note that as in Scenarios B and C, there are two crates with a banana but no crate with *only* a banana, so O-Exhaust makes the correct prediction that (1) will be unacceptable.

5.2 Granularity

The second response to the overgeneration problem for O-Rationality is based on the idea that our sensitivity to probabilities (at least when deriving strengthening) is not as fine-grained as assumed in PSI. In Scenario A, where PSI successfully allowed for ‘crate with a banana’ to identify crate II in the first iteration, there was a big difference between a naïve speaker’s probability of uttering ‘crate with a banana’ in crates II and III: 1 in crate II and 0.5 in crate III. In Scenarios C and D, where PSI incorrectly allowed for ‘crate with a banana’ to identify crate II (also in the first iteration), the difference between a naïve speaker’s probability of uttering ‘crate with a banana’ in the two crates was much smaller: 0.5 in crate II and 0.33 in crate III. To block identification by ‘crate with a banana’ in Scenarios C and D, then, one might propose a less fine-grained version of PSI that allows a message m to identify a state t only when $P(mlt)$ is sufficiently bigger than $P(mlt')$ for any other state, and one can imagine various statements of what counts as sufficiently bigger.

We think, however, that no such statement will succeed. To see why, consider the following variant of Scenario D: instead of using only single fruit names in our messages as we have done so far, we will now allow also for *conjunctions* of fruit names. In this setting, ‘Pick the crate with a banana and a pear’ is a good message and is understood as referring to crate II.⁸ ‘Pick the crate with a banana’, on the other hand, is a bad message, just as it was when we considered only single fruit names, a fact that O-Exhaust correctly predicts for reasons that are familiar by now. The problem for any attempt to modify PSI based on probability differences by a naïve speaker is that the two messages have the exact same probabilities in crates II and III: 0.33 in crate II and 0.14 in crate III. No way of restricting identification based on probability differences will therefore succeed in allowing ‘Pick the crate with a banana and a pear’ to identify crate II while preventing ‘Pick the crate with a banana’ from doing the same.⁹

6 Strengthening in matrix positions

Above we focused on strengthenings in embedded positions. This allowed us to rely on the effect of strengthening on acceptability (given the presuppositions of the definite article), which resulted in a clear pattern. We believe, however, that very similar facts obtain in matrix positions as well.

First, consider the exchanges in (7)-(9), where B’s response is only appropriate if it identifies a difference between two individuals; that is, it needs to attribute non-overlapping predicates to each of the individuals. While B’s response in (7) is felicitous, B’s responses in (8) and (9) are not. We claim that the explanation of this contrast must rely on strengthening. To understand why this is the case, we first note that the utterance by Speaker B in each exchange should not be acceptable unless a strengthening is derived. In (7), while the unstrengthened meaning of ‘Mary did some of the homework’ is compatible with her having done all of it (a meaning that would presumably be odd given the continuation of B’s response), the sentence can be strengthened to mean that she did some but not all of the homework, which identifies an appropriate difference between Mary and John. B’s response in this case is therefore felicitous. In (8), on the other hand, ‘ate an apple or a banana’ cannot be strengthened to deny eating a banana, and in (9) ‘lives in France’ cannot be strengthened to deny living in Paris. Consequently, neither case provides a difference between Mary and John, and in both cases B’s response is predictably odd. The candidates for strengthening in (7)-(9) have been discussed extensively in the literature, and the success or failure in each case is readily predicted both by O-Exhaust and by O-Rationality.

⁸ For example, one can imagine a context in which all the possible messages are explicitly mentioned at the beginning of the conversation, before any of them is used in a request.

⁹ As far as we can tell, this property of O-Rationality persists even if one assigns different costs to different messages, e.g., by message length, as has occasionally been suggested within the Rational Speech Act framework (see Bergen et al., 2016, Scontras et al., 2018, Cremers et al., 2023) and elsewhere.

7. **Speaker A:** What is the difference between Mary and John?
Speaker B: Mary did some of the homework. John did all of it.
8. **Speaker A:** What is the difference between Mary and John?
Speaker B: #Mary ate an apple or a banana. John ate a banana.
9. **Speaker A:** What is the difference between Mary and John?
Speaker B: #Mary lives in France. John lives in Paris.

With this context in place, we can now return to a case in which, differently from (7)–(9), the predictions of O-Exhaust and O-Rationality diverge. Consider the exchange in (10) between two speakers who look at crate II and crate III in Scenarios A–D above.

10. **Speaker A:** What is the difference between crate II and crate III?
Speaker B: In crate II there is a banana. In crate III there is an apple and a banana.

The exchange seems natural in Scenario A but not in Scenarios B–D. Without strengthening, what is said of crate II is that it has a banana in it and possibly other fruit as well. This is also true of crate III in all four scenarios, and hence does not in any way identify a difference between the two crates. The pattern of felicity of Speaker B’s response in (10) across the different scenarios is as expected if strengthening in matrix positions corresponds to O-Exhaust (which succeeds in strengthening in Scenario A but not in Scenarios B–D) and is unexpected if strengthening in matrix positions corresponds to one of the variants of O-Rationality discussed above (all of which succeed in strengthening in at least one of B–D). Thus, the same effect observed in the scope of the definite article in the preceding sections also holds in matrix position.

7 On the division of labor between semantics and pragmatics in deriving strengthenings

We saw that, across a range of scenarios, strengthening is possible exactly when ‘only’ yields strengthening. Specifically, we considered scenarios in which there are two crates with a banana, so without strengthening, the uniqueness presupposition of the definite article ‘the’ in ‘the crate with a banana’ is not satisfied. In such scenarios, “Pick the crate with a banana” is judged acceptable exactly when ‘crate with *only* a banana’ has just one individual in its extension. We then showed that a similar pattern obtains in matrix positions: strengthening corresponds to the entailments of ‘only’. This is, of course, expected under O-Exhaust. On the other hand, it is surprising for the view of strengthening as rational inference, as expressed in the variants of O-Rationality that we considered. One can of course imagine other notions of rational inference, and it is possible that a different conceptualization would lead to a version of O-Rationality that does get the pattern of acceptability right. As far as we can tell, however, reformulating O-Rationality according to

prominent theories of pragmatics in the literature—including the neo-Gricean approach (Horn, 1972), the Rational Speech Act model (Frank & Goodman, 2012), Iterated Best Response (Franke, 2009, 2011), and the elimination of weakly dominated strategies (Rothschild, 2013)—does not change the picture materially. We conclude that strengthening, both in embedded positions as in (1) and in matrix positions as in (10) follows the logic of exhaustification and not that of rational inference.

Obviously, our conclusion does not rule out that speakers use rational inference to reason about linguistic expressions. If we are right, though, such inferences do not lead to strengthening of the kind examined above. Rather, strengthening is computed as exhaustification along the lines of ‘only’, leaving the task of disambiguating between parses—deciding where in the structure exhaustification applies, among other things—to pragmatic reasoning. This view is in line with recent work such as Champollion et al. (2019), Franke and Bergen (2020), Fox and Katzir (2021), and Cremers et al. (2023).

Acknowledgements We thank Moysh Bar-Lev, Itai Bassi, Adi Behar Medrano, Emmanuel Chemla, Anton Benz, Alma Frischhoff, Yosef Grodzinsky, Martin Hackl, Roger Levy, Eyal Marco, Paul Portner, Viola Schmitt, Greg Scontras, Benjamin Spector, Tue Trinh, and several anonymous reviewers, as well as the audiences at MIT, the LSA, ZAS, ESSLLI, Tel Aviv University, and the University of Göttingen.

Funding Open Access funding provided by the MIT Libraries.

Declarations

Conflict of interest The authors have no competing interests to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bergen, L., Levy, R., & Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*. <https://doi.org/10.3765/sp.9.20>
- Champollion, L., Alsop, A., & Grosu, I. (2019). Free choice disjunction as a rational speech act. *Proceedings of SALT*, 29, 238–257.
- Cremers, A., Wilcox, E., & Spector, B. (2023). Exhaustivity and anti-exhaustivity in the RSA framework: Testing the effect of prior beliefs. *Cognitive Science*, 47(5), e13286.
- Degen, J., Tessler, M. H., & Goodman, N. D. (2015). Wonky worlds: Listeners revise world knowledge when utterances are odd. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimimi, T. Matlock, C. J. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual conference of the cognitive science society* (pp. 548–553). New York: Springer.

- Fox, D. (2007). Free choice disjunction and the theory of scalar implicatures. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 71–120). London: Palgrave-Macmillan.
- Fox, D., & Katzir, R. (2021). Notes on iterated rationality models of scalar implicatures. *Journal of Semantics*, 38(4), 571–600.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998.
- Franke, M. (2009). *Signal to act: Game theory in pragmatics*. Amsterdam: Institute for Logic, Language and Computation.
- Franke, M. (2011). Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics and Pragmatics*, 4, 1–82.
- Franke, M., & Bergen, L. (2020). Theory-driven statistical modeling for semantics and pragmatics: A case study on grammatically generated implicature readings. *Language*, 96(2), e77–e96.
- Geurts, B., & van Tiel, B. (2013). Embedded scalars. *Semantics and Pragmatics*, 6(9), 1–37.
- Horn, L. R. (1972). *On the semantic properties of logical operators in English*. UCLA PhD thesis.
- Horn, L. R. (1985). Metalinguistic negation and pragmatic ambiguity. *Language*, 61(1), 121–174.
- Horn, L. R. (1989). *A natural history of negation*. Chicago: University of Chicago Press.
- Rothschild, D. (2013). Game theory and scalar implicatures. *Philosophical Perspectives*, 27, 438–478.
- Schreiber, A., & Onea, E. (2021). Are narrow focus exhaustivity inferences Bayesian inferences? *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2021.677223>
- Scontras, G., Tessler, M. H., & Franke, M. (2018). Probabilistic language understanding: An introduction to the Rational Speech Act framework. Retrieved 2021-7-8 from <https://www.problang.org>.
- Spector, B. (2016). Comparing Exhaustivity Operators. *Semantics and Pragmatics*, 9, 1–33.
- Stiller, A., Goodman, N., & Frank, M. C. (2011). Ad-hoc scalar implicature in adults and children. In: *Proceedings of the 33rd annual meeting of the Cognitive Science Society* (pp. 2134–2139).
- Stiller, A., Goodman, N., & Frank, M. C. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development*, 11(2), 176–190.
- van Rooij, R. (2010). Conjunctive interpretation of disjunction. *Semantics and Pragmatics*, 3, 1–28.
- Vogel, A., Emilsson, A. G., Frank, M. C., Jurafsky, D., & Potts, C. (2014). Learning to reason pragmatically with cognitive limitations. In *Proceedings of the 36th annual meeting of the cognitive science society* (pp. 3055–3060)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.