Free Choice Disjunction as a Rational Speech Act

It is well known that (1) has the Free Choice inference (FCI) (1a),(1b). More controversially, (1) may also lead to the exclusivity inference (EI) (1c). As (2) shows, EI is easier to cancel than FCI; as (3) shows, FCI disappears under negation. Unembedded disjunction lacks an analogue of FCI but may give rise to EI (4). Our work is situated within the literature that takes these facts to motivate nonsemantic accounts of FCI and EI, broadly speaking (e.g. Fox, 2007; Franke, 2011).

(1) You may take an apple or a pear. $\diamond(A \vee B)$

a. $\rightsquigarrow$ You may take an apple. $\diamond A$

b. $\rightsquigarrow$ You may take a pear. $\diamond B$

c. $\rightsquigarrow$ You may not take both. $\neg\diamond(A \wedge B)$

(2) a. You may take an apple or a pear. #In fact, you may not take an apple.

b. You may take an apple or a pear. In fact, you may take both.

(3) You may not take an apple or a pear. $\neg\diamond(A \vee B)$

a. $\not\approx$ You don't have both permissions; I leave open whether you have one. $\neg(\diamond A \wedge \diamond B)$

(4) John took an apple or a pear. $A \vee B$

a. $\not\rightsquigarrow$ John took an apple and a pear. $A \wedge B$

b. $\rightsquigarrow$ John did not take both an apple and a pear. $\neg(A \wedge B)$

We show that FCI and EI fall out of inference over LFs in a cooperative language game without any assumptions about speaker ignorance. We derive both FCI and EI using a game-theoretic model in the Rational Speech Acts framework (RSA, Frank and Goodman, 2012). Our work, like e.g. Potts et al. (2016), reconciles exhaustification-based models (Fox, 2007) with game-theoretic accounts in the style of iterated best response (IBR, Franke, 2011). On our account, when the speaker utters (1), the listener reasons about why the speaker did not choose alternatives such as (1a). A crucial ingredient in our explanation is uncertainty about LFs (cf. lexical uncertainty in Bergen et al., 2016). We assume that the speaker is unsure whether the listener might take (1a) as entailing a prohibition against taking a pear; this is analogous to Fox's optional exhaustification operator *Exh*. Uttering (1) as opposed to (1a) or (1b) is a way to prevent the listener from concluding about any fruit that it is forbidden to take it. Knowing this, the listener concludes that (1) signals FCI. Whether EI arises as well depends mainly on its prior probability.

Our model assumes a state space $\{\textsf{Only A}, \textsf{Only B}, \textsf{Only 1}, \textsf{Any \#}, \textsf{Only 2}\}$ where in Only A, A is allowed but B is forbidden; in Only 1 FCI and EI hold (any one fruit is allowed); in Any #, FCI holds but not EI, thus taking both fruit is allowed as well; and in Only 2, the only thing allowed is taking both fruit. Our utterances $\{u_{\diamond A}, u_{\diamond B}, u_{\diamond(A\vee B)}, u_{\diamond(A\wedge B)}\}$ are labeled with their meanings in the absence of *Exh*; however, we assume that there is uncertainty in the sense of Bergen et al. (2016), for example about whether the semantic meaning of (1a) = $u_{\diamond A}$ in the given context is $\{\textsf{Only A}, \textsf{Only 1}, \textsf{Any \#}, \textsf{Only 2}\}$ or $\{\textsf{Only A}, \textsf{Only 1}, \textsf{Any \#}\}$ or $\{\textsf{Only A}\}$. This uncertainty stems from different LFs that are available for (1a) due to optional insertion of *Exh* in the sense of Fox (2007): either simply $\diamond A$, or $\diamond Exh(A)$, or $Exh\diamond(A)$. Likewise, (1) = $u_{\diamond(A\vee B)}$ has at least the LFs $\diamond(A \vee B)$, $\diamond Exh(A \vee B)$, and $\diamond(Exh(A) \vee Exh(B))$ (the last two are equivalent). With Fox (2007) and similar approaches, we see insertion of *Exh* into an LF as a grammaticalized operation that is distinct from Gricean/Bayesian reasoning; like Potts et al. (2016), we go beyond Fox in explicitly modeling the coordination problem that arises from a silent *Exh* operator. We represent uncertainty about LFs via the form-meaning mappings $\mathscr{L}_1$, $\mathscr{L}_2$, $\mathscr{L}_3$ in (5)-(7) (cf. lexica in Bergen et al. 2016).

(5) $[\![u_{\diamond A}]\!]^{\mathscr{L}_1} = \{\mathsf{Only\ A}, \mathsf{Only\ 1}, \mathsf{Any\ \#}, \mathsf{Only\ 2}\}$, $[\![u_{\diamond B}]\!]^{\mathscr{L}_1} = \{\mathsf{Only\ B}, \mathsf{Only\ 1}, \mathsf{Any\ \#}, \mathsf{Only\ 2}\}$, $[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_1} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ 1}, \mathsf{Any\ \#}, \mathsf{Only\ 2}\}$, $[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_1} = \{\mathsf{Any\ \#}, \mathsf{Only\ 2}\}$

(6) $[\![u_{\diamond A}]\!]^{\mathscr{L}_2} = \{\mathsf{Only\ A}, \mathsf{Only\ 1}, \mathsf{Any\ \#}\}$, $[\![u_{\diamond B}]\!]^{\mathscr{L}_2} = \{\mathsf{Only\ B}, \mathsf{Only\ 1}, \mathsf{Any\ \#}\}$, $[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_2} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ 1}, \mathsf{Any\ \#}\}$, $[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_2} = \{\mathsf{Any\ \#}, \mathsf{Only\ 2}\}$

(7) $[\![u_{\diamond A}]\!]^{\mathscr{L}_3} = \{\mathsf{Only\ A}\}$, $[\![u_{\diamond B}]\!]^{\mathscr{L}_3} = \{\mathsf{Only\ B}\}$, $[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_3} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ 1}, \mathsf{Any\ \#}\}$, $[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_3} = \{\mathsf{Only\ 2}\}$

Our model is robust to certain changes in these assumptions. E.g., dropping $\mathscr{L}_1$ or $\mathscr{L}_2$ still generates FCI, as does adding $\mathscr{L}_i$ that mix elements of (5)-(7). — The RSA framework represents listeners as rational Bayesian interpreters and speakers as soft-max rational agents. Speakers' and listeners' reasoning obeys the recursive probabilistic functions in (8) for worlds w, utterances u, greedy optimality parameter α, and mappings $\mathscr{L}$. Here, $\mathscr{L}^u(w) = 1$ if $w \in [\![u]\!]^{\mathscr{L}}$, and 0 otherwise.

(8) a. $P_{listener\,0}(w|u,\mathscr{L}) \propto \mathscr{L}^u(w)P(w)$ b. $P_{speaker\,1}(u|w,\mathscr{L}) \propto \left[P_{listener\,0}(w|u,\mathscr{L})\right]^{\alpha}$
c. $P_{listener\,1}(w|u) \propto P(w)\sum_{\mathscr{L}} P_{speaker\,1}(u|w,\mathscr{L})$
d. For $n > 1$: $P_{speaker\,n}(u|w) \propto \left[P_{listener\,(n-1)}(w|u)\right]^{\alpha}$ e. $P_{listener\,n}(w|u) \propto P(w)P_{speaker\,n}(u|w)$

This model derives FCI for the level-1 pragmatic listener, in that for uniform priors $P(w)$, $\mathscr{L}_i$ as above, and sufficiently large α, the posterior distribution $P_{listener\,1}(\cdot|u_{\diamond(A\vee B)})$ splits its probability mass almost evenly between the FCI+EI world $\mathsf{Only\ 1}$ and the FCI-EI world $\mathsf{Any\ \#}$, with virtually no mass assigned to the non-FCI worlds. (For lower values of α, FCI arises only for $n > 1$-listeners.) For nonuniform priors $P(w)$ that assign some FCI world a high prior probability, $P_{listener\,1}$ upon hearing $u_{\diamond(A\vee B)}$ also assigns it a higher posterior probability; this derives the optionality of EI as a matter of prior knowledge, at least when using just (5)-(7). The low posterior probabilities of non-FCI worlds remain virtually unaffected by shifting a comparable amount of probability mass to any of them in the prior; this explains why FCI is a stronger inference than EI.

Our model captures the absence of FCI under negation under plausible assumptions about the possible meanings of the utterances involved; specifically, we assume that one of the LFs for (3) is equivalent to the classical $\neg\diamond(A\vee B)$, with no *Exh* inserted, which semantically entails $\neg\diamond A$ and $\neg\diamond B$. Other, weaker LFs which contain *Exh* lack this entailment; we show that their presence does not lead the listener to conclude FCI. By contrast, Fox (2007) relies on a stipulation that prevents *Exh* insertion into LFs whose semantic meaning would be weakened as a result (Chierchia, 2013).

Our model has much in common with IBR, a precursor of RSA which Franke (2011) applies to free choice. IBR is similar to RSA with $\alpha = \infty$, but there are crucial differences. First, Franke relies on the fact that $u_{\diamond(A\vee B)}$ emerges as a surprise message at higher levels. This turns out to prevent his analysis of free choice from working in RSA, where due to $\alpha < \infty$, any speaker will choose $u_{\diamond(A\vee B)}$ with nonzero probability. Second, Franke uses only $\mathscr{L}_1$; in RSA, this would cause level-1 listeners to put so little credence in the FCI worlds that FCI fails to emerge even at higher levels. Finally, to avoid predicting that plain disjunctions like (4) are surprise messages, Franke moves to a more complex setting that models speaker ignorance. We do not have to rely on this: For a state space $\{\mathsf{A}, \mathsf{B}, \mathsf{Both}\}$, utterances $\{u_A, u_B, u_{A\vee B}, u_{A\wedge B}\}$, with two $\mathscr{L}_i$ that disagree only on whether u_A and u_B are true at Both, a level-1 listener who hears $u_{A\vee B}$ will assign most of the probability mass to A and to B. To convey A, speakers will prefer u_A over $u_{A\vee B}$ but not entirely avoid the latter.

References

Bergen, L., Levy, R., and Goodman, N. D. (2016). Pragmatic reasoning through semantic inference. *Semantics and Pragmatics*, 9(20).

Chierchia, G. (2013). *Logic in grammar: polarity, free choice, and intervention*. Oxford University Press, Oxford, UK.

Fox, D. (2007). Free choice and the theory of scalar implicatures. In Sauerland, U. and Stateva, P., editors, *Presupposition and Implicature in Compositional Semantics*, pages 71–120. Palgrave Macmillan, London, UK.

Frank, M. C. and Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.

Franke, M. (2011). Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics and Pragmatics*, 4(1):1–82.

Potts, C., Lassiter, D., Levy, R., and Frank, M. C. (2016). Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. *Journal of Semantics*, 33(4):755–802.