Free Choice Disjunction as a Rational Speech Act

Lucas Champollion, Anna Alsop, and Ioana Grosu
{champollion,aalsop,ig950}@nyu.edu

January 8, 2019
Paris, Institut Jean Nicod
minor changes on May 15, 2019

1 Introduction

- Main fact to explain: Free Choice Inference
- (1) has the Free Choice inference (FCI) (1a),(1b).

(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. $\diamond B$

- More controversially, (1) may also lead to the exclusivity inference (EI) (2).

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

- The disjunction in (1) compares to unembedded disjunctions, which lack an analogue of FCI but may give rise to EI (3).

(3) 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. $\neg(A \wedge B)$

- Additional facts to explain
  - EI is easier to cancel than FCI.

(4) 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.

  - FCI disappears under negation.

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

a. $\not\approx$ You don't have both permissions, and it is open whether you have one. $\neg(\diamond A \wedge \diamond B)$

- We use these facts to motivate a nonsemantic account of FCI and EI which combines elements from Fox (2007) and Franke (2011).

2 The main idea

- We derive both FCI and EI using a game-theoretic model in the Rational Speech Acts framework (RSA, Frank and Goodman, 2012).
- We show that FCI and EI both arise from inference over LFs in a cooperative language game.
  - We show this without any assumptions about speaker ignorance.
- Our work reconciles exhaustification-based models (Fox, 2007) with game-theoretic accounts in the style of iterated best response (IBR, Franke, 2011) like Potts et al. (2016).
  - On our account, when the speaker utters (1), the listener reasons about why the speaker did not choose alternative utterances such as (1a).
  - We make use of 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 interpretation 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 that any fruit is forbidden to take.
  - Knowing this, the listener concludes that (1) signals FCI.
  - Whether EI arises as well depends mainly on its prior probability.

3 Modeling Pragmatics with the RSA framework

- Models in the RSA framework see communication as a speaker and a listener recursively reasoning about each other's goals and behavior (See Fig. 1). They optimize their utterances and interpretations based on this reasoning.
- In the model, a "literal" listener (L0) only has access to the literal semantic denotations of each utterance. If an utterance may correctly describe multiple worlds, the listener chooses according to their prior beliefs, or if they have none, at random. We refer to this base level of reasoning as level-0.
- We next build a pragmatic speaker (S1) who reasons about the literal listener. The pragmatic speaker usually chooses whatever utterance leads the listener to arrive at the intended interpretation most of the time.

- The optimality parameter α makes a speaker more optimal when increased. We will informally refer to α as "temperature".
- Next, we build a pragmatic listener (L1) that reasons about a pragmatic speaker. Upon hearing an utterance, the listener tries to infer which world they are in. Among the possible worlds, they put the highest probability on whatever world best explains why the speaker chose the utterance they did.
- The recursion may proceed indefinitely over an infinite set of levels.

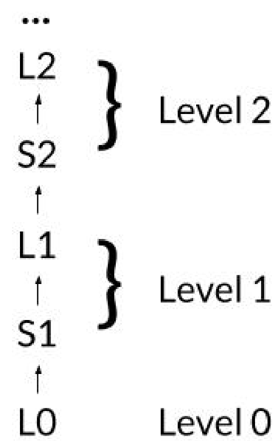

Figure 1: Levels of Recursion in RSA

- For reference, our official model is defined in the following equations:

(6)
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. $P_{speaker\,n}(u|w) \propto \left[P_{listener\,(n-1)}(w|u)\right]^\alpha$ $(n > 1)$
e. $P_{listener\,n}(w|u) \propto P(w)P_{speaker\,n}(u|w)$

- Here, $\mathscr{L}^u(w) = 1$ if $w \in [\![u]\!]^{\mathscr{L}}$, and 0 otherwise.

3.1 World States

- Our RSA model assumes the following world states: $\{\textsf{Only A}, \textsf{Only B}, \textsf{Only One}, \textsf{Any Number}, \textsf{Only Both}\}$ where:
  - in Only A, taking an apple is allowed but taking a pear is forbidden;
  - in Only One FCI and EI hold (any one fruit is allowed);
  - in Any Number, FCI holds but not EI, thus taking both fruit is allowed as well;
  - and in Only Both, the only thing allowed is taking both fruit (as a package deal).

3.2 Alternatives

- Our utterances $\{u_{\diamond A}, u_{\diamond B}, u_{\diamond(A\vee B)}, u_{\diamond(A\wedge B)}\}$ are labeled with the meanings they get in the absence of *Exh*.

(7) a. You may take an apple. $u_{\diamond A}$
b. You may take a pear. $u_{\diamond B}$
c. You may take an apple or a pear. $u_{\diamond(A\vee B)}$
d. You may take an apple and a pear. $u_{\diamond(A\wedge B)}$

- However, we assume that there is uncertainty in the sense of Bergen et al. (2016) about whether and where a given utterance contains the silent operator *Exh*.
- This uncertainty stems from different LFs that are available due to optional insertion of *Exh* in the sense of Fox (2007).
  - E.g., For the utterance (7a), the following parses are available: simply $\diamond A$, or $\diamond Exh(A)$ ($\approx$ *You may take only an apple*), or $Exh \diamond (A)$ ($\approx$ *You may only take an apple*).
- The utterance (7a) = $u_{\diamond A}$ might be parsed as $\diamond A$, in which case it denotes the following set of worlds: $\{\text{Only A}, \text{Only One}, \text{Any Number}, \text{Only Both}\}$
- But it might also be parsed as $\diamond Exh(A)$ ($\approx$ *You may take only an apple*), in which case it denotes the set $\{\text{Only A, Only One, Any Number}\}$.
- Or it might be parsed as $Exh \diamond (A)$ ($\approx$ *You may only take an apple*), in which case it denotes the set $\{\text{Only A}\}$.
- (7c) = $u_{\diamond(A\vee B)}$ has the LFs $\diamond(A \vee B)$ and $\diamond Exh(A \vee B)$ (this one is equivalent to $\diamond(Exh(A) \vee Exh(B))$, and is used twice in the model).
- 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 three "lexica" ($\mathscr{L}$) in (8)-(10).

(8) $[\![u_{\diamond A}]\!]^{\mathscr{L}_1} = \{\text{Only A}, \text{Only One}, \text{Any Number}, \text{Only Both}\}$,
$[\![u_{\diamond B}]\!]^{\mathscr{L}_1} = \{\text{Only B}, \text{Only One}, \text{Any Number}, \text{Only Both}\}$,
$[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_1} = \{\text{Only A}, \text{Only B}, \text{Only One}, \text{Any Number}, \text{Only Both}\}$,
$[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_1} = \{\text{Any Number}, \text{Only Both}\}$

(9) $[\![u_{\diamond A}]\!]^{\mathscr{L}_2} = \{\text{Only A}, \text{Only One}, \text{Any Number}\}$,
$[\![u_{\diamond B}]\!]^{\mathscr{L}_2} = \{\text{Only B}, \text{Only One}, \text{Any Number}\}$,
$[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_2} = \{\text{Only A}, \text{Only B}, \text{Only One}, \text{Any Number}\}$,
$[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_2} = \{\text{Any Number}, \text{Only Both}\}$

(10) $[\![u_{\diamond A}]\!]^{\mathscr{L}_3} = \{\text{Only A}\}$,
$[\![u_{\diamond B}]\!]^{\mathscr{L}_3} = \{\text{Only B}\}$,
$[\![u_{\diamond(A\vee B)}]\!]^{\mathscr{L}_3} = \{\text{Only A}, \text{Only B}, \text{Only One}, \text{Any Number}\}$,
$[\![u_{\diamond(A\wedge B)}]\!]^{\mathscr{L}_3} = \{\text{Only Both}\}$

- 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 lexica that mix elements of (8)-(10).

4 Deriving free choice

- Deriving free choice means assigning (near-)zero probability to the worlds Only A, Only B, and Only Both upon hearing the disjunction.
- This model derives FCI for the level-1 pragmatic listener.
- We assume that the literal listener (L0) needs to be told what lexicon is at play, but the pragmatic listeners and the speakers average/marginalize over lexica.
- For uniform priors $P(w)$, lexica as above, and sufficiently large α, the level-1 pragmatic listener $P_{listener\,1}(\cdot|u_{\diamond(A\vee B)})$ splits its probability mass almost evenly between the FCI+EI world Only One and the FCI-EI world Any Number, with virtually no mass assigned to the non-FCI worlds Only A, Only B, Only Both:

```
    temperature: 100

L0 literal listener for lexicon 1:
      Only A    Only B    Any Number    Only One    Only Both
a     0.25      0         0.25          0.25        0.25
b     0         0.25      0.25          0.25        0.25
disj  0.2       0.2       0.2           0.2         0.2
conj  0         0         0.5           0           0.5

L0 literal listener for lexicon 2:
      Only A    Only B    Any Number    Only One    Only Both
a     0.33      0         0.33          0.33        0
b     0         0.33      0.33          0.33        0
disj  0.25      0.25      0.25          0.25        0
conj  0         0         0.5           0           0.5

L0 literal listener for lexicon 3:
      Only A    Only B    Any Number    Only One    Only Both
a     1         0         0             0           0
```

b	0	1	0	0	0
disj	0.25	0.25	0.25	0.25	0
conj	0	0	0	0	1

L1 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.75	0	0	0.25	0
b	0	0.75	0	0.25	0
disj	0	0	0.5	0.5	0
conj	0	0	0.4	0	0.6

L2 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0	0	0.5	0.5	0
conj	0	0	0	0	1

- Reading the tables:
  - The rows correspond to the possible utterances which the listeners take as their input.
  - The columns correspond to the world states.
  - The three literal listener outputs ($\mathscr{L}_1$-$\mathscr{L}_3$) are the outputs given the three lexica used.
  - The pragmatic listener outputs are the outputs from each recursive pragmatic listener layer.
  - For each utterance and world state, we display the probability that the given listener layer would assign to that world state given the utterance.
    - E.g., In the above output for the level-1 pragmatic listener, the probability that the listener would assign to the Only A world given the utterance "You may take an apple" (A) is 0.75.
- For lower values of α, FCI arises at higher levels of recursion.

```
temperature: 2
```

L1 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.63	0	0.1	0.22	0.04
b	0	0.63	0.1	0.22	0.04
disj	0.18	0.18	0.28	0.33	0.02
conj	0	0	0.29	0	0.71

L2 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.76	0	0.05	0.19	0

b	0	0.76	0.05	0.19	0
disj	0.07	0.07	0.38	0.48	0
conj	0	0	0.32	0	0.68

L3 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.88	0	0.01	0.11	0
b	0	0.88	0.01	0.11	0
disj	0.01	0.01	0.43	0.56	0
conj	0	0	0.29	0	0.71

L4 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.97	0	0	0.03	0
b	0	0.97	0	0.03	0
disj	0	0	0.42	0.58	0
conj	0	0	0.24	0	0.76

L5 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0	0	0.43	0.57	0
conj	0	0	0.2	0	0.8

4.1 Exclusivity Inference vs Free Choice Inference

- Now let's turn to the exclusivity inference (EI, "not both"). This corresponds to hearing the disjunction and assigning a low probability to the Any Number world.
- Recall that this inference is not as strong as FCI. This means that it should be possible to derive free choice and still assign a high probability to the Any Number world.
- We suggest that this happens, e.g., when there are plenty of fruit on the table and so it is a priori likely that one has permission to take more than one. In such situation, EI does not arise.
- For nonuniform priors $P(w)$ that assign Any Number a high prior, $P_{listener\,1}$ and upwards also assign it a high posterior upon hearing the disjunction $u_{\diamond(A\vee B)}$.

```
"worlds":        [w.a, w.b, w.any_number, w.only_one, w.only_both],
"probabilities": [.1,  .1,    .6,           .1,          .1]

  temperature: 100
```

L0 literal listener for lexicon 1:

	Only A	Only B	Any Number	Only One	Only Both
a	0.11	0	0.67	0.11	0.11

```
b     0         0.11      0.67          0.11         0.11
disj  0.1       0.1       0.6           0.1          0.1
conj  0         0         0.86          0            0.14
L0 literal listener for lexicon 2:
      Only A    Only B    Any Number    Only One     Only Both
a     0.13      0         0.75          0.13         0
b     0         0.13      0.75          0.13         0
disj  0.11      0.11      0.67          0.11         0
conj  0         0         0.86          0            0.14
L0 literal listener for lexicon 3:
      Only A    Only B    Any Number    Only One     Only Both
a     1         0         0             0            0
b     0         1         0             0            0
disj  0.11      0.11      0.67          0.11         0
conj  0         0         0             0            1
L1 pragmatic listener
      Only A    Only B    Any Number    Only One     Only Both
a     0.75      0         0             0.25         0
b     0         0.75      0             0.25         0
disj  0         0         0.86          0.14         0
conj  0         0         0.8           0            0.2
L2 pragmatic listener
      Only A    Only B    Any Number    Only One     Only Both
a     0.67      0         0             0.33         0
b     0         0.67      0             0.33         0
disj  0         0         1             0            0
conj  0         0         0.01          0            0.99
```

- Likewise, a high prior on the Only One world will lead to EI at L2 (i.e. a low posterior on the Any Number world). This could be a scenario where there are only two fruit on the table and it is a priori impolite to take both.

```
"worlds":        [w.a, w.b, w.any_number, w.only_one, w.only_both],
"probabilities": [.1,  .1,    .1,            .6,           .1]
```

```
temperature: 100
L0 literal listener for lexicon 1:
      Only A    Only B    Any Number    Only One     Only Both
a     0.11      0         0.11          0.67         0.11
b     0         0.11      0.11          0.67         0.11
disj  0.1       0.1       0.1           0.6          0.1
conj  0         0         0.5           0            0.5
L0 literal listener for lexicon 2:
```

	Only A	Only B	Any Number	Only One	Only Both
a	0.13	0	0.13	0.75	0
b	0	0.13	0.13	0.75	0
disj	0.11	0.11	0.11	0.67	0
conj	0	0	0.5	0	0.5

L0 literal listener for lexicon 3:

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0.11	0.11	0.11	0.67	0
conj	0	0	0	0	1

L1 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	0.33	0	0	0.67	0
b	0	0.33	0	0.67	0
disj	0	0	0.14	0.86	0
conj	0	0	0.4	0	0.6

L2 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0	0	0	1	0
conj	0	0	0.5	0	0.5

- This shows we can derive the optionality of EI as a matter of prior knowledge, at least when using just (8)-(10).
- Crucially, 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.
  - Assigning a high prior probability ton the Only A world does not lead to a high posterior; instead, leads to the same L2 lis the same astener as it was ifor the uniform priorprior was uniform:

```
"worlds":        [w.a, w.b, w.any_number, w.only_one, w.only_both],
"probabilities": [.6,  .1,    .1,           .1,           .1]

  temperature: 100
```

L0 literal listener for lexicon 1:

	Only A	Only B	Any Number	Only One	Only Both
a	0.67	0	0.11	0.11	0.11
b	0	0.25	0.25	0.25	0.25
disj	0.6	0.1	0.1	0.1	0.1
conj	0	0	0.5	0	0.5

L0 literal listener for lexicon 2:

	Only A	Only B	Any Number	Only One	Only Both
a	0.75	0	0.13	0.13	0
b	0	0.33	0.33	0.33	0
disj	0.67	0.11	0.11	0.11	0
conj	0	0	0.5	0	0.5

L0 literal listener for lexicon 3:

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0.67	0.11	0.11	0.11	0
conj	0	0	0	0	1

L1 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	0.6	0	0.4	0
disj	0	0	0.5	0.5	0
conj	0	0	0.4	0	0.6

L2 pragmatic listener

	Only A	Only B	Any Number	Only One	Only Both
a	1	0	0	0	0
b	0	1	0	0	0
disj	0	0	0.5	0.5	0
conj	0	0	0	0	1

  - Assigning high prior probability to the Only B world (or to the Only Both world) also leads to this L2 listener. (We do not show this here.)

- This explains why FCI is a stronger inference than EI.
  - FCI does not show the same optionality as EI when the probabilities over world priors are shifted.

4.2 No free choice without lexical uncertainty

- Lexical uncertainty is crucial for modeling FCI in RSA. With justOnly oOne lexicon for interpretation, the model does not derive FCI.

- For example, aE.g. mModel with just the unexhaustified lexicon $\mathscr{L}_1$ places most of the posterior probability on non-FCI Only A and Only B words upon hearing $u_{\diamond(A \vee B)}$:

```
temperature: 100
```

L0 literal listener:

	Only A	Only B	Any Number	Only One	Only Both
a	0.25	0	0.25	0.25	0.25

b	0	0.25	0.25	0.25	0.25
disj	0.2	0.2	0.2	0.2	0.2
conj	0	0	0.5	0	0.5

```
L1 pragmatic listener (and upwards)
```

	Only A	Only B	Any Number	Only One	Only Both
a	0.67	0	0	0.33	0
b	0	0.67	0	0.33	0
disj	0.4	0.4	0	0.2	0
conj	0	0	0.5	0	0.5

5 No free choice under negation

- Our model captures the absence of FCI under negation under plausible assumptions about the lexical meanings of the utterances involved.

(11) You may not take an apple or a pear.

- 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).
- Specifically, we assume that one of the LFs for (11) is equivalent to the classical $\neg\diamond(A \vee B)$, with no *Exh* inserted. This semantically entails $\neg\diamond A$ and $\neg\diamond B$.
- Without lexical uncertainty, we derive absence of free choice under negation, unsuprisingly.
- We assume the following world states: $\{\textsf{Only A}, \textsf{Only B}, \textsf{Only One}, \textsf{Neither}, \textsf{Only Both}\}$.
- In the Neither world, nothing is allowed.
- Semantically, $\neg\diamond(A \vee B)$ denotes $\{$ Neither $\}$.
- The world Any Number is incompatible with $\neg\diamond(A \vee B)$ and all of its alternatives, so we don't consider it.
- Free choice under negation would obtain if L1 or upwards interpreted $\neg\diamond(A \vee B)$ as $\neg(\diamond A \wedge \diamond B)$, that is, as $\{\textsf{Only A}, \textsf{Only B}\}$.
- But they don't:

```
    temperature: 100

L0 literal listener
```

	Only A	Only B	Neither	Only One
not-a	0	0.5	0.5	0
not-b	0.5	0	0.5	0
not-or	0	0	1	0

```
not-and  0.25      0.25      0.25       0.25
L1 pragmatic listener (and upwards)
         Only A    Only B    Neither    Only One
not-a    0         1         0          0
not-b    1         0         0          0
not-or   0         0         1          0
not-and  0         0         0          1
```

- We have assumed here that the space of world states is limited to those in which the non-exhaustified utterance or one of its alternatives is true.
- So we exclude Any Number and Only Both from consideration.
- Given this, the presence of uncertainty (i.e. the availability of other, weaker LFs which contain *Exh*) doesn't affect this result.
- For example, the LF $Exh\neg\diamond(A\vee B)$, we assume that $\neg\diamond(A\vee B)$ has the alternatives $\neg\diamond A, \neg\diamond B, \neg\diamond(A\wedge B)$ (all of which it entails). In that case this LF is equivalent to its non-exhaustified version, $\neg\diamond(A\vee B)$, which denotes {Neither}.
- For the LF $\neg Exh(\diamond(A\vee B))$, we assume that $\diamond(A\vee B)$ has the alternatives $\diamond A$, $\diamond B$, and $\diamond(A\wedge B)$. In that case this LF is equivalent to $\neg[\diamond(A\vee B)\wedge\neg\diamond(A\wedge B)]$, which is in turn equivalent to $(\neg\diamond(A\vee B))\vee\diamond(A\wedge B)$, i.e. again {Neither} (Recall that we exclude Any Number and Only Both from consideration.)
- As for the LF $\neg\diamond(Exh(A)\vee Exh(B))$, it is presumably equivalent to $\neg\diamond((A\wedge\neg B)\vee(B\wedge\neg A))$, i.e. {Neither} (recall we do not include Only Both).
- So it doesn't matter if and how disjunction is exhaustified.
- Still, there is uncertainty as to whether $\neg\diamond A$ rules out Only One or not.
- Within the states in play, $\neg\diamond A$ and $Exh\neg\diamond A$ are true at Only B and Neither.
- $\neg Exh\diamond(A)$ is true at Only B, Neither, and Only One.
- It turns out that from the level-2 listener upwards, this uncertainty has no effect:

```
temperature: 100

L0 literal listener for lexicon 1:
         Only A    Only B    Neither    Only One
not-a    0         0.5       0.5        0
not-b    0.5       0         0.5        0
not-or   0         0         1          0
not-and  0.25      0.25      0.25       0.25
L0 literal listener for lexicon 2:
         Only A    Only B    Neither    Only One
```

not-a	0	0.33	0.33	0.33
not-b	0.33	0	0.33	0.33
not-or	0	0	1	0
not-and	0.25	0.25	0.25	0.25

L1 pragmatic listener

	Only A	Only B	Neither	Only One
not-a	0	0.8	0	0.2
not-b	0.8	0	0	0.2
not-or	0	0	1	0
not-and	0	0	0	1

L2 pragmatic listener

	Only A	Only B	Neither	Only One
not-a	0	1	0	0
not-b	1	0	0	0
not-or	0	0	1	0
not-and	0	0	0	1

6 Conclusion

- We have derived the free choice inference and the exclusivity inference within the framework of the RSA model.
  - We show, using uncertainty over LFs, that RSA can derive free choice.
  - The exclusivity inference arises as a result of changes in the prior probabilities over worlds.
- We also demonstrate in our model that free choice does not arise under negation.
- Further steps: clarify the selection of the space of world states (see Franke 2011); consider additional states; consider scopal ambiguity of connectives wrt modals as an additional source of uncertainty.

7 Further Work

- Clarify the selection of the space of world states (see Franke 2011).
- Consider additional states (e.g. you may take the pear only if you take the apple).
- Add states in which the speaker is uncertain about the nature of the permissions in question (as in Franke 2011's epistemic models)
- Consider scopal ambiguity of connectives wrt modals as an additional source of uncertainty.
- Adding states in which the speaker is uncertain about the nature of the permissions in question (as in Franke 2011's epistemic models)

- Alex Warstadt, p.c.: “He is rich enough to buy a yacht or a plane.” (FC arises, EI is strong.) vs. “He is tall enough to ride the Cyclone or the Thunderbolt” (FC arises, EI does not. Height is not a consumable resource in the way that money is). These are not DE. (Look at Zimmermann 2000 under epistemic models, and Dan Harris’s NYPLW talk. Similar behavior with ‘want’.)

- How might we model a L2 speaker or child or an otherwise cognitively different participant in conversation?

- “You don’t have to clean your room and wash the dishes.” (Fox)

- Do we still get FCI in the absence of conjunction as an utterance alternative?

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.

Zimmermann, T. E. (2000). Free choice disjunction and epistemic possibility. *Natural Language Semantics*, 8(4):255–290.