Proceedings of SALT 29: 238–257, 2019

Free choice disjunction as a rational speech act*

Lucas Champollion
New York University

Anna Alsop
New York University

Ioana Grosu
New York University

Abstract The so-called free choice inference (from *You may take an apple or a pear* to *You may take an apple*) is mysterious because it does not follow from ordinary modal logic. We show that this inference arises in the Rational Speech Act framework (Frank & Goodman 2012). Our basic idea is inspired by exhaustification-based models of free choice (Fox 2007) and by game-theoretic accounts based on iterated best response (Franke 2011). We assume that when the speaker utters *You may take an apple or a pear*, the hearer reasons about why the speaker did not choose alternative utterances such as *You may take an apple*. A crucial ingredient in our explanation is the idea of semantic uncertainty (Bergen, Levy & Goodman 2016). Specifically, we assume that the speaker is uncertain whether or not the hearer will interpret *You may take an apple* as forbidding them from taking a pear. This uncertainty can be thought of as resulting from Fox's (2007) optional covert exhaustification. Uttering the disjunction is a way for the speaker to prevent the hearer from concluding that any fruit is forbidden to take. Knowing this, the hearer concludes that they may choose either fruit.

Keywords: free choice, disjunction, Gricean reasoning, Bayesian inference, game theory, RSA, quantity implicature, utterance ambiguity

1 Introduction

When a disjunction is embedded under a modal as in (1a), it conveys (1b) and (1c):

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

b. $\rightsquigarrow$ You may take an apple (by itself). $\diamond(A \wedge \neg B)$

c. $\rightsquigarrow$ You may take a pear (by itself). $\diamond(B \wedge \neg A)$

* We thank Chris Barker, Leon Bergen, Alexandre Cremers, Salvador Mascarenhas, Chris Potts, Ciyang Qing, Jacopo Romoli, Philippe Schlenker, Benjamin Spector, Anna Szabolcsi, and audiences at NYU, ENS Paris, and SALT.

This phenomenon, which we will call *free choice inference (FCI)*, is unexpected given standard modal logic. In fact, as is well known, not even $\diamond A$ follows from $\diamond(A \vee B)$, nor does $\diamond B$ (Kamp 1974).

(1a) is also taken to generally convey (2), which we call the *exclusivity inference (EI)*. Given EI, $\diamond(A \wedge \neg B)$ is equivalent to $\diamond A$, which is how (1b) is often displayed.

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

EI appears to be cancelable, in contrast to FCI (Simons 2005). For example, canceling FCI in (3a) is infelicitous, while canceling EI in (3b) is acceptable.

(3) You may take an apple or a pear. . . .

a. . . . #In fact, you may not take an apple { ∅ / by itself }.

b. . . . In fact, you may take both.

FCI tends to disappear in downward-entailing environments such as negation. For example, (4a) is typically interpreted roughly as (4b), which is equivalent to $\neg\diamond(A \vee B)$ via de Morgan's law. The negation of FCI in (4c), which would be weaker than (4b) and compatible with $\diamond(A \wedge B)$, is not an available reading of (4a).

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

b. $\approx$ You may not take an apple and you may not take a pear. $\neg\diamond A \wedge \neg\diamond B$

c. $\not\approx$ You do not get to choose. $\neg[\diamond(A \wedge \neg B) \wedge \diamond(B \wedge \neg A)]$

This paper develops a theory of free choice disjunction that accounts for each of these observations. In particular, we provide an explicit mechanism that derives both FCI and EI in a unified way. We also explain the difference in cancelability of FCI vs. EI, and the disappearance of FCI in downward-entailing environments.

Broadly speaking, accounts of free choice fall into two classes (see Meyer to appear for an overview). Semantic accounts attribute FCI to a nonstandard semantics of either disjunction or the modal (e.g. Zimmermann 2000; Aloni 2007; Simons 2005). Pragmatic accounts derive FCI by invoking general mechanisms of implicature computation. This paper focuses on pragmatic accounts and sets aside semantic ones, which typically struggle to explain why FCI is absent under negation.

Pragmatic accounts can be further subcategorized into grammatical and Neo-Gricean accounts, corresponding to the current divide among pragmatic theories (Schlenker 2016). Grammatical accounts derive scalar implicatures by an exhaustification operator *Exh*, which can be inserted into the LF (logical form) at various places. For example, Fox (2007) derives FCI through the recursive application of *Exh* to (1a). Fox additionally explains the optionality of EI by deriving the weaker reading (1b) using a different LF that applies *Exh* to each of the disjuncts.

Neo-Gricean accounts derive implicatures as well as FCI by relying on Gricean reasoning. We focus here on game-theoretic formalizations of Gricean reasoning. Franke (2011) presents such an account, formalizing the pragmatics of free choice disjunction within the iterated best response (IBR) model. In IBR, a listener's goal is to determine what the state of the world is, and a speaker's goal is to choose an utterance which will help the listener do so. Speakers and listeners reason recursively about each other's behavior and take each other to behave optimally, e.g. speakers will always select the most informative true utterances. With all states equally likely *a priori*, an utterance's informativity is higher whenever it rules out more states.

Each of these accounts faces challenges. While recursive exhaustification is able to derive FCI, Fox's (2007) formulation of the pragmatics does not fully explain why FCI arises reliably. This is because Fox generates many LFs for an utterance like (1a), of which many lack FCI. Free choice is therefore derived only to the extent that hearers rule out such LFs, and Fox's account is arguably not complete without an explanation of why they do so (Alsop, Champollion & Grosu to appear).

Turning to game-theoretic accounts such as Franke (2011), the challenge is that the disjunction is always less informative than the disjuncts as far as literal meaning is concerned. Given the preference for more informative utterances, a listener would not expect a speaker to ever utter the disjunction. In IBR as described so far, a listener who hears the disjunction anyway lacks the resources to interpret it. Franke (2011) stipulates that hearers interpret such *surprise messages* by falling back on their literal interpretations. As we will see, this stipulation turns out to be crucial.

We propose an account of free choice disjunction that incorporates components from both Fox (2007) and Franke (2011). Our hybrid account derives FCI and EI, while also explaining the asymmetry in cancelability between these two inferences, as well as the disappearance of FCI under negation. Crucially, we include multiple literal interpretation functions. To preview our account, the speaker will be uncertain whether the hearer would interpret *You may take a pear* as forbidding them from taking an apple. This reflects the fact that many utterances are semantically ambiguous and that hearers do not know *a priori* which interpretation speakers have in mind. One may take exhaustification-based accounts such as Fox (2007) to characterize the source of such ambiguities. To model the process by which speakers and listeners reason about this uncertainty, we use the Rational Speech Act (RSA) framework (Frank & Goodman 2012; Bergen et al. 2016), which is similar to IBR but avoids the problems surrounding surprise messages. We also explain the asymmetry in cancelability between FCI and EI by exploiting the varying sensitivity of RSA to prior probabilities. In this respect, we go beyond Fox (2007) and Franke (2011), neither of which sets out to account for this asymmetry.

Our paper is structured as follows. We start by criticially reviewing Franke (2011) in Section 2. In Section 3 we describe our model, its key assumptions, and its predictions. Section 4 shows that our model predicts the disappearance of FCI under negation. Section 5 concludes and points out avenues for further research.

2 Free Choice in IBR

We start by reviewing the four-state, four-utterance model of free choice disjunction in IBR proposed in Franke (2011: 47). We choose this as opposed to Franke's other models of free choice because it takes conjunction into account and sets aside speaker uncertainty; we will make analogous choices in our own model.

In IBR and RSA, an utterance is interpreted against the backdrop of a predetermined set of alternatives, the *utterance set*, which is assumed to be mutually known. Franke's four-utterance model uses the following set, which we adopt as well:

(5) a. You may take an apple. $\diamond A$ — A
b. You may take a pear. $\diamond B$ — B
c. You may take an apple or a pear. $\diamond(A \vee B)$ — Or
d. You may take an apple and a pear. $\diamond(A \wedge B)$ — And

A listener's goal is to determine what *state* the world is in. A state can be thought of as something that contains enough information to determine, for each member of the utterance set, whether it is true. For example, Franke models free choice disjunction by interpreting the alternatives in his utterance set in (5) according to classical modal logic (as indicated by the formulas) and then constructing all the states not ruled out by any utterance. Due to the entailment relations between the alternatives, this leads him to four states, which we will refer to as Only A, Only B, Only One, and All True. In Only A, the listener only has permission to take an apple, but not a pear; in Only B, it is the other way around. In Only One, the listener may choose an apple or a pear, but not both (that is, both FCI and EI hold). At the state we have labeled All True, all utterances in (5) are true (we will come back to this state below). Table 1, adapted from Franke (2011: 47), displays the state space along with the truth values of the utterances under consideration and the status of FCI and EI at each state as applicable (we return to All True in Section 3.1).[1]

By default, hearers in both IBR and RSA are assumed to take each state to be equally likely *a priori*, though in Section 3 we will manipulate this parameter. A level-0 listener (L0) interprets every utterance literally, in our case according to Table 1. If the utterance is true at two or more states, L0 randomly samples among them

1 To remain consistent with our subsequent terminology and with that of the RSA literature, we have renamed Franke's listeners, speakers, and states. He uses the sequence R_0, S_1, R_2, S_3, R_4 etc. where we have L_0, S_1, L_1, S_2, L_2, etc. Our states Only A, Only B, Only One, and All True correspond to his states $t_{[1,0,0,1]}$, $t_{[0,1,0,1]}$, $t_{[1,1,0,1]}$, and $t_{[1,1,1,1]}$, in that order. We have also renamed his utterances.

	A	B	Or	And	Does FCI hold?	Does EI hold?
Only A	T	F	T	F	no	yes
Only B	F	T	T	F	no	yes
Only One	T	T	T	F	yes	yes
All True	T	T	T	T	?	no

Table 1: Truth table for utterances, using states from Franke (2011) (state and utterance labels our own)

$$\underset{\text{Level 0}}{\text{L0}} \longrightarrow \underbrace{\text{S1} \longrightarrow \text{L1}}_{\text{Level 1}} \longrightarrow \underbrace{\text{S2} \longrightarrow \text{L2}}_{\text{Level 2}} \quad \cdots$$

Figure 1: Conversational agents in the IBR and RSA models

according to their prior beliefs. A speaker chooses an utterance from the utterance set with the goal of helping the listener determine the actual state—be it by conveying what that state is, or at least by ruling out as many nonactual states as possible.

Since the modal logic formula corresponding to (5c) lacks the FCI and EI inferences, that utterance is true at all four states. These inferences will arise when speakers and listeners at higher levels reason recursively about each other's behavior, which leads them to enrich the meanings of utterances in various ways. Enriched meanings can capture scalar implicatures and the like, and will be used here to account for FCI and EI. This is where the explanatory value of IBR and RSA lies. A level-1 speaker (S1) expects to be facing L0, and chooses an utterance accordingly. A level-1 listener (L1) expects S1 and takes this into account while determining the most likely explanation for the speaker's utterance. To this end, L1 rules out not only those states at which the utterance is false, but also those at which S1 would have been more likely to choose another utterance. In this way, Gricean reasoning is formalized as Bayesian inference. IBR and RSA also allow for higher-level agents, who reason analogously about each other (see Figure 1). While Franke (2011) crucially relies on these higher levels, the model we present in this paper does not.

In IBR, listeners and speakers are assumed to always make the choice that is optimal from their point of view (this assumption is the chief distinction between IBR and RSA). Given this, Franke's model runs into an issue: the disjunction is a *surprise message* for L1, i.e. a message that L1 never expects S1 to choose. The following explanation is based on Franke (2011: see Fig. 12 and Appendix B.3).

The reason S1 would never choose *You may A or B* is that there is no state for which it is an optimal choice. To convey Only A, of the true utterances S1 prefers

You may A, since this rules out the nonactual state Only B while *You may A or B* would not rule out any state. Similarly for Only B. To convey Only One, S1 should utter either *You may A* or *You may B* because both allow L0 to rule out one nonactual state and are thus more informative than *You may A or B*. And to convey All True, the best choice is *You may A and B* because it rules out all nonactual states.

Franke stipulates that any listener who hears a surprise message falls back to L0 behavior and interprets it literally. This means that, just as for L0, for L1 *You may A or B* is the only message that does not rule out any states. In this sense, for L1 the disjunction is not only surprising but useless.

How should a speaker act who thinks of the listener as L1? That is, how should S2 act in each state? Consider first Only A, Only B, and All True. As we have seen, L1 expects each of these states to trigger a specific utterance; so in each case the best choice for S2 is to choose that utterance. For example, S2 should use *You may A* to convey Only A even though L1 knows it is also compatible with Only One. This is because L1 thinks of the speaker as S1. As we have seen, if S1 is in Only A, then *You may A* was their only optimal choice, whereas in Only One it was one of two optimal choices since S1 could just as well have chosen *You may B*. Thus *You may A* is a more reliable indicator for Only A than it is for Only One. Assuming equal priors, L1 therefore ranks Only A above Only One as more likely given this utterance. Here the Bayesian character of IBR can be observed particularly clearly.

Now suppose S2 wants to convey Only One. The best choice is *You may A or B*, since any other message would lead L1 to zero in on another state. While for L1 the disjunction will be useless, at least it will not rule out the actual state.

To derive free choice, Franke (2011) assumes that the listener is not L1 but L2, i.e. someone who takes the speaker to be S2. If S2 utters *You may A or B*, the intended state must be Only One, otherwise S2 would have chosen another message. FCI arises because Only A and Only B are ruled out, and EI arises because All True is ruled out. So for L2, unlike for L0 and L1, *You may A or B* is not useless after all.

To sum up, in the IBR model, *You may A or B* is a surprise message for L1, and the free choice and exclusivity inferences arise only at the level of L2. In other words, free choice arises based on two assumptions: first, that the speaker takes the listener to never expect them to utter a disjunction, and takes them to interpret unexpected messages literally; and second, that the listener takes the speaker to behave this way. These assumptions strike us as conceptually implausible.

To be sure, these assumptions are not inherent in IBR. As it turns out, one can derive free choice within IBR while avoiding surprise messages and without resorting to the L2 level; but this comes at a cost. As Franke (2011: 48-50) shows, the tradeoff involves assuming that for all the listener knows, the speaker might be partially ignorant about the true state of the world (see also Jäger (2011) for a similar model). This also strikes us as an implausible assumption, as free choice inferences can arise even when it is mutually known that the speaker is fully knowledgeable. In what follows, we use only the L1 level and a knowledgeable speaker.

3 Modeling the Basic Free Choice Effect in RSA

We now turn to the RSA framework, which is similar to IBR in many respects. It models communication as recursive reasoning, with the same set of conversational agents schematized in Fig. 1. One key feature of the RSA framework which differentiates it from IBR is that it does not assume that speakers always make optimal choices in the game. Recall that the shared goal of the speaker and the listener is to help the listener determine the state of the world. In RSA, speakers tend towards rational behavior, but occasionally still choose sub-optimal utterances that do not maximize a listener's odds of arriving at the actual state. Formally, speakers are modeled using a soft-max rather than hard-max function. A speaker's behavior can be adjusted using the optimality parameter $\alpha > 0$; increasing α increases the probability that a speaker will choose the optimal utterance(s). For example, whereas in IBR S1 would never use *You may A or B*, in RSA S1 will sometimes do so, even though in every state there is at least one message that S1 is more likely to pick. Therefore the disjunction is no longer a surprise message for L1. Even so, a speaker who faces L1 (i.e., S2) will tend to avoid using it as they would in IBR. This is because no matter the state, at least one other utterance will rule out more states.

This means that a vanilla RSA framework does not derive free choice. The classical pragmatic account of free choice in Kratzer & Shimoyama (2002) ("if the speaker had said *You may A*, this would have implicated *You may not B*, and vice versa; since the speaker chose not to use these utterances, it must be that these implicatures are false") cannot be implemented. This is because no matter what S2 says, L1 will never consider Only One to be the most likely state. All S2 can achieve is to make Only One only somewhat unlikely (by using a disjunct), as opposed to very unlikely (by using the disjunction). So there is no reason for S2 to use the disjunction; but due to the soft-max function, S2 will sometimes use it anyway, and there are no surprise messages for L2. Even at higher levels, the system never recovers from this predicament. Varying the value of α does not affect this pattern.

The core issue is that for L0, disjunction is less specific and thus leads to more uncertainty than any other message. We address this imbalance by increasing the uncertainty associated with the other messages. We do so by assuming that each of the other utterances is semantically ambiguous. For consistency, we extend this assumption to disjunction itself, although this will not be crucial in our account.

In the following sections, we lay out the specifications for an RSA model that meets the challenges we have described and derives both FCI and EI. The key ingredient we add is semantic ambiguity, which we implement by defining multiple

	A	B	Or	And	Does FCI hold?	Does EI hold?
Only A	T	F	T	F	no	yes
Only B	F	T	T	F	no	yes
Only One	T	T	T	F	yes	yes
Any Number	T	T	T	T	yes	no
Only Both	T	T	T	T	no	no

Table 2: Truth table for utterances, using our expanded state space

interpretation functions over which our conversational agents then reason.

3.1 Utterances and states

We begin by assuming the set of four utterances introduced in (5). These are a standard set of pragmatic alternatives for disjunction assumed by Gazdar (1979), Sauerland (2004), and Fox (2007), among others. As in Franke (2011), we ignore implicatures from *may* to *not must*; accordingly, we do not include the scalar alternative *must* even though it is arguably evoked by the modal *may*.

The standard modal logic denotations of our utterances give rise to the four states shown in Table 1; this is also the state space assumed by Franke. For deriving basic FCI, this state space is sufficient. All we need to show is that upon hearing Or, a listener in our model concludes that Only One is the most likely state.

However, this state space is arguably not granular enough to represent a state of affairs in which FCI holds but EI does not. This is why we included a question mark in Table 1. Suppose that the utterances in (5) are interpreted as in classical modal logic, as $\diamond A$, $\diamond B$, $\diamond(A \vee B)$, and $\diamond(A \wedge B)$, and that the All True state is fully characterized by the fact that it makes each of them true. This rules out EI but does not specify whether FCI holds as we have defined it. It might be that you can take either fruit by itself or both, in which case FCI holds. Or it might be that you are only allowed to take both fruit together, as a package deal (van Rooij 2000). Then even though $\diamond A$ and $\diamond B$ are true, $\diamond(A \wedge \neg B)$ and $\diamond(B \wedge \neg A)$ are false, and FCI does not hold.

We therefore replace the All True state in Table 1 by two new states, which we call Any Number and Only Both. In Any Number, we may choose either fruit by itself or both, i.e. $\diamond(A \wedge \neg B) \wedge \diamond(B \wedge \neg A) \wedge \diamond(A \wedge B)$. This encodes FCI without EI. In Only Both, we may not choose either fruit by itself, but choosing both is permitted, i.e. $\diamond(A \wedge B) \wedge \neg \diamond (A \wedge \neg B) \wedge \neg \diamond (B \wedge \neg A)$. Here we have neither FCI nor EI.[2]

2 While Any Number and Only Both are subsets of All True, they do not exhaust it. What we have left out includes, for example, the state in which A is allowed, both A and B are allowed together, but

We follow Franke (2011) in ignoring states at which no utterance is true. For example, we do not include a state in which the listener has no permissions. In such a state, every utterance in our set is false; therefore, it is left out of our state space.

3.2 The Key Ingredient: Semantic Ambiguity

As mentioned, the Franke (2011) model ported directly into RSA does not derive free choice. Our newly expanded state space leads to the same undesirable result.

We will remedy this by adding the crucial ingredient of semantic ambiguity to our free choice model. Bergen et al. (2016) assumes that listeners are uncertain about what *lexicon*, or set of word-meaning mappings, a speaker has in mind. A pragmatic listener must reason over possible lexica in addition to utterances and states. Here we recast this idea as a form of uncertainty over interpretation functions, in which one utterance may be mapped to multiple propositions.

Incorporating semantic ambiguity brings our account closer to grammatical theories of scalar implicature along the lines of Fox (2007). The optional insertion of exhaustivity operators in that work makes multiple LFs available for a single utterance, which a listener must then reason about. Fox (2007) invokes the notion of an economical parser to explain why a listener chooses certain LFs over others; our model simply folds this reasoning process into the pragmatics.

Here, we provide a minimal model with two interpretation functions to illustrate how RSA augmented with semantic ambiguity derives free choice. The two interpretation functions ($\mathscr{I}$) in (6) and (7) map each utterance to a proposition. $\mathscr{I}_1$ corresponds to the classical modal logic denotations of each utterance. In $\mathscr{I}_2$, the meanings of each utterance have been strengthened. While our utterance interpretations in the latter are stipulated, each can be understood as stemming from covert exhaustification following the innocent exclusion algorithm in Fox (2007).[3]

(6) Interpretation Function 1

a. $[\![\mathsf{A}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ A}, \mathsf{Only\ One}, \mathsf{Any\ Number}, \mathsf{Only\ Both}\}$

b. $[\![\mathsf{B}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Any\ Number}, \mathsf{Only\ Both}\}$

c. $[\![\mathsf{Or}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Any\ Number}, \mathsf{Only\ Both}\}$

d. $[\![\mathsf{And}]\!]^{\mathscr{I}_1} = \{\mathsf{Any\ Number}, \mathsf{Only\ Both}\}$

B is not allowed on its own. This asymmetrical set of permissions may seem puzzling in the fruit scenario, but it is natural in cases where one action is more desirable than the other. A parent may allow a child to take a dessert (B) only if she takes a vegetable (A) with it, for example. We have left out this type of state because it is not crucial to the free choice inferences at hand.

3 $[\![\mathsf{A}]\!]^{\mathscr{I}_2} = \mathrm{Exh}_C(\diamond A) = \diamond A \wedge \neg\diamond B$; $[\![\mathsf{B}]\!]^{\mathscr{I}_2} = \mathrm{Exh}_C(\diamond B) = \diamond B \wedge \neg\diamond A$; $[\![\mathsf{Or}]\!]^{\mathscr{I}_2} = \diamond \mathrm{Exh}_C(A \vee B) = \diamond((A \vee B) \wedge \neg(A \wedge B))$; $[\![\mathsf{And}]\!]^{\mathscr{I}_2} = \mathrm{Exh}_{C'}(\diamond \mathrm{Exh}_C(A \wedge B)) = \mathrm{Exh}_{C'}(\diamond(A \wedge B)) = \diamond(A \wedge B) \wedge \neg\diamond(A \wedge \neg B) \wedge \neg\diamond(B \wedge \neg A) \wedge \neg\diamond((A \vee B) \wedge \neg(A \wedge B)) = \diamond(A \wedge B) \wedge \neg\diamond((A \vee B) \wedge \neg(A \wedge B))$.

	Only A	Only B	Only One	Any Number	Only Both
A	0.25	0	0.25	0.25	0.25
B	0	0.25	0.25	0.25	0.25
Or	0.2	0.2	0.2	0.2	0.2
And	0	0	0	0.5	0.5

Table 3: L0 given $\mathscr{I}_1$ in (6)

	Only A	Only B	Only One	Any Number	Only Both
A	1	0	0	0	0
B	0	1	0	0	0
Or	0.25	0.25	0.25	0.25	0
And	0	0	0	0	1

Table 4: L0 given $\mathscr{I}_2$ in (7)

(7) Interpretation Function 2

a. $[\![\mathsf{A}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ A}\}$

b. $[\![\mathsf{B}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ B}\}$

c. $[\![\mathsf{Or}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Any\ Number}\}$

d. $[\![\mathsf{And}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ Both}\}$

As in Bergen et al. (2016), L0 and S1 are fixed-lexicon agents and must be provided with an utterance and an interpretation function; they do not reason over multiple interpretation functions. The probabilities that L0 assigns to each state for each interpretation function upon hearing an utterance are outlined in Tables 3 and 4. L0 interprets each utterance by sampling a state according to the probabilities in the row determined by the utterance. At higher levels of recursion, speakers and listeners are not explicitly provided with an interpretation function. These pragmatic actors in the model instead reason over interpretation functions, jointly with utterances and states. Our model, with states w, utterances u, interpretation functions $\mathscr{I}$, levels of recursion n, and optimality parameter α, consists of the probability functions given in (8). Here, $\mathscr{I}(u,w) = 1$ if $w \in [\![u]\!]^{\mathscr{I}}$, and 0 otherwise.

(8) a. $P_{listener\,0}(w|u,\mathscr{I}) \propto \mathscr{I}(u,w)P(w)$

b. $P_{speaker\,1}(u|w,\mathscr{I}) \propto \left[P_{listener\,0}(w|u,\mathscr{I})\right]^{\alpha}$

	Only A	Only B	Only One	Any Number	Only Both
A	0.8	0	0.2	0	0
B	0	0.8	0.2	0	0
Or	0	0	0.5	0.5	0
And	0	0	0	0.33	0.67

Table 5: L1 with uniform state prior, $\alpha = 100$.

c. $P_{listener\,1}(w|u) \propto P(w) \sum_{\mathscr{I}} P_{speaker\,1}(u|w, \mathscr{I})$

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

e. $P_{listener\,n}(w|u) \propto P(w) P_{speaker\,n}(u|w)$

The behavior of S1 that these equations describe can be read off any table for L0 as follows. Within the column determined by the state S1 wants to communicate, S1 first re-weights the numbers according to the parameter α, thereby accentuating or attenuating differences between them. After renormalizing these numbers so they sum up to 1, S1 then samples an utterance accordingly.

3.3 Model Predictions

In our state space, the listener has free choice in the Only One and Any Number states but not in the other states. Therefore, in the context of our model, deriving FCI means assigning essentially no posterior probability given Or to those other states. Assuming uniform prior probabilities over states, interpretation functions as defined in (6) and (7), and sufficiently high α, L1 derives FCI as expected. As shown in Table 5, given Or, L1 splits posterior probability evenly between the FCI states Only One and Any Number, with virtually none assigned to the non-FCI states Only A, Only B, and Only Both. This is because the uncertainty associated with the A and B utterances makes S1 more reluctant to use them to convey any states other than Only A and Only B; that is, if L0 chooses the wrong interpretation function these utterances will be false in such states. In the Only A and Only B states this is not a problem because they are not ruled out under either interpretation function. This makes disjunction a reasonable choice in the Only One and Any Number states again. Essentially, in these two states uttering the disjunction rather than, say, *You may A* allows S1 to avoid the risk that the hearer might interpret *You may A* as Only A. L1 is aware of this and rules out Only A (as well as Only B for analogous reasons).

For lower values of α, the conversational agents are less optimal in their behavior. They cannot rely on each other as much, and as might be expected, FCI becomes

	Only A	Only B	Only One	Any Number	Only Both
A	0.8	0	0.2	0	0
B	0	0.8	0.2	0	0
Or	0	0	0.08	0.92	0
And	0	0	0	0.86	0.14

Table 6: L1 with asymmetric state prior (Any Number is 75% likely), $\alpha = 100$

less robust under these conditions. At $\alpha = 2$, for instance, L1 assigns only 70% probability to the FCI states Any Number and Only One upon hearing Or.

This basic model not only successfully derives FCI; it also suggests an explanation of why EI is not as strong as FCI. To show this, we will manipulating the listener's prior beliefs about states. As we change prior probabilities, the part of the posterior that corresponds to FCI remains essentially unchanged while the part which corresponds to EI varies. As a baseline, we assign uniform prior probability across all states as we have done so far. Here FCI is already stronger than EI (see Table 5): None of L1's posterior probability upon hearing Or is on non-FCI states but fully half of it is on the non-EI state Any Number.

This difference in strength becomes even more marked as we move to a context in which it is a priori more likely that one may take more than one fruit (e.g. because there is an excess of fruit available). Assume a nonuniform prior that assigns the non-EI state Any Number a higher probability than any other state. Upon hearing Or, listeners will assign this state a higher posterior than any other state. For example, if the prior assigns 75% probability to Any Number and equal probability to each of the other states, L1 assigns 92% posterior probability to Any Number (see Table 6). Suppose instead that it is the Only One state that is assigned 75% prior probability; in that case 92% posterior probability goes to Only One as well. Thus the posterior tracks and accentuates the prior.

These model results suggest that the optionality of EI is due to the fact that prior beliefs concerning EI remain present in the posterior and in fact become more marked there. This is in line with intuitions put forth by Geurts (2010) and Franke (2011) to the effect that EI arises from world knowledge. Conversely, the optionality of EI in our model also provides clues as to why FCI might be a stronger inference than EI. As it turns out, assigning a comparable amount of prior probability to non-FCI states does not destabilize FCI in the posterior as derived in our model. Rather, the low posterior probabilities of non-FCI states given Or that we have observed for flat priors remain virtually unaffected even with this modification. For example, assigning 75% prior probability to Only A and keeping α at 100 leads to virtually

	Only A	Only B	Only One	Any Number
A	0.67	0	0.17	0.17
B	0	0.67	0.17	0.17
Or	0	0	0.5	0.5

Table 7: L1 with And and Only Both removed, $\alpha = 100$

	Only A	Only B	Only One	Any Number	Only Both
A	0.57	0	0.14	0.14	0.14
B	0	0.57	0.14	0.14	0.14
Or	0	0	0.5	0.5	0
Null	0	0	0	0	1

Table 8: L1 with And removed and null utterance added, $\alpha = 100$

the same results for disjunction as uniform priors did (as was shown in Table 5). Thus in our model FCI, unlike EI, is robust to changes in world knowledge.

While we have presented a minimal working model here, many of our specific assumptions are not crucial to deriving FCI. For example, one might question whether the conjunctive utterance And should be considered an alternative. As it turns out, FCI still arises if we remove it from the utterance set; however, this requires some adjustments to the model. Without And, $\mathscr{I}_2$ no longer contains an utterance that is true in the Only Both state; but the probability function in (8b) is only well-defined if every state can be truthfully described by at least one utterance under every interpretation function. One solution is to remove the problematic Only Both state, in which case FCI straightforwardly arises in L1 (see Table 7). Alternatively, we can leave Only Both in the model, and add a null utterance that is true at every state under both interpretations and that represents the option of saying nothing. This approach also derives FCI because saying nothing now becomes the most effective strategy to convey the Only Both state, as can be seen in Table 8.

We might also choose to modify the set of interpretation functions supplied to the RSA model. The model is robust to the addition of new interpretation functions; we have found that many different combinations of such functions result in FCI. Ideally, one would determine the set of interpretation functions in a principled manner, perhaps by including all the LFs that can be generated by various exhaustification algorithms. For reasons of space, we do not pursue this further here.

4 No Free Choice Under Negation

As noted in the introduction, free choice inferences tend to disappear under negation and other downward-entailing environments. Here we show that our model correctly accounts for the absence of FCI under negation without any further stipulations. While our account postulates semantic ambiguity as a part of the explanation of FCI in upward-entailing environments, we demonstrate here that this ambiguity does not lead to FCI under negation. The reason for this is that RSA derives FCI via pragmatic strengthening of the original utterance, while in the scope of negation deriving FCI would amount to weakening rather than strengthening the utterance (Kratzer & Shimoyama 2002). By contrast, a system like Fox (2007), in which the explanatory burden rests exclusively on insertion of silent exhaustification operators, has to appeal to a stipulation that prevents insertion of these operators into LFs whose semantic meaning would be weakened as a result (such as the *Maximize Strength* rule in Chierchia 2013: 25). Like other pragmatic accounts, our theory can be seen as providing a principled explanation for why the effect of such a rule should be present. Semantic systems like Zimmermann (2000) that attribute FCI to nonstandard lexical entries likewise typically struggle with explaining why it goes away in the scope of negation. In this respect, our model shares the advantage of Gricean accounts more generally.

To capture the absence of FCI under negation, we assume that one of the LFs for an utterance like (4a), repeated here as (9), is equivalent to the classical $\neg\diamond(A \vee B)$:

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

This semantically entails $\neg\diamond(A \wedge \neg B)$ and $\neg\diamond(B \wedge \neg A)$ (as well as the more general $\neg\diamond A$ and $\neg\diamond B$). The unavailable FCI reading, $\neg[\diamond(A \wedge \neg B) \wedge \diamond(B \wedge \neg A)]$, lacks these entailments. Hence the absence of free choice under negation is already correctly predicted by the semantics. All that is left to do is show that, given natural assumptions, our pragmatic account does not upset this result.

4.1 Utterances and states

The four utterances in our negation model are simply the negated forms of (5):

(10) a. You may not take an apple. Not A
b. You may not take a pear. Not B
c. You may not take an apple or a pear. Not Or
d. You may not take an apple and a pear. Not And

Given the utterance set in (10), only four states can be differentiated: Only A, Only B, Only One, and Neither. The first three states are the same as in the basic model. The Any Number and Only Both states cannot be truthfully described by any of the utterances in (10), and are therefore not carried over from the basic model. The only addition is the Neither state, in which no fruits may be taken at all. We add this state because it can be truthfully described by every utterance.

4.2 Interpretation Functions

In analogy with the model we used to derive FCI in the absence of negation, we assume that two interpretation functions $\mathscr{I}_1$ and $\mathscr{I}_2$ apply to our utterances. As before, $\mathscr{I}_1$ corresponds to the classical modal logic denotations of each utterance. We extend it to the negated utterances by letting $\mathscr{I}_1(\neg\phi)$ denote the complement of $\mathscr{I}_1(\phi)$ for each utterance ϕ, except that we remove Any Number and Only Both and add Neither as just described. We extend $\mathscr{I}_2$ analogously. In this way, we arrive at the interpretation functions in (11) and (12). In the basic model, $\mathscr{I}_2$ strengthens utterance meanings as compared to $\mathscr{I}_1$; here, due to negation $\mathscr{I}_2$ weakens meanings.

(11) Interpretation Function 1

a. $[\![\mathsf{Not\ A}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ B}, \mathsf{Neither}\}$

b. $[\![\mathsf{Not\ B}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ A}, \mathsf{Neither}\}$

c. $[\![\mathsf{Not\ Or}]\!]^{\mathscr{I}_1} = \{\mathsf{Neither}\}$

d. $[\![\mathsf{Not\ And}]\!]^{\mathscr{I}_1} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Neither}\}$

(12) Interpretation Function 2

a. $[\![\mathsf{Not\ A}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Neither}\}$

b. $[\![\mathsf{Not\ B}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ A}, \mathsf{Only\ One}, \mathsf{Neither}\}$

c. $[\![\mathsf{Not\ Or}]\!]^{\mathscr{I}_2} = \{\mathsf{Neither}\}$

d. $[\![\mathsf{Not\ And}]\!]^{\mathscr{I}_2} = \{\mathsf{Only\ A}, \mathsf{Only\ B}, \mathsf{Only\ One}, \mathsf{Neither}\}$

The only uncertainty in this negation model emerges with utterances Not A and Not B; depending on which interpretation function is used, these utterances may or may not rule out the Only One state. The utterances Not Or and Not And have the same interpretation under both $\mathscr{I}_1$ and $\mathscr{I}_2$. This lack of ambiguity is due to two reasons. First, the states Any Number and Only Both have been excluded; as mentioned above, this exclusion turns out not to affect our results. Second, as can be seen in (6) and (7), it is only at these states that the two interpretation functions disagree on the meaning of the nonnegated utterances Or and And. The upshot is that there is in effect no uncertainty regarding the literal meaning of Not Or and of Not And in our model. Furthermore, under both interpretations Not Or is already maximally strong since it is true at only one state. Since RSA always strengthens

	Only A	Only B	Only One	Neither
Not A	0	0.8	0.2	0
Not B	0.8	0	0.2	0
Not Or	0	0	0	1
Not And	0	0	1	0

Table 9: L1 in negation model, $\alpha = 100$

meanings and never weakens them, it will not come as a surprise that it does not compute any FCI or other inferences in this case.[4] Thus, it does not matter if the negated disjunction utterance is interpreted using $\mathscr{I}_1$ or $\mathscr{I}_2$. In every case, we arrive at the same denotation {Neither}. Without uncertainty regarding the truth values of this utterance, absence of free choice under negation at L0 and all higher levels is straightforwardly derived, as we now show.

4.3 Negation Model Predictions

What does deriving the absence of free choice amount to? Recall that in the context of our model, deriving FCI for the nonnegated disjunction utterance amounts to predicting that the listener assigns essentially zero probability to the Only A and Only B states upon hearing that utterance, i.e. L1 takes Or to be false at these states. Now consider the utterance Not Or. If a free choice interpretation was present under negation, then due to the negation, the utterance as a whole would be true at these states. That free choice is absent under negation therefore means that Not Or is not used to convey these states. So, deriving the absence of FCI for the negated disjunction amounts to predicting that L1 assigns zero probability to these states upon hearing that utterance.

Given the lack of ambiguity for Not Or, listeners at all levels take it to denote

4 As before, we may think of $\mathscr{I}_2$ as being derived by covert exhaustification following the innocent exclusion algorithm in Fox (2007). As it turns out, no matter where exhaustification is applied in the LF, the resulting meaning is the same given our restricted state space. The unexhaustified LF (i) $\neg\diamond(A \vee B)$ denotes {Neither}. The LF (ii) $\text{Exh}_C(\neg\diamond(A \vee B))$ is logically equivalent. The LF (iii) $\neg\text{Exh}_C(\diamond(A \vee B)) = \neg[\diamond(A \vee B) \wedge \neg\diamond(A \wedge B)]$ rules out the Only A, Only B, and Only One states because they do not grant permission to take both fruit; as for the LF (iv) $\neg\diamond(\text{Exh}_C(A \vee B)) = \neg\diamond((A \vee B) \wedge \neg(A \wedge B))$, it rules out any states where one has permission to take one fruit and leave the other. This rules out the same states as before, as well as Any Number. The remaining states that can be truthfully described are {Any Number, Only Both, Neither} for LF (iii) and {Only Both, Neither} for LF (iv). Since Any Number and Only Both are not in the restricted state space, they are removed from consideration, leaving us with {Neither} for LFs (iii) and (iv) as well.

	Only A	Only B	Only One	Any Number	Only Both	Neither
Not A	0	0.67	0.17	0.17	0	0
Not B	0.67	0	0.17	0.17	0	0
Not Or	0	0	0	0	0.33	0.67
Not And	0	0	1	0	0	0
Null	0	0	0	0.5	0.5	0

Table 10: L1 in expanded negation model with null utterance, $\alpha = 100$

{Neither}. Table 9 shows the behavior of L1 in this model. We thus have a straightforward semantic explanation for the absence of FCI under negation. The pattern is stable at higher levels. L2 and upwards are identical to L1 except that values above 0.5 are replaced by 1 and values below 0.5 by 0.

We have assumed here that the space of states is limited to those in which the utterance or one of its alternatives is true on its classical modal logic interpretation. Given that interpretation, the states Any Number and Only Both are incompatible with all utterances, which is why we have excluded them. One might ask what would happen if these two states remained included. In that case, since according to $\mathscr{I}_1$, at Any Number and Only Both there are no true messages, we need to include a null utterance that is true everywhere, as described in Section 3.3. As Table 10 shows, the results remain essentially unchanged; in particular, we still correctly derive the absence of FCI for negated disjunctions. As before, the basic pattern remains the same for higher levels, with values above 0.5 replaced by 1, values below 0.5 by 0, and values at 0.5 unchanged.[5]

We end this section with a remark on the alternative utterances Not A and Not B, which do introduce uncertainty at the level of literal interpretation. We see in Tables 9 and 10 that even for high values of α, L1 assigns some probability to Only One for these utterances. However, L2 and above exclude Only One given these utterances. In fact, LF uncertainty does not affect interlocutor behavior. Pragmatic listeners in a model that only uses the unexhaustified LF, i.e. only $\mathscr{I}_1$, converge on the same pattern as those that assume $\mathscr{I}_1$ and $\mathscr{I}_2$ as illustrated here. Thus, uncertainty does not interfere with pragmatic interpretation under negation.

5 Given Footnote 4, one may wonder how adding Any Number and Only Both back into the model would affect interpretation functions. If we take $\mathscr{I}_2$ to represent interpretations that result from silent exhaustification, we may want to set $\mathscr{I}_2(\text{Not Or}) = \{\text{Neither}, \text{Only Both}\}$ since $\mathscr{I}_2(\text{Or})$ includes Only Both. This does not affect the results reported in this section as long as we keep $\mathscr{I}_1(\text{Not Or}) = \{\text{Neither}\}$. Even though for L0, Not Or no longer unambiguously rules out the Only Both state, this does not matter for L1, as the null utterance remains a more reliable way to convey this state.

5 Conclusion

In this paper we have proposed a pragmatic model of free choice disjunction within the Rational Speech Act (RSA) framework, which formalizes the way in which rational and cooperative agents reason in conversation (Frank & Goodman 2012). We have additionally accounted for the exclusivity inference (EI), i.e. the inference from $\diamond(A \vee B)$ to $\neg\diamond(A \wedge B)$. We have explained why this inference is easier to cancel than the free choice inference (FCI), i.e. the inference to $\diamond A$ and $\diamond B$. Our model incorporates aspects of Fox (2007), Franke (2011), and Bergen et al. (2016), combining game-theoretic modeling with semantic uncertainty that can be thought of as resulting from different LFs derived by the optional insertion of silent exhaustification operators. We have formalized this semantic uncertainty via multiple interpretation functions, or mappings of utterances to propositions. Given these competing interpretations, the speaker is uncertain whether or not, in a context where apples and pears are under discussion, the hearer would interpret *You may take an apple* as forbidding them from taking a pear. Unlike e.g. in Kratzer & Shimoyama (2002), this uncertainty is semantic and distinct in nature from scalar implicatures. The speaker consequently uses the disjunction to prevent the hearer from concluding that any fruit is forbidden. The hearer concludes that they can choose either fruit.

We have also accounted for the asymmetry in cancelability between FCI and EI. In our model, a listener's prior beliefs about the world are the driving force for EI. In other words, whether a listener concludes that they may not take both fruit at once depends on how likely they take this permission to be in the first place. FCI is not as sensitive to these prior beliefs because it arises primarily as a function of the conversational agents' pragmatic reasoning process. Thus, it is possible to use the same mechanism in order to derive inferences which are not of the same strength.

Finally, as pragmatic models in general do, our model derives the disappearance of FCI under negation due to the already maximally informative meaning of the negated disjunction. We have shown that the additional assumptions we have used to account for FCI in non-negated utterances do not get in the way of this explanation.

Assuming multiple available interpretation functions in RSA allows us to model alternatives at both the semantic and pragmatic level. At the semantic level, the conversational agents must reason about competing LFs for a single utterance, as in Fox (2007). At the pragmatic level, they must also reason about alternative utterances a speaker may choose from, as in Franke (2011). While Fox provides a mechanism for the generation of semantic alternatives, his account does not fully explain why listeners prefer certain LFs over others (Alsop et al. to appear). The flexibility of RSA to model both types of alternatives supplements Fox (2007) in accounting for the way speakers and listeners successfully communicate using a system in which the silent nature of exhaustification operators leads to massive ambiguities.

References

Aloni, Maria. 2007. Free choice, modals, and imperatives. *Natural Language Semantics* 15(1). 65–94. doi:10.1007/s11050-007-9010-2.

Alsop, Anna, Lucas Champollion & Ioana Grosu. to appear. A problem for Fox's (2007) account of free choice disjunction. *Snippets* .

Bergen, Leon, Roger Levy & Noah D. Goodman. 2016. Pragmatic reasoning through semantic inference. *Semantics and Pragmatics* 9(20). doi:10.3765/sp.9.20.

Chierchia, Gennaro. 2013. *Logic in grammar: polarity, free choice, and intervention*. Oxford, UK: Oxford University Press. doi:10.1093/acprof:oso/9780199697977.001.0001.

Fox, Danny. 2007. Free choice and the theory of scalar implicatures. In Uli Sauerland & Penka Stateva (eds.), *Presupposition and implicature in compositional semantics*, 71–120. London, UK: Palgrave Macmillan. doi:10.1057/9780230210752_4.

Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998–998. doi:10.1126/science.1218633.

Franke, Michael. 2011. Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics and Pragmatics* 4(1). 1–82. doi:10.3765/sp.4.1.

Gazdar, Gerald. 1979. *Pragmatics: Implicature, presupposition and logical form*. New York, NY: Academic Press.

Geurts, Bart. 2010. *Quantity implicatures*. Cambridge, UK: Cambridge University Press. doi:10.1017/CBO9780511975158.

Jäger, Gerhard. 2011. Game theory in semantics and pragmatics. In Klaus von Heusinger, Claudia Maienborn & Paul Portner (eds.), *Semantics: An international handbook of natural language meaning*, vol. 3 Handbücher zur Sprach- und Kommunikationswissenschaft / Handbooks of Linguistics and Communication Science (HSK), chap. 93, 2487–2516. de Gruyter. doi:10.1515/9783110253382.2487.

Kamp, Hans. 1974. Free choice permission. *Proceedings of the Aristotelian Society* 74(1). 57–74. doi:10.1093/aristotelian/74.1.57.

Kratzer, Angelika & Junko Shimoyama. 2002. Indeterminate pronouns: The view from Japanese. In Yukio Otsu (ed.), *3rd Tokyo conference on psycholinguistics*, 1–25. doi:10.1007/978-3-319-10106-4_7.

Meyer, Marie-Christine. to appear. An apple or a pear: free choice disjunction. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann & Thomas Ede Zimmermann (eds.), *Wiley's companion to semantics*, Hoboken, NJ: Wiley.

van Rooij, Robert. 2000. Permission to change. *Journal of Semantics* 17(2). 119–143. doi:10.1093/jos/17.2.119.

Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and*

Philosophy 27(3). 367–391. doi:10.1023/b:ling.0000023378.71748.db.
Schlenker, Philippe. 2016. The semantics–pragmatics interface. In Maria Aloni & Paul Dekker (eds.), *Cambridge handbook of formal semantics* Cambridge Handbooks in Language and Linguistics, chap. 22, 664–727. Cambridge, UK: Cambridge University Press. doi:10.1017/cbo9781139236157.023.
Simons, Mandy. 2005. Dividing things up: The semantics of *or* and the modal/*or* interaction. *Natural Language Semantics* 13(3). 271–316. doi:10.1007/s11050-004-2900-7.
Zimmermann, Thomas Ede. 2000. Free choice disjunction and epistemic possibility. *Natural Language Semantics* 8(4). 255–290. doi:10.1023/a:1011255819284.

Lucas Champollion
10 Washington Place
New York, NY 10003
champollion@nyu.edu

Anna Alsop
10 Washington Place
New York, NY 10003
aalsop@nyu.edu

Ioana Grosu
10 Washington Place
New York, NY 10003
ig950@nyu.edu