# Free Choice from Iterated Best Response

Michael Franke

Universiteit van Amsterdam & ILLC
Amsterdam, The Netherlands

**Abstract.** This paper summarizes the essence of a recent game theoretic explanation of free choice readings of disjunctions under existential modals ([8]). It introduces principles of game model construction to represent the context of utterance, and it spells out the basic mechanism of iterated best response reasoning in signaling games.

**Key words:** conversational implicatures, game theoretic pragmatics, free choice disjunction, iterated best response

## 1 Free Choice Disjunctions & Game Theory

Contrary to their logical semantics, disjunctions under modal operators as in (1a) may receive free-choice readings (FC-readings) as in (1b) ([12]).

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

b. You may take an apple and you may take a pear. $\Diamond A \wedge \Diamond B$

This inference is not guaranteed by the standard logical semantics which treats disjunction as truth-functional connective and the modal as an existential quantifier over accessible worlds. Of course, different semantics of disjunctions or modals are conceivable and have been proposed by, for instance, [12], [18] or [1]. But, all else being equal, a *pragmatic* solution that retains the logical semantics and treats FC-readings as Gricean inferences seems preferable (cf. the arguments in [16]).

Unfortunately, a naïve approach to Gricean scalar reasoning does not suffice. If we assume that the set of expression alternatives with which to compare an utterance of (1a) contains the simple expressions in (2), we run into a problem.

(2) a. You may take an apple. $\Diamond A$

b. You may take a pear. $\Diamond B$

Standard scalar reasoning tells us that all semantically stronger alternatives are to be inferred *not* to be true. This yields that $\Box\neg A$ and that $\Box\neg B$ are true, which contradicts (1a) itself.

This particular problem has a simple solution. [14] observe that the FC-reading follows from naïve scalar reasoning based on the alternatives in (2) if we use the already exhaustified readings of the alternatives as in (3).

(3) a. You may take an apple, but no pear. $\Diamond A \wedge \neg\Diamond B$

b. You may take a pear, but no apple. $\Diamond B \wedge \neg\Diamond A$

Truth of (1a) together with the falsity of both sentences in (3) entails the FC-reading in (1b).

There is clearly a certain intuitive appeal to this idea: when reasoning about expression alternatives it is likely that potential pragmatic enrichments of these may at times be taken into account as well. But when and how exactly? Standard theories of scalar reasoning do not integrate such nested pragmatic reasoning. This has been taken as support for theories of local implicature computation in the syntax where exhaustifity operators can apply, if necessary, several times ([5, 7]). But the proof that such nested or iterated reasoning is very much compatible with a systematic, global, and entirely Gricean approach amenable to intuitions about economic language use is still up in the air.

Enter game theory. Recent research in game theoretic pragmatics has produced a number of related models of agents' step-by-step pragmatic reasoning about each others' hypothetical behavior ([17, 2, 10]). This is opposed to the more classical equilibrium-based solution concepts which merely focus on stable outcomes of, mostly, repeated play or evolutionary dynamics. The main argument of this paper is that such step-by-step reasoning, which is independently motivated, explains free-choice readings along the lines sketched above: early steps of such reasoning establish the exhaustive readings of alternative forms, while later steps of the same kind of global reasoning can pick on previously established readings.

In order to introduce and motivate this game theoretical approach, two sets of arguments are necessary.[1] Firstly, we need to settle on what kind of *game model* is required in order to represent conversational moves and their interpretation. This is to be addressed in section 2. Secondly, we need to spell out a *solution concept* by means of which pragmatic language use can be explained in the chosen game models. This is the topic of section 3. Finally, section 4 reviews briefly how this approach generalizes.

## 2 Interpretation Games as Context Models

It is standard in game-theoretic pragmatics to assume that an informative assertion and its uptake can reasonably be modelled as a *signaling game*. More specifically then, the pragmatic interpretation of assertions can be modelled by a particular kind of signaling game, which I will call *interpretation game*. These latter games function as representations of the context of utterance (as conceived by the receiver) and are constructed from a given target expression whose interpretation we are interested in, together with its natural Neo-Gricean alternatives and their logical semantics. Let me introduce both signaling games and interpretation games one after the other.

[1] These arguments can only be given in their bare essentials here (see [8] for the full story).

*Signaling Games.* A signaling game is a simple dynamic game with imperfect information between a sender and a receiver. The sender has some private information about the state of the world $t$ which the receiver lacks. The sender chooses a message $m$ from a given set of alternatives, all of which we assume to have a semantic meaning commonly known between players. The receiver observes the sent message $m$ and chooses an action $a$ based on this observation. An *outcome* of playing a signaling game for one round is given by the triple $t$, $m$ and $a$. Each player has his own preferences over such outcomes.

More formally speaking, a signaling game (with meaningful signals) is a tuple

$$\langle \{S, R\}, T, \Pr, M, [\![\cdot]\!], A, \mathrm{U}_S, \mathrm{U}_R \rangle$$

where sender $S$ and receiver $R$ are the players of the game; $T$ is a set of states of the world; $\Pr \in \Delta(T)$ is a probability distribution over $T$, which represents the receiver's uncertainty which state in $T$ is actual;[2] $M$ is a set of messages that the sender can send; $[\![\cdot]\!] : M \rightarrow \mathcal{P}(T) \setminus \emptyset$ is a denotation function that gives the predefined semantic meaning of a message as the set of all states where that message is true; $A$ is the set of response actions available to the receiver; and $\mathrm{U}_{S,R} : T \times M \times A \rightarrow \mathbb{R}$ are utility functions for both sender and receiver.

*Interpretation Games.* For models of natural language interpretation a special class of signaling games is of particular relevance. To explain pragmatic inferences like implicatures we should look at *interpretation games.* I assume here that these games can be constructed generically from a set of alternatives to the to-be-interpreted expression, together with their logical semantics. Here are the assumptions and the construction steps.

Firstly, the set of receiver actions is equated with the set of states $A = T$ and the receiver's utilities model his interest in getting to know the true state of affairs, i.e., getting the right *interpretation* of the observed message:

$$\mathrm{U}_R(t, m, a) = \begin{cases} 1 & \text{if } t = a \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, in the vein of [9], we assume that conversation is a *cooperative* effort —at least on the level of such generic context models— so that the sender shares the receiver's interest in correct interpretation:[3]

$$\mathrm{U}_S(t, m, a) = \mathrm{U}_R(t, m, a).$$

The set $T$ of state distinctions is to be derived from the set $M$ of messages given by some (normal, natural, Neo-Gricean) set of alternative forms to the

[2] As for notation, $\Delta(X)$ is the set of all probability distributions over set $X$, $Y^X$ is the set of all functions from $X$ to $Y$, $X : Y \rightarrow Z$ is alternative notion for $X \in Z^Y$, and $\mathcal{P}(X)$ is the power set of $X$.

[3] Notice that this implicitly also commits us to the assumption that all messages are equally costly, or, if you wish, costless.

target sentence whose implicatures we are interested in. Clearly, not every possible way the world could be can be distinguished with any set $M$. So we should restrict ourselves to only those states that can feasibly be expressed with the linguistic means at hand. What are those distinctions? Suppose $M$ contains only logically independent alternatives. In that case, we could in principle distinguish $2^M$ possible states of the world, according to whether some subset of messages $X \subseteq M$ is such that all messages in $X$ are true, while all messages in its complement are false. (This is what happens in propositional logic, when we individuate possible worlds by all different valuations for a set of proposition letters.) But for normal pragmatic applications the expressions in $M$ will not all be logically independent. So in that case we should look at states which can be *consistently* described by a set of messages $X \subseteq M$ all being true while all expressions in its complement are false. Moreover, since at least the target message may be assumed true for pragmatic interpretation, we should define the set of states of the interpretation game as given by the set of all subsets $X \subseteq M$ containing the target message such that the formula

$$\bigwedge X \wedge \neg \bigvee M \setminus X$$

is consistent. With this, also the semantic denotation function $[\![\cdot]\!]$ is then straightforwardly defined as:

$$[\![m]\!] = \{t \in T \mid m \in t\} \,.$$

Finally, since we are dealing with general models of utterance interpretation, we should not assume that the receiver has biased beliefs about which specific state obtains. This simply means that in interpretation games $\Pr(\cdot)$ is a flat probability distribution.

*Example.* To give a concrete example, here is how to construct an interpretation game for the target expression in (1a). Everything falls into place once a set of alternatives is fixed. To keep the exposition extremely simple, let us first only look at the set of messages in (4). (See section 4 for more discussion.)

(4) a. You may take an apple or a pear. $m_{\Diamond(\mathtt{A}\vee\mathtt{B})}$
b. You may take an apple. $m_{\Diamond\mathtt{A}}$
c. You may take a pear. $m_{\Diamond\mathtt{B}}$

Based on these alternatives, there are three states we need to distinguish:

$$\begin{aligned} t_{\mathtt{A}} &= \left\{m_{\Diamond\mathtt{A}}, m_{\Diamond(\mathtt{A}\vee\mathtt{B})}\right\} \\ t_{\mathtt{B}} &= \left\{m_{\Diamond\mathtt{B}}, m_{\Diamond(\mathtt{A}\vee\mathtt{B})}\right\} \\ t_{\mathtt{AB}} &= \left\{m_{\Diamond\mathtt{A}}, m_{\Diamond\mathtt{B}}, m_{\Diamond(\mathtt{A}\vee\mathtt{B})}\right\} . \end{aligned}$$

Here, $t_{\mathtt{A}}$ is a state where the hearer may take an apple but no pear, and $t_{\mathtt{AB}}$ is a state where the hearer may take both an apple and a pear. These states yield the interpretation game in figure 1. Notice that we consider only these states,

| | Pr($t$) | $a_{\mathtt{A}}$ | $a_{\mathtt{B}}$ | $a_{\mathtt{AB}}$ | $m_{\Diamond\mathtt{A}}$ | $m_{\Diamond\mathtt{B}}$ | $m_{\Diamond(\mathtt{A}\vee\mathtt{B})}$ |
|---|---|---|---|---|---|---|---|
| $t_{\mathtt{A}}$ | 1/3 | 1,1 | 0,0 | 0,0 | $\surd$ | – | $\surd$ |
| $t_{\mathtt{B}}$ | 1/3 | 0,0 | 1,1 | 0,0 | – | $\surd$ | $\surd$ |
| $t_{\mathtt{AB}}$ | 1/3 | 0,0 | 0,0 | 1,1 | $\surd$ | $\surd$ | $\surd$ |

**Fig. 1.** Interpretation game constructed from (1a) and (4)

because these are the only distinctions we can make between worlds where the target message (1a) is true that can be expressed based on consistent valuations of all alternatives. Certainly, in the present case, this is nearly excessively simple, but it is not trivial and, most importantly, there is still room for pragmatic interpretation: there are still many ways in which sender and receiver could coordinate on language use in this game. What is needed is a solution concept that singles out uniquely the player behavior that explains the free choice inference.

## 3 Iterated Best Response Reasoning

Generally, behavior of players is represented in terms of *strategies*. A *pure sender strategy* $s \in \mathsf{S} = M^T$ is a function from states to messages and a *pure receiver strategy* $r \in \mathsf{R} = A^M$ is a function from messages to actions. A *pure strategy profile* $\langle s, r \rangle$ is then a characterization of the players' *joint behavior* in a given signaling game. For instance, the tuple:

$$s = \left\{ \begin{array}{ll} t_{\mathtt{A}} & \mapsto m_{\Diamond\mathtt{A}} \\ t_{\mathtt{B}} & \mapsto m_{\Diamond\mathtt{B}} \\ t_{\mathtt{AB}} & \mapsto m_{\Diamond(\mathtt{A}\vee\mathtt{B})} \end{array} \right\} \qquad r = \left\{ \begin{array}{ll} m_{\Diamond\mathtt{A}} & \mapsto t_{\mathtt{A}} \\ m_{\Diamond\mathtt{B}} & \mapsto t_{\mathtt{B}} \\ m_{\Diamond(\mathtt{A}\vee\mathtt{B})} & \mapsto t_{\mathtt{AB}} \end{array} \right\} \tag{1}$$

is a strategy profile for the game in figure 1. And a special one, indeed. It corresponds to the intuitive way of using the corresponding natural language expressions: the interpretation of $m_{\Diamond\mathtt{A}}$, for instance, is the exhaustive reading that only $A$, but not $B$ is allowed; and the interpretation of $m_{\Diamond(\mathtt{A}\vee\mathtt{B})}$ is the free choice inference that both taking $A$ and taking $B$ are allowed. This is therefore what a solution concept is required to predict in order to explain FC-readings based on the game in figure 1.

But the strategy profile in (1) is not the only one there is. Also, the rather unintuitive pooling strategy profile

$$s = \left\{ \begin{array}{ll} t_{\mathtt{A}} & \mapsto m_{\Diamond(\mathtt{A}\vee\mathtt{B})} \\ t_{\mathtt{B}} & \mapsto m_{\Diamond(\mathtt{A}\vee\mathtt{B})} \\ t_{\mathtt{AB}} & \mapsto m_{\Diamond(\mathtt{A}\vee\mathtt{B})} \end{array} \right\} \qquad r = \left\{ \begin{array}{ll} m_{\Diamond\mathtt{A}} & \mapsto t_{\mathtt{AB}} \\ m_{\Diamond\mathtt{B}} & \mapsto t_{\mathtt{AB}} \\ m_{\Diamond(\mathtt{A}\vee\mathtt{B})} & \mapsto t_{\mathtt{AB}} \end{array} \right\} \tag{2}$$

is conceivable. What is worse, both strategy profiles describe an equilibrium state: given the behavior of the opponent neither player has an incentive to deviate. But, clearly, to explain the FC-reading, the profile in (1) should be selected, while the profile in (2) should be ruled out. In other words, we need a mechanism with which to select one equilibrium and rule out others.

*IBR Models.* One way of looking at an *iterated best response model* (IBR model) is exactly that: a plausible mechanism with which reasoners (or a population) may arrive at an equilibrium state (rather than another). An IBR model assumes that agents reason about each other's behavior in a step-by-step fashion. The model is anchored in naïve behavior of level-0 players that do not take opponent behavior into account, but that may be sensitive to other non-strategic, psychological factors, such as, in our case, the semantic meaning of messages. Players of level-$(k+1)$ assume that their opponent shows level-$k$ behavior and play a best response to this belief.[4]

Here is a straightforward IBR sequence as a solution concept for signaling games. Naïve players of level-0 are defined as playing some arbitrary strategy that conforms to semantic meaning. For the sender, this yields:

$$S_0 = \{s \in \mathsf{S} \mid \forall t \in T: \; t \in [\![s(t)]\!]\}\,.$$

Level-0 senders are characterized by the set of all pure strategies that send only true messages. For interpretation games, naïve receiver types receive a similarly straightforward characterization:

$$R_0 = \{r \in \mathsf{R} \mid \forall m \in M: \; r(m) \in [\![m]\!]\}\,.$$

Level-0 receivers are characterized by the set of all pure strategies that interpret messages as true.

In order to define level-$(k+1)$ types, it is necessary to define the notion of a *best response* to a belief in level-$k$ behavior. There are several possibilities of defining beliefs in level-$k$ behavior.[5] The most convenient approach is to assume that agents have *unbiased beliefs* about opponent behavior. Unbiased beliefs in level-$k$ behavior do not favor any one possible level-$k$ behavior, if there are several, over any other, and can therefore be equated simply with a flat probability distribution over the set of level-$k$ strategies.

Turning first to higher-level sender types, let us write $R_k(m,a)$ for the probability that a level-$k$ receiver who is believed to play a random strategy in $R_k$ will play $a$ after observing $m$. Then level-$(k+1)$ senders are defined by

$$S_{k+1} = \left\{ s \in \mathsf{S} \mid s(t) \in \arg\max_{m \in M} \sum_{a \in A} R_k(m,a) \times \mathrm{U}_S(t,m,a) \right\}$$

as the set of all best responses to that unbiased belief.

For higher-level receiver types the same standard definition applies once we have characterized the receiver's *posterior beliefs*, i.e., beliefs the receiver holds about the state of the world after he observed a message. These need to be

[4] Models of this kind are good predictors of laboratory data on human reasoning (see, for instance [3]), but also solve conceptual issues with equilibrium solution concepts (see [6]). Both of these aspects make IBR models fit for use in linguistic applications.

[5] This is the crucial difference between various IBR models such as given by [4], [11] and [8], for instance.

derived, again in entirely standard fashion, from the receiver's prior beliefs $\Pr(\cdot)$ and his beliefs in sender behavior as given by $S_k$. Let $S_k(t,m)$ be the probability that a level-$k$ sender who is believed to play a random strategy in $S_k$ will send $m$ in state $t$. A level-$(k+1)$ receiver has posterior beliefs $\mu_{k+1} \in (\Delta(T))^M$ calculated by Bayesian conditionalization, as usual:

$$\mu_{k+1}(t|m) = \frac{\Pr(t) \times S_k(t,m)}{\sum_{t' \in T} \Pr(t') \times S_k(t',m)}.$$

Level-$(k+1)$ receivers are then defined as best responding to this posterior belief:

$$R_{k+1} = \left\{ r \in \mathsf{R} \mid r(m) \in \arg\max_{a \in A} \sum_{t \in T} \mu_{k+1}(t|m) \times \mathrm{U}_R(t,m,a) \right\}.$$

This last definition is incomplete. Bayesian conditionalization is only defined for messages that are not *surprise messages.* A surprise message for a level-$(k+1)$ receiver is a message that is not used by any strategy in $S_k$ in any state. A lot can be said about the proper interpretation of surprise messages (see the discussion in [11, 8, 15]). This is the place where different *belief revision strategies* of the receiver could be implemented, if needed or wanted. For the purposes of this paper it is sufficient to assume that whatever else the receiver may come to believe if he observes a surprise message, he will stick to the belief that it is true. So, if for some message $m$ we have $S_k(t,m) = 0$ for all $t$, then define $\mu_{k+1}(t|m) = \Pr(t| [\![m]\!])$.

*Example.* The simple IBR model sketched here does what we want it to: it uniquely singles out the intuitive equilibrium state in equation (1) for the game in figure 1. To see how this works, and to see where IBR may rationalize the use of exhaustified alternatives in Gricean reasoning, let us calculate the sequence of reasoning starting with $R_0$ for the simple game in figure 1 (the case starting with $S_0$ is parallel):[6]

$$R_0 = \left\{ \begin{array}{ll} m_{\Diamond\mathtt{A}} & \mapsto t_{\mathtt{A}}, t_{\mathtt{AB}} \\ m_{\Diamond\mathtt{B}} & \mapsto t_{\mathtt{B}}, t_{\mathtt{AB}} \\ m_{\Diamond(\mathtt{A}\vee\mathtt{B})} & \mapsto t_{\mathtt{A}}, t_{\mathtt{B}}, t_{\mathtt{AB}} \end{array} \right\} \qquad S_1 = \left\{ \begin{array}{ll} t_{\mathtt{A}} & \mapsto m_{\Diamond\mathtt{A}} \\ t_{\mathtt{B}} & \mapsto m_{\Diamond\mathtt{B}} \\ t_{\mathtt{AB}} & \mapsto m_{\Diamond\mathtt{A}}, m_{\Diamond\mathtt{B}} \end{array} \right\}$$

$$R_2 = \left\{ \begin{array}{ll} m_{\Diamond\mathtt{A}} & \mapsto t_{\mathtt{A}} \\ m_{\Diamond\mathtt{B}} & \mapsto t_{\mathtt{B}} \\ m_{\Diamond(\mathtt{A}\vee\mathtt{B})} & \mapsto t_{\mathtt{A}}, t_{\mathtt{B}}, t_{\mathtt{AB}} \end{array} \right\} \qquad S_3 = \left\{ \begin{array}{ll} t_{\mathtt{A}} & \mapsto m_{\Diamond\mathtt{A}} \\ t_{\mathtt{B}} & \mapsto m_{\Diamond\mathtt{B}} \\ t_{\mathtt{AB}} & \mapsto m_{\Diamond\mathtt{A}\vee\mathtt{B}} \end{array} \right\}$$

$$R_4 = \left\{ \begin{array}{ll} m_{\Diamond\mathtt{A}} & \mapsto t_{\mathtt{A}} \\ m_{\Diamond\mathtt{B}} & \mapsto t_{\mathtt{B}} \\ m_{\Diamond(\mathtt{A}\vee\mathtt{B})} & \mapsto t_{\mathtt{AB}} \end{array} \right\}.$$

Naïve receiver behavior only takes semantic meaning into account and this is what $S_1$ plays a best response to. Given $S_1$, message $m_{\Diamond\mathtt{A}}$ is interpreted exhaustively by $R_2$, as meaning "you may do $A$, but not $B$" (and similarly for $m_{\Diamond\mathtt{B}}$),

[6] Sets of pure strategies $Z \subseteq X^Y$ are represented by listing for each $x \in X$ the set of all $y \in Y$ such that for some strategy $z \in Z$ we have $z(x) = y$.

while message $m_{\Diamond(\texttt{A}\vee\texttt{B})}$ is a surprise message, and will be interpreted merely as true. This makes $m_{\Diamond(\texttt{A}\vee\texttt{B})}$ the only rational choice for $S_3$ to send in $t_{\texttt{AB}}$, so that in one more round of iteration we reach a fixed point equilibrium state in which $R_4$ assigns to $m_{\Diamond(\texttt{A}\vee\texttt{B})}$ the FC-reading that he may do $A$ and that he may do $B$. In sum, the FC-reading of $m_{\Diamond(\texttt{A}\vee\texttt{B})}$ is derived in two steps of receiver reasoning by first establishing an exhaustive interpretation of the alternatives, and then reasoning with this exhaustive interpretation to arrive at the FC-reading.

## 4 IBR Reasoning: The Bigger Picture

The previous two sections have tried to give, as short and yet accessible as possible, the main mechanism of IBR reasoning and the demonstration that IBR reasoning can account for FC-readings of disjunctions. Many assumptions of this approach could not have possibly been spelled out sufficiently, and so the impression may arise that IBR reasoning, as outlined here, is really only arbitrarily designed to deal with a small problem of linguistic interest. This is, decidedly, not so. There are good and independent motivations for both game model construction and solution concept, and both in tandem do good explanatory work, both conceptually and empirically (see [2, 11, 8]).

Moreover, it should be stressed that the IBR approach also handles more complex cases than the easy example discussed above, of course. Most importantly, it predicts well also when other scalar contrasts, such as given by (5a) or (5b), are taken into account as well.

(5) a. You must take an apple or a pear. $m_{\Box(\texttt{A}\vee\texttt{B})}$

b. You may take an apple and a pear. $m_{\Diamond(\texttt{A}\wedge\texttt{B})}$

Including more alternative messages results in bigger context models that include more state distinctions. But still IBR reasoning gives intuitive results. For instance, [8] spells out the IBR reasoning based on a set of alternatives that includes (4) and the conjunctive alternative in (5b). Doing so, we derive that (1a) is taken to implicate that $\Diamond(A \wedge B)$ is false. This is as it should be: in a context where the conjunctive alternative (5b) is salient, this inference should be predicted, but for the FC-reading alone the simple alternatives as in (4) should suffice. Similar considerations apply to the stronger modal alternative in (5a).

Generalizing the result further, it is possible to show that for any $n$-place case of the form $\Diamond(A_1 \vee \cdots \vee A_n)$ we derive the inference that $\Diamond A_i$ under IBR logic. The argument that establishes this result is a so-called *unravelling argument* which I can only sketch here: in the first step (of receiver reasoning) all "singleton" messages of the form $\Diamond A_i$ are associated with their exhaustive readings; in the second step all two-place disjunctions $\Diamond(A_i \vee A_j)$, $i \neq j$, are associated with states in which exactly two actions are allowed one of which must be $A_i$ or $A_j$;[7]

[7] In order to make this inference more specific, as it clearly should be, a slightly more careful setup of the reasoning sequence is necessary than given here. But this is a technical problem that does not disturb the conceptual point that is of relevance.

continuing in this way, after $n$ rounds of reasoning the form $\Diamond(A_1 \vee \cdots \vee A_n)$ gets the right interpretation that all actions $A_i$ are allowed.

Interestingly, IBR does *not need* to assume conjunctive alternatives even for the general $n$-place case, while [14]'s approach *has to*.[8] To see this, look at the three-placed case $\Diamond(A \vee B \vee C)$ with only alternatives $\Diamond A$, $\Diamond B$ and $\Diamond C$. The exhaustive readings of these are given in (6).

(6) a. $\Diamond A \wedge \neg\Diamond B \wedge \neg\Diamond C$

b. $\Diamond B \wedge \neg\Diamond A \wedge \neg\Diamond C$

c. $\Diamond C \wedge \neg\Diamond A \wedge \neg\Diamond B$

But truth of $\Diamond(A \vee B \vee C)$ together with the falsity of all sentences in (6) does not yield the FC-reading that any of $A$, $B$ or $C$ are allowed. To establish the FC-reading, we also need the alternatives $\Diamond(A \wedge B)$, $\Diamond(A \wedge C)$ and $\Diamond(B \wedge C)$ with their exhaustive readings in (7).

(7) a. $\Diamond(A \wedge B) \wedge \neg\Diamond C$

b. $\Diamond(A \wedge C) \wedge \neg\Diamond B$

c. $\Diamond(B \wedge C) \wedge \neg\Diamond A$

If we then want to account for the presence of the FC-reading in the absence of the scalar inference that $\Diamond(A \wedge B \wedge C)$ is false, we need to assume that all alternatives with two-placed conjunctions are given, but *not* the three-placed conjunctive alternative. This is not impossible, but also not very plausible.

Finally, let me also mention for the sake of completeness that the IBR approach also deals with free choice readings of disjunctions under universal modals in the exact same fashion as outlined here. A parallel account also deals with the structurally similar inference called *simplification of disjunctive antecedents* as exemplified in (8).

(8) a. If you take an apple or a pear, that's okay.

b. If you take an apple, that's okay. And if you take a pear, that's also okay.

The IBR model is also capable of dealing with epistemic ignorance readings such as forced by (9).

(9) You may take an apple or a pear, but I don't know which.

To capture these, however, the game models have to be adapted to include also possible sender uncertainty (see [8] for details).

[8] And with it, in slightly amended form, the syntactic account of [7].

## References


1. Asher, N. and Bonevac, D.: Free choice permission as strong permission. Synthese 145(3), 303–323 (2005)
2. Benz, A. and van Rooij, R.: Optimal assertions and what they implicate. Topoi 26, 63–78 (2007)
3. Camerer, C. F.: Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press (2003)
4. Camerer, C. F., Ho, T.-H., and Chong, J.-K.: A cognitive hierarchy model of games. The Quarterly Journal of Economics 119(3), 861–898 (2004)
5. Chierchia, G.: Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. In: Belletti, A. (ed.) Structures and Beyond, pp. 39–103. Oxford University Press (2004)
6. Crawford, V. P.: Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. American Economic Review 93(1), 133–149 (2003)
7. Fox, D.: Free choice and the theory of scalar implicatures. In: Sauerland, U. and Stateva, P. (eds.) Presupposition and Implicature in Compositional Semantics, pp. 71–120. Palgrave MacMillan, Hampshire (2007)
8. Franke, M.: Signal to Act: Game Theory in Pragmatics. PhD thesis, Universiteit van Amsterdam (2009)
9. Grice, P. H.: Studies in the Ways of Words. Harvard University Press (1989)
10. Jäger, G.: Game dynamics connects semantics and pragmatics. In: Pietarinen, A.-V. (ed.) Game Theory and Linguistic Meaning, pp. 89–102. Elsevier (2007)
11. Jäger, G. and Ebert, C.: Pragmatic rationalizability. In: Riester, A. and Solstad, T. (eds.) Proceedings of Sinn und Bedeutung 13, pp. 1–15 (2009)
12. Kamp, H.: Free choice permission. Proceedings of the Aristotelian Society 74, 57–74 (1973)
13. Kamp, H.: Semantics versus pragmatics. In: Guenthner, F. and Schmidt, S. J. (eds.) Formal Semantics and Pragmatics for Natural Languages, pp. 255–287. Reidel, Dordrecht (1978)
14. Kratzer, A. and Shimoyama, J.: Indeterminate pronouns: The view from Japanese. In: Otsu, Y. (ed.) Proceeding of the 3rd Tokyo Conference on Psycholinguistics, pp. 1–25 (2002)
15. Mühlenbernd, R.: Kommunikationsmodell für den Entwicklungsprozess von Implikaturen. Master's thesis, University of Bielefeld (2009)
16. Schulz, K.: A pragmatic solution for the paradox of free choice permission. Synthese 147, 343–377 (2005)
17. Stalnaker, R.: Saying and meaning, cheap talk and credibility. In: Benz, A., Jäger, G., and van Rooij, R. (eds.) Game Theory and Pragmatics, pp. 83–100. Palgrave MacMillan, Hampshire (2006)
18. Zimmermann, T. E.: Free choice disjunction and epistemic possibility. Natural Language Semantics 8, 255–290 (2000)