

Game theory in semantics and pragmatics

Gerhard Jäger
University of Tübingen
Institute of Linguistics
Wilhelmstr. 19, 72074 Tübingen, Germany
phone: +49-7071-29-77303, fax: +49-7071-29-5212
gerhard.jaeger@uni-tuebingen.de

submitted: January 2011
last revision: September 2012

1 Introduction

The interpretation of linguistic utterances is determined by the words involved and the way they are combined, but not exclusively so. Establishing the content that is communicated by an utterance is inextricably intertwined with the communicative context where the utterance is made, including the expectations of the interlocutors about each other. Clark and Marshall (1981) make a good case that even the reference of a definite description depends on the reasoning of the speaker and hearer about each other's knowledge state. Likewise, computing the implicatures of a sentence requires reasoning about the knowledge states and intentions of the communication partners. To use a worn-out example, Grice (1975) points out that a sentence like (1b), if uttered to the owner of an immobilized car by a passerby, carries much more information than what is literally said.

- (1) a. *car owner*: I am out of petrol.
b. *passerby*: There is a garage round the corner.

The listener will legitimately infer from this exchange that the garage is open and sells petrol (or at least so the passerby believes). This follows from his expectations that the passerby understands his situation and is cooperative. But this is not sufficient. The passerby must be aware of these expectations, and he must believe that the car owner is capable to make the required inference, the car owner must be able to assign these epistemic qualities to the passerby etc.

The standard tool to model this kind of inferences about mutual belief states is epistemic logic, and it has been extensively used to model the interface between semantics and pragmatics (see for instance Stone 1998; Gerbrandy 1999). *Game Theory* is another

intellectual tradition that aims at a mathematical model of the inferences that are involved in the interaction of rational agents. Unlike epistemic logic, game theory puts its focus on the decisions and preferences of the interacting agents, rather than on explicit models of their internal states and reasoning processes.¹ As communicative goals and decisions play a crucial role for pragmatics, the use of game theoretic methods in pragmatics suggests itself.

The application of game theory to linguistic issues traces back to David Lewis' dissertation (Lewis 1969), where he showed that the game theoretic notion of a *Nash equilibrium* is apt to explain how a linguistic convention can be self-sustaining in a community. While the game theoretic investigation of communication was mostly pursued by economists and biologists in the quarter century after Lewis' work (see for instance Spence 1973; Crawford and Sobel 1982; Maynard Smith 1991), since the mid-nineties formal linguists and philosophers of language have paid increasing attention to the potential of game theory to the analysis of language use and language evolution (see for instance Dekker and van Rooij 2000; Parikh 2001; van Rooij 2004; Skyrms 2010 or the papers in Benz et al. 2005; Pietarinen 2007).

In this article, I will not even attempt to give a comprehensive overview on the various applications of game theory to semantics and pragmatics that have emerged in recent years. Rather, I will describe a particular approach in some detail. The *Iterated Best Response Model* (*IBR Model* for short) of game theoretic pragmatics is an implementation of the neo-Gricean program of pragmatics to derive the mapping from *what is said* to *what is meant* from first principles of rational communication. Various incarnations of it have been proposed in recent years (cf. Jäger 2007c; Franke 2008; Jäger and Ebert 2009; Franke 2009; see also Benz and van Rooij 2007 for a related approach), drawing on earlier work in economics like Rabin (1990) that was not specifically directed at the analysis of natural language. In this article I will largely follow the version from Franke (2009). In particular, the analyses of disjunction, free choice permissions and embedded implicatures given below are taken from there.

2 Game construction I: Interpretation games

Following the lead of Lewis (1969), the present approach models linguistic communication as a *signaling game*. This is a class of dynamic games involving two players, a *sender* and a *receiver*. The game proceeds in three stages:

1. The sender is assigned some piece of private information, its *type*, that is not revealed to the receiver.
2. The sender transmits a *signal* to the receiver. The choice of the signal may depend on the type.

¹The two approaches are by no means independent of each other. There is a long history of mutual inspiration of the two communities.

3. The receiver chooses a certain *action*, possibly dependent on the observed signal.

To keep the mathematical analysis manageable, I will assume that the number of possible types, signals and actions are finite. The type of the sender is assigned according to a certain *prior* probability distribution that is common knowledge between the players.

A *history* of such a game is a type-signal-action triple. Both players have certain preferences over histories. These preferences are modeled as numerical *utilities*. It is important to keep in mind that the numerical values used are to some degree arbitrary. What is important is only the relative preferences between histories (or probability distributions over histories) that are captured by such a utility function. Unless otherwise noted, it is always tacitly assumed that both players are *rational*. This entails that they have a consistent epistemic state, and that they make their choices in a way that maximizes their *expected utilities*, given their epistemic state.

Unlike Lewis (1969) and much subsequent work, the IBR model is not concerned with the question how signals acquire meaning in a signaling game. Instead it is presupposed that signals already have a conventionalized meaning that is common knowledge between the players, and the model serves to study how the conventionalized meaning affects the way signals are used in particular strategic situations. Therefore we augment signaling games with a relation of truth that holds between types and signals.

Let us make the model formally precise. A signaling game consists of the following components:

- a finite set \mathcal{S} of *signals* that the sender has at her disposal,
- a finite set \mathcal{T} of *types* (information states) the sender might be in,
- a *prior probability distribution* p^* over \mathcal{T} ,
- a *truth* relation \models between \mathcal{T} and \mathcal{S} ,
- a set \mathcal{A} of *actions* that the receiver may take, and
- utility functions u_s and u_r for the sender and the receiver respectively that both map triples from $\mathcal{T} \times \mathcal{S} \times \mathcal{A}$ to real numbers.

Intuitively, a type $t \in \mathcal{T}$ specifies the information that the sender wants to communicate in a given situation, and \mathcal{S} is the set of linguistic expressions that she could, in principle, utter. It would be unwieldy though to identify \mathcal{S} with all well-formed expressions of English (or some other natural language), and to pin down the set of possible information states of the sender seems to be outright hopeless. To analyze a particular linguistic example, we have to construct a local game that captures the relevant alternatives that the interlocutors take into consideration. In the sequel, I will present a recipe how to construct such a game. It closely follows the proposal from Franke (2009).

To illustrate the recipe, I will use the standard example for a scalar implicature:

- (2) Some boys came to the party.

If the listener is interested in the question how many boys came to the party, (2) carries the implicature that not all boys came to the party. Neo-Gricean pragmaticists usually assume (cf. for instance Horn 1984) that this effect arises because the word *some* in (2) is element of a scale $\langle no, some, all \rangle$. When interpreting (2), the listener takes into consideration that the speaker could have uttered (3a) or (b):

- (3) a. No boy came to the party.
 b. All boys came to the party.

When constructing a game for a particular expression s that is uttered in a particular context c , the set of signals \mathcal{S} is simply $ALT_c(s)$, the set of expression alternatives to s in c . This set may be determined by lexical properties of the linguistic material involved (as in the case of scalar expressions), by information structure (cf. Rooth 1985; Büring 1995) by an explicit or implicit question and perhaps by other factors. The issue has been discussed extensively in the literature and I will not dwell further on it here.² For the running example, $\mathcal{S} = \{(2), (3a), (3b)\}$, mnemonically named {SOME, NO, ALL} in the sequel.

The set of \mathcal{T} of *types* or information states that the sender may have can be identified with the set of answers to the current *question under discussion* that — according to the listener — the speaker might know. For the time being, I will assume the *strong competence assumption*³, i.e. the idea that the speaker is perfectly informed about the issue at hand.⁴ This means that the speaker is assumed to know the truth value of all elements of \mathcal{S} .⁵ If the speaker is also rational, we can identify her possible information states with the set of maximally consistent sets of the elements of \mathcal{S} and their negations that are consistent with the context and contain at least one non-negated element. To avoid the formal complication of defining an appropriate negation for whatever representation language we use, information states are defined as non-empty sets of sentences that are deductively closed within \mathcal{S} , with the implicit assumption that a sentence which is not a member of an information state is false in this state (the so-called *closed world assumption*).

In the example at hand, these would be (under the assumptions that {NO, SOME} is inconsistent and $ALL \vdash SOME$, where \vdash indicates entailment):

$$\mathcal{T} = \{\{NO\}, \{SOME\}, \{SOME, ALL\}\}.$$

Types are thus essentially *possible worlds* (provided the competence assumption holds). In our running example, I will denote the three types with $w_{\neg\exists}$, $w_{\exists\rightarrow\forall}$, and w_{\forall} respectively.

The *prior probability distribution* p^* over \mathcal{T} captures the subjective probabilities that the receiver assigns to the possible states of the sender. It has no obvious counterparts

²See for instance Fox (2007); Katzir (2007) for some recent discussion.

³The name “competence assumption” is due to van Rooij and Schulz 2004, but the notion has been used by many neo-Gricean researchers under varying names. I call it “strong competence assumption” here because I will later introduce a weaker variant thereof.

⁴The consequences of lifting this assumptions are explored in Section 6

⁵It is actually sufficient to assume that the speaker has a definite opinion about these truth values. I will continue to pretend that the speaker only has correct information, but nothing hinges on this.

in the epistemic models that are commonly used in formal semantics and pragmatics. Unless any specific information is given, we have no reason to assume that the receiver considers a certain type as being more likely than any other type. In such a situation the *principle of insufficient reason* (cf. Jaynes 2003) applies: If two possibilities a and b are indistinguishable except for their names, they have the same subjective probability. This entails that $p^*(t) = 1/|\mathcal{T}|$ for all types t . In our example, this means:

$$p^*(w_{\neg\exists}) = p^*(w_{\exists\neg\forall}) = p^*(w_{\forall}) = \frac{1}{3}.$$

A signal s is *true* in a type t iff $s \in t$. We thus have:

$$\begin{aligned} w_{\neg\exists} & \models \text{NO} \\ w_{\exists\neg\forall} & \models \text{SOME} \\ w_{\forall} & \models \text{SOME, ALL} \end{aligned}$$

In an interpretation game, the receiver’s task is to figure out which type the sender has. The set \mathcal{A} of receiver *actions* can thus be identified with \mathcal{T} . (I will continue to use the letters a and \mathcal{A} when talking about receiver’s options, and t/\mathcal{T} when referring to information states.)

It is furthermore assumed that both players have an overarching interest in successful communication. Formally this cashes out as the postulate that

$$t' \neq t \Rightarrow \forall s. u_{s/r}(t, s, t) > u_{s/r}(t, s, t').$$

Players have a secondary interest in avoiding complexity. For the sender this means that she prefers short or otherwise unmarked expressions over long and marked ones. I will refrain from spelling this out in more detail. The fact that there is differential signal complexity should be uncontroversial; it underlies Grice’s Maxim of Manner as well as more recent implementations like Bidirectional Optimality Theory (cf. Blutner 2001). Pinning it down in a general way is not trivial, but I will confine myself to clear examples. Formally this is captured by a cost function c_s that maps type-signal pairs to real numbers. The assumption to make signal costs dependent on types is motivated by regularities like TOPIC PRECEDES COMMENT that indicate that preferences over expressions may depend on the meaning that the speaker wants to express.

Likewise, the receiver may have, *ceteris paribus*, differential preferences between different interpretations of an expression — like a dispreference for coercion or type shift, a preference for presupposition binding over accommodation etc. (cf. van der Sandt 1992). This is implemented by a cost function c_r that maps signal-action pairs to real numbers.⁶ In all examples discussed in this chapter, receiver costs do not properly depend on signals but only on actions.

⁶To my knowledge, receiver’s costs as well as type-dependent sender’s costs were first introduced in Franke and Jäger (2010).

These assumptions are implemented by the following utility functions:

$$u_s(t, s, a) = -c_s(t, s) + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else,} \end{cases} \quad (1)$$

$$u_r(t, s, a) = -c_r(s, a) + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else.} \end{cases} \quad (2)$$

Unless otherwise noted, it is always assumed that costs are *nominal*. This means that costs are extremely low in comparison to the payoff that is to be gained from getting the type across correctly. They only ever play a role when comparing options that are otherwise equally successful for communication. Formally this can be ensured by the following postulate:⁷

(4) NOMINAL COSTS

$$\begin{aligned} 0 &\leq c_s(t, s) < |\mathcal{A}|^{-2}, \\ 0 &\leq c_r(s, a) < |\mathcal{S}|^{-2}. \end{aligned}$$

In our running example differential costs play no role, and we thus assume that all costs are 0.

The game construction for an example sentence, represented as signal s , is summarized in the following definition. (An inference relation is assumed to be given, and a set of sentences is consistent if not everything can be inferred from it.)

Definition 1 ((Strong) Interpretation Game) *Let s be an expression that is uttered in a context ct (where a context is a set of sentences the truth of which is common knowledge between the interlocutors), and that has a set of alternatives $ALT(s)$ such that each element thereof is consistent with ct . The strong interpretation game $\mathcal{G}_s^* = \langle \mathcal{T}, p^*, \mathcal{S}, \mathcal{A}, u_s, u_r \rangle$ for*

⁷The reason for this particular definition is the following: It will be motivated in the next section that we only need to consider strategies where the probability that a sender has to achieve successful information transmission with a certain signal s is either 0 or $1/n$, for some positive integer $n \leq |\mathcal{A}|$. The latter occurs if the receiver maps s to one out of n different actions. The minimal difference between the success chance of two different signals is thus $1/n - 1/(n-1)$, which is always larger than $|\mathcal{A}|^{-2}$. The argument for the receiver is analogous.

s is constructed as follows, where c_s and c_r are nominal cost functions:

$$\begin{aligned}
\mathcal{T} &= \{t \subseteq \mathcal{S} \mid t \neq \emptyset \wedge t \cap ct \text{ is consistent} \wedge \\
&\quad \forall s' \in \mathcal{S}. t \vdash s \Rightarrow s \in t\}, \\
p^*(t) &= \frac{1}{|\mathcal{T}|}, \\
\mathcal{S} &= ALT(s), \\
t \models s' &\Leftrightarrow s' \in t, \\
\mathcal{A} &= \mathcal{T}, \\
u_s(t, s', a) &= -c_s(t, s') + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else,} \end{cases} \\
u_r(t, s', a) &= -c_r(s', a) + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else.} \end{cases}
\end{aligned}$$

3 Strategies and best responses

A (behavioral) strategy⁸ for the sender is a probabilistic function from types to signals. Formally, the set of sender strategies Σ is defined as

$$\Sigma = \mathcal{T} \mapsto \Delta(\mathcal{S}),$$

(where $\Delta(X)$ is the set of probability distributions over the set X , i.e. $p \in \Delta(X) \Leftrightarrow p \in X \mapsto [0, 1] \wedge \sum_{x \in X} p(x) = 1$). We write $\sigma(s|t)$ rather than $\sigma(t, s)$ to stress the intuition that this is the probability that the sender uses signal s provided she is in type t .

Likewise, the set P (speak: *Rho*) of receiver strategies consists of the probabilistic functions from \mathcal{S} to \mathcal{A} :

$$P = \mathcal{S} \mapsto \Delta(\mathcal{A}),$$

where $\rho(a|s)$ is the probability that, upon observing signal s , the receiver will perform action a .

It is important to note that these probabilities are again subjective probabilities. $\sigma \in \Sigma$ does not model the plan of the sender, but rather the expectations the receiver has about the sender's behavioral dispositions (and likewise for $\rho \in P$).

If the receiver observes signal s and has a model σ of the receiver, she can estimate the utility that she can expect for each signal. This is captured by the notion of *expected utility*:

$$EU_s(s|t; \rho) = \sum_{a \in \mathcal{A}} \rho(a|s) u_s(t, s, a).$$

⁸The game theoretic literature distinguishes *pure*, *mixed*, and *behavioral* strategies for such games. For our purposes behavioral strategies are the most relevant variety, and I will use the term *strategy* in this sense throughout.

For the utility function defined above, this simplifies to

$$EU_s(s|t; \rho) = \rho(t|s) - c_s(t, s).$$

To estimate the receiver's expected utility upon observing a signal s given a model σ of the sender, one needs the probability of the types in \mathcal{T} conditional on observing s . σ provides the exact opposite, the conditional probability of signals given types. These two conditional probabilities are related via *Bayes' Rule*:

$$\sigma(t|s) = \frac{\sigma(s|t)p^*(t)}{\sum_{t' \in \mathcal{T}} \sigma(s|t')p^*(t')}.$$

If p^* is a uniform distribution, this simplifies to

$$\sigma(t|s) = \frac{\sigma(s|t)}{\sum_{t' \in \mathcal{T}} \sigma(s|t')}.$$

This is only defined if the denominator is > 0 . If it does equal 0, σ assigns probability 0 to the signal s in all types t' . This would be a scenario where the receiver observes a signal that contradicts his model of the sender's behavior. Such signals are called *surprise signals*. They force the receiver to *revise his beliefs* in a way that accommodates the unexpected observation. The issue of the appropriate belief revision policy is a major issue in epistemic game theory, and I will not discuss it here at any length (see for instance Battigalli and Siniscalchi 2002). In the context of this chapter, I will adopt a rather simple version. If the receiver observes a surprise signal, he gives up all his assumptions about the sender's behavior and about costs of signals. He does preserve his knowledge about the literal interpretation of signals and about prior probabilities though. This is implemented in the following way:

$$\sigma(t|s) = \begin{cases} \frac{\sigma(s|t)}{\sum_{t' \in \mathcal{T}} \sigma(s|t')} & \text{if } \sum_{t' \in \mathcal{T}} \sigma(s|t') > 0 \\ \text{undefined} & \text{else.} \end{cases}$$

The receiver's expected utility given a model σ of the sender thus comes out as

$$EU_r(a|s; \sigma) = \begin{cases} \sum_{t \in \mathcal{T}} \sigma(t|s)u_r(t, s, a) & \text{if defined,} \\ \text{undefined} & \text{else.} \end{cases}$$

For the utility function given above, this simplifies to

$$EU_r(a|s; \sigma) = \begin{cases} \sigma(a|s) - c_r(s, a) & \text{if defined,} \\ \text{undefined} & \text{else.} \end{cases}$$

Rational players will play a *best response* to their model of the other player, i.e. they will always choose an option that maximizes their expected utility. If the expected utility is not defined (i.e. if the receiver observes a surprise signal), the receiver takes resort to

hypothesizing that the observed signal is true. This taken into account, best responses are defined as:

$$br_s(t, \rho) = \{s | EU_s(s|t; \rho) = \max_{s' \in \mathcal{S}} EU_s(s'|t; \rho)\},$$

$$br_r(s, \sigma) = \begin{cases} \{a | EU_r(a|s; \sigma) = \max_{a' \in \mathcal{A}} EU_r(a'|s; \sigma)\} & \text{if } EU_r(\cdot|s; \sigma) \text{ is defined,} \\ \{a | a \models s\} & \text{else.} \end{cases}$$

The *best response* (singular definite) of a player to a model of the other player is the behavioral strategy that assigns equal probabilities to all best responses (due to the principle of insufficient reason), and 0 to all sub-optimal responses.

$$BR_s(\rho) = \sigma \quad \text{iff} \quad \sigma(s|t) = \begin{cases} \frac{1}{|br_s(t, \rho)|} & \text{if } s \in br_s(t, \rho), \\ 0 & \text{else,} \end{cases}$$

$$BR_r(\sigma) = \rho \quad \text{iff} \quad \rho(a|s) = \begin{cases} \frac{1}{|br_r(s, \sigma)|} & \text{if } a \in br_r(s, \sigma), \\ 0 & \text{else.} \end{cases}$$

Let me illustrate these notions with the running example. Suppose the receiver assumes the sender to be absolutely honest, and this is the only assumption he makes about her. A model of such a sender would be σ_0 , who sends some true signal in each state. Applying the principle of insufficient reason again, a receiver who expects the sender to be honest but makes no further assumptions about her will assign her the following strategy:

$$\sigma_0(s|t) = \begin{cases} \frac{1}{|\{s|t \models s\}|} & \text{if } t \models s, \\ 0 & \text{else.} \end{cases}$$

In our little example, σ_0 would look as in Table 1.

	NO	SOME	ALL
$w_{\neg\exists}$	1	0	0
$w_{\exists\neg\forall}$	0	1	0
w_{\forall}	0	$\frac{1}{2}$	$\frac{1}{2}$

Table 1: Honest sender σ_0

The expected utilities of the receiver against the honest sender's strategy σ_0 come out as given in Table 2 (which is the result of flipping Table 1 along the diagonal and normalizing row-wise).

Consider $\rho_0 = BR_r(\sigma_0)$, the best response of the receiver against σ_0 from Table 1. In each type, the receiver will maximize his expected utility. So in each row in Table 2, he will choose a cell that is maximal within this row. If there were several such cells, ρ_0 would put

	$w_{\neg\exists}$	$w_{\exists\neg\forall}$	w_{\forall}
NO	1	0	0
SOME	0	$\frac{2}{3}$	$\frac{1}{3}$
ALL	0	0	1

Table 2: $EU_r(\cdot|\cdot; \sigma_0)$

	$w_{\neg\exists}$	$w_{\exists\neg\forall}$	w_{\forall}
NO	1	0	0
SOME	0	1	0
ALL	0	0	1

Table 3: $\rho_0 = BR_r(\sigma_0)$

equal probability to each of it. In the example, the maxima are unique and thus receive a probability 1, as shown in Table 3.

Computing the sender's expected utility against ρ_0 also amounts to simply flipping the table along the diagonal, and the best response to ρ_0 puts all probability mass into the cells with a maximal expected utility within its row (see Table 4).

	NO	SOME	ALL		NO	SOME	ALL
$w_{\neg\exists}$	1	0	0	$w_{\neg\exists}$	1	0	0
$w_{\exists\neg\forall}$	0	1	0	$w_{\exists\neg\forall}$	0	1	0
w_{\forall}	0	0	1	w_{\forall}	0	0	1

Table 4: $EU_s(\cdot|\cdot; \rho_0), BR_s(\rho_0)$

It is easy to see that $BR_r(BR_s(\rho_0)) = \rho_0$, so iterating the best response operation will not alter the strategies anymore.

4 Iterated Best Response

The notion of a best response models the behavior of rational players, given certain (perhaps probabilistic) assumptions about the behavior of the other player. But how are these assumptions formed? It might seem reasonable to assume that one's opponents are also rational. This means that they play a best response to their model of one's own behavior.

This answer apparently only shifts the problem one step further, and this could be repeated *ad infinitum*.

The standard solution concept in classical game theory is agnostic with regard to the question where beliefs about other players come from. It is only required that beliefs are consistent with the assumption that everybody is rational (and that this fact is common knowledge). If I believe that my opponent has a correct model of my behavior and is rational, I will ascribe to them a best response to my behavior and they ascribe to me a best response to my behavior. Such a state is self-reinforcing. It is called a *Nash equilibrium*. There is no simple answer to the question though how players find such an equilibrium in the first place. This problem is aggravated by the fact that many games have many equilibria. In our example, the 1-1 map between types and signals from Table 3 (together with the best response to it) constitutes an equilibrium, but so do the other five 1-1 maps, plus an infinity of properly stochastic strategy pairs.

In the present context the equilibrium concept is thus problematic because it disregards the literal meaning of signals — any 1-1 map between meanings and signals is as good as any other. Also, results from behavioral experiments suggest that people are actually not very good at finding Nash equilibria, even if there is only one. This can be illustrated with the game “*p*-beauty contest”.⁹ In this game, each of n test persons is asked to pick a (real) number between 0 and 100. The winner is the one whose choice is closest to p times the average of all numbers from the group, where p is some number that is announced in advance. Suppose $p = 0.7$. Then the only Nash equilibrium of the game is that everybody picks 0 (because for any other average, one could benefit by undercutting this average). This game has been tested experimentally many times, with very revealing results. For instance, Camerer and Fehr (2006) report that participants chose numbers from the entire interval $[0, 100]$, with peaks around 50, 35 and 25. Their interpretation is that the participants fall into three groups. The irrational players simply choose a random number without any strategic thinking. (The authors do not interpret the peak at 50, but it seems to me that this is just a variant of irrational behavior were 50 is chosen because it is located nicely in the center of the interval). Then there are *level-1 players*. They anticipate that many people are irrational, and that their average choice is 50. The best response to this is to pick $50 \times 0.7 = 35$. Finally, *level-2 players* expect there to be level-1 players and play a best response to their choice. This would be $35 \times 0.7 = 24.5$ and corresponds to the peak at 25.

Similar results have been found time and again by other researchers. (See Camerer 2003 for further information and references.) It is also interesting that the average choice quickly converges towards 0 if the game is repeated and the players get feedback about the average outcome after each round (as reported in Slonim 2005). So repeated playing seems to foster level- k thinking for levels > 2 .

These considerations lead to the following model: There is a default strategy that people

⁹It is named after a remark by John Maynard Keynes whereas the stock market functions like a beauty contest, where investors try to anticipate which stocks the other investors will find attractive to buy them early and sell them with a gain.

follow if they do not deliberate. In the beauty contest game, this is the strategy to pick randomly from the available numbers. Furthermore, there is a *cognitive hierarchy* of rational player types of ever higher sophistication. *Level-1* players play a best response to the default strategy, *level-2* players a best response to *level-1* players and so on. (Note that the level of a player does not necessarily measure their ability to perform nested reasoning. It is also bounded by their assumption about the level of their opponent.)

This architecture is readily applicable to interpretation games. The default strategy of a non-deliberating player is the honest sender σ_0 that sends each true signal with the same probability in each type.¹⁰ A *level-0 receiver* plays the best response ρ_0 to σ_0 , a *level-1 sender* plays the best response to ρ_0 and so on. This is captured by the notion of an *IBR sequence* (where *IBR* abbreviates *Iterated Best Response*).

Definition 2 (IBR sequence)

$$\begin{aligned} \sigma_0(s|t) &= \begin{cases} \frac{1}{|\{s|t \models s\}|} & \text{if } t \models s, \\ 0 & \text{else,} \end{cases} \\ \rho_n &= BR_r(\sigma_n), \\ \sigma_{n+1} &= BR_s(\rho_n). \end{aligned}$$

In the running example σ_1 is given in the left panel of Table 4. For all $n > 1$, $\sigma_n = \sigma_1$ and $\rho_n = \rho_0$. The strategy pair (σ_1, ρ_0) is a fixed point of the IBR sequence and thus a Nash equilibrium. It can be shown that with a utility function as given in Equation (1) and nominal costs, the IBR sequence always reaches a fixed point (see Appendix).

Let s be an expression and $\mathcal{G}(s)$ the interpretation game constructed for s . Let (σ^*, ρ^*) be the fixed point of the IBR sequence. The IBR model then predicts that s receives as possible readings the set $\{a | \rho^*(a|s) > 0\}$. For the running example, the prediction is thus that *Some boys came to the party* will be interpreted as entailing that not all boys came to the party. As this game is constructed from expression (2), it does not serve to make predictions about the interpretation of the alternative expressions (even though in this example these predictions would be correct). Also, the model is not designed to make predictions about speaker behavior, so it is not claimed that a speaker who wants to express that some but not all boys came will use expression (2).

5 Examples

The previous example illustrated how scalar implicatures are derived in the IBR model. Here is another example of a quantity implicature:

- (5) (Who of Ann and Bert came to the party?) Ann came to the party.

¹⁰Jäger and Ebert (2009) and Franke (2009) also consider a default strategy for the receiver where each signal is assigned a uniform distribution over types where this signal is true. This version of the model leads to unwelcome predictions in connection with I-implicatures though. In example (7), starting with the receiver would lead to a fixed point where the I-implicature does not emerge.

The question defines a set of alternative answers?

- (6) a. Ann came to the party. (= ANN)
 b. Bert came to the party. (= BERT)
 c. Neither came to the party. (= NEITHER)
 d. Both came to the party. (= BOTH)

They define four possible worlds (i.e. maximally consistent sets of alternatives or their negations): Only Ann came (w_a), only Bert came (w_b), neither came (w_\emptyset) or both came ($w_{a,b}$).

σ_0 is constructed in the following way:

- Draw a table with types as rows and expressions as columns.
- Put a 1 in each cell where the column-expression is true in the row-world, and 0 everywhere else.
- Divide each row by its sum.

The result is shown in Table 5.

	ANN	BERT	NEITHER	BOTH
w_a	1	0	0	0
w_b	0	1	0	0
w_\emptyset	0	0	1	0
w_{ab}	$\frac{1}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$

Table 5: σ_0

The table for the best response $BR_r(\sigma)$ to some sender strategy σ can be constructed by the following operations:

- Flip the table for σ along the main diagonal.
- If a row consists only of 0s, replace each cell corresponding to a true type-signal association with 1.
- Otherwise subtract receiver costs from each cell, replace each cell that is maximal within its row by 1, and every other cell by 0.
- Divide each row by its sum.

Note that the best response to a surprise signal (a column with only 0s in a receiver strategy) only makes use of the literal meaning of that signal and disregards costs. This reflects the fact that surprise signals enforce belief revision. In the present model, speakers only preserve their knowledge about the literal meanings of signals under belief revision and delete all other information.

The table for the best response $BR_s(\rho)$ to some receiver strategy ρ can be computed by the following steps:

- Flip the table for ρ along the main diagonal.
- Subtract sender costs from each cell.
- Replace each cell that is maximal within its row by 1, and all other cells by 0.
- Divide each row by its sum.

The strategy $\sigma_1 = BR_s(\rho_0)$ is given in Table 6. (σ_1, ρ_0) form a fixed point. Hence the

ρ_0	w_a	w_b	w_\emptyset	w_{ab}	σ_1	ANN	BERT	NEITHER	BOTH
ANN	1	0	0	0	w_a	1	0	0	0
BERT	0	1	0	0	w_b	0	1	0	0
NEITHER	0	0	1	0	w_\emptyset	0	0	1	0
BOTH	0	0	0	1	w_{ab}	0	0	0	1

Table 6:

model predicts that (5) is interpreted exhaustively, i.e. in the sense that Ann but not Bert came to the party.

The two examples illustrate the general pattern of *quantity implicatures*. If an expression A has an alternative B that is logically stronger than A , then in the fixed point A will be interpreted exhaustively, i.e. as entailing “ A and not B ”.

This inference may be blocked though if the alternative set already contains an expression which literally denotes “ A and not B ”. This may be illustrated with the next example.

- (7) a. John opened the door. (= OPEN)
 b. John opened the door using the handle. (= OPEN-H)
 c. John opened the door with an axe. (= OPEN-A)

Normally (7a) will be interpreted as meaning (7b) rather than (c). This is an instance of an inference type that Levinson (2000) calls *I-implicatures* — underspecified expressions tend to be interpreted as the prototypical elements of their literal denotation.

c_s, c_r	OPEN	OPEN-H	OPEN-A
w_h	0, 0	$\frac{1}{20}, 0$	$\frac{1}{20}, 0$
w_a	$0, \frac{1}{20}$	$\frac{1}{20}, \frac{1}{20}$	$\frac{1}{20}, \frac{1}{20}$

Table 7: Cost functions

Suppose (b) and (c) are expression alternatives in the interpretation game for (a), and suppose that the only ways to open a door is with the handle or with an axe. (This example is a bit more abstract than the previous two, first because the set of alternatives along the lines of (7c) is actually open-ended, and second because these alternatives are arguably not activated in the same way when interpreting (a) as for instance the scalar alternatives are in the first example. Nevertheless the example is instructive to illustrate the logic of the IBR model.) Under these assumptions, there are just two possible worlds: w_h where John opens the door with the handle, and w_a where he uses an axe. (Recall that the possible worlds are maximally consistent sets of alternatives or negations thereof *that contain at least one non-negated expression*. Therefore the alternatives do not generate a world where John does not open the door at all.)

In this example costs play a role. The implicature arises because (7a) is shorter than both (b) and (c), and because w_h is more prototypical than w_a . Let us thus say that the costs are as given in Table 7, where the first number in each cell give the sender’s cost and the second the receiver’s. The actual choice of numbers is inessential for the IBR sequence. As long as costs are sufficiently small, only the relative ranking of the different alternative types and expressions according to their costs matters.

The construction of σ_0 follows the same procedure as explained above.

The IBR sequence for example (7) is given in Table 8. The fixed point is (σ_1, ρ_1) . The signal OPEN is interpreted as the proposition $\{w_h\}$, i.e. as expressing that John opened the door using the handle.

Note that this outcome depends both on differential costs between the different signals and the fact that the more specific signals jointly exhaust the meaning of the more general signal. If OPEN would not entail OPEN-H \vee OPEN-A, there would be a third type w^* where John opened the door, but neither with the handle nor with an axe, and OPEN would be interpreted as $\{w^*\}$, i.e. we would observe a quantity implicature rather than an I-implicature.

A similar but slightly more complex inference pattern has been called *M-implicature* or *division of pragmatic labor* (Horn 1984). An example is given in (8).

- (8) a. John stopped the car. (= STOP)
b. John made the car stop. (= MAKE-STOP)

Normally (8a) is interpreted as referring to an event where John stopped the car in the usual way, i.e. using the foot brake, while (b) rather conveys that he stopped it in a less

σ_0	OPEN	OPEN-H	OPEN-A	ρ_0	w_h	w_a
w_h	$\frac{1}{2}$	$\frac{1}{2}$	0	OPEN	1	0
w_a	$\frac{1}{2}$	0	$\frac{1}{2}$	OPEN-H	1	0
				OPEN-A	0	1

σ_1	OPEN	OPEN-H	OPEN-A	ρ_1	w_h	w_a
w_h	1	0	0	OPEN	1	0
w_a	0	0	1	OPEN-H	1	0
				OPEN-A	0	1

Table 8: IBR sequence

conventional way, like using the hand brake or driving into the ditch.

The following game is again more of an illustration of the working of the IBR model than a serious attempt to analyze this particular linguistic example. Suppose there are two possible worlds: in w_1 John stops the car using the foot brake, and in w_2 he drives it into the ditch. These are the only ways to stop the car. Both STOP and MAKE-STOP are true in both worlds. The receiver has, *ceteris paribus*, a preference for w_1 over w_2 , which is mirrored by the fact the latter has higher receiver costs than the former. Likewise, the sender has a preference for STOP over MAKE-STOP, which is mirrored in the cost function. The cost functions are given in Table 9.

c_s, c_r	STOP	MAKE-STOP
w_1	0, 0	$\frac{1}{10}, 0$
w_2	$0, \frac{1}{10}$	$\frac{1}{10}, \frac{1}{10}$

Table 9: Cost functions

Table 10 shows the IBR sequence for this game. In the fixed point (σ_2, ρ_2) , STOP is interpreted as $\{w_1\}$ and MAKE-STOP as $\{w_2\}$.

6 Game construction II: Lifted games

In its present version, our model assumes a very strong version of the competence assumption: it is hard-wired into the model that the speaker knows the complete answer to the question under discussion. This is not just unrealistic, it also leads to unwarranted predictions. Consider (9), the standard example for a quantity implicature.

σ_0	STOP	MAKE-STOP		ρ_0	w_1	w_2
w_1	$\frac{1}{2}$	$\frac{1}{2}$		STOP	1	0
w_2	$\frac{1}{2}$	$\frac{1}{2}$		MAKE-STOP	1	0
σ_1	STOP	MAKE-STOP		ρ_1	w_1	w_2
w_1	1	0		STOP	1	0
w_2	1	0		MAKE-STOP	$\frac{1}{2}$	$\frac{1}{2}$
σ_2	STOP	MAKE-STOP		ρ_2	w_1	w_2
w_1	1	0		STOP	1	0
w_2	0	1		MAKE-STOP	0	1

Table 10: IBR sequence

- (9) a. Ann or Bert showed up. (= OR)
b. Ann showed up. (= A)
c. Bert showed up. (= B)
d. Ann and Bert showed up. (= AND)

Since Grice (1975) it is a common assumption that *or* literally denotes inclusive disjunction, and that the exclusive interpretation is due to a quantity implicature, arising from the competition with *and*. If we assume that the alternative set to (9) is (9a–d), we end up with a game that is represented by σ_0 in Table 11. In the fixed point (σ_1, ρ_1) , OR receives the inclusive interpretation. What is even more odd is that OR is a surprise signal, because it is never used in σ_1 . This is a consequence of the strong competence assumption. For each of the three maximally informative information states assumed here, there is a more specific signal, so OR is in fact superfluous in this game.

To adequately deal with examples like this, we have to take the possibility into account that the sender is underinformed. This can be accommodated if we identify information states, i.e. types, with non-empty *sets* of possible worlds (as it is standard practice since Stalnaker 1978). Franke (2009) calls this kind of game *epistemically lifted games*. So if W is the set of worlds that is constructed from a set of alternative expressions,

$$\mathcal{T} = POW(W) - \{\emptyset\}.$$

An expression is considered true in a state if it is true in all possible worlds in this state, i.e.

$$t \models s \text{ iff } \forall w \in t. w \models s.$$

σ_0	OR	A	B	AND	ρ_0	w_a	w_b	w_{ab}
w_a	$\frac{1}{2}$	$\frac{1}{2}$	0	0	OR	$\frac{1}{2}$	$\frac{1}{2}$	0
w_b	$\frac{1}{2}$	0	$\frac{1}{2}$	0	A	1	0	0
w_{ab}	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	B	0	1	0
					AND	0	0	1
σ_1	OR	A	B	AND	ρ_1	w_a	w_b	w_{ab}
w_a	0	1	0	0	OR	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
w_b	0	0	1	0	A	1	0	0
w_{ab}	0	0	0	1	B	0	1	0
					AND	0	0	1

Table 11: IBR sequence

The competence assumption is not completely given up though. A weaker version of it is implemented via the receiver’s cost function. Everything else being equal, the receiver has a preference for more specific interpretations (cf. Dalrymple et al’s 1994 *Strongest Meaning Hypothesis*). This corresponds to the *weak competence assumption*

$$a_1 \subsetneq a_2 \Rightarrow \forall s. c_r(s, a_1) < c_r(s, a_2).$$

For concreteness’ sake, let us say that

$$\forall s. c(s, a) \propto |a|.$$

Applied to example (9), this leads to the following cost functions (It is assumed that AND and OR are more complex than A and B for the sender, but nothing hinges on this in this example.):

$$\begin{aligned} \forall t. c_s(t, A) = c_s(t, B) &= 0, \\ \forall t. c_s(t, \text{AND}) = c_s(t, \text{OR}) &= \frac{1}{100}, \\ \forall s. c_r(s, a) &= \frac{|a|}{1000}. \end{aligned}$$

The IBR sequence is given in Table 12.

The fixed point here is (σ_1, ρ_0) . Two things are noteworthy in this example. First, the construction correctly predicts that OR receives an exclusive interpretation. Second, the usage of OR triggers an IGNORANCE IMPLICATURE. In the fixed point strategy, the receiver infers from observing OR that the sender does not know A for sure, and she doesn’t know B for sure either.

The game construction for lifted interpretation games is summarized in the following definition. These games are called *weak interpretation games* because they implement the

σ_0	OR	A	B	AND
$\{w_a\}$	$\frac{1}{2}$	$\frac{1}{2}$	0	0
$\{w_b\}$	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$\{w_{ab}\}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
$\{w_a, w_b\}$	1	0	0	0
$\{w_a, w_{ab}\}$	$\frac{1}{2}$	$\frac{1}{2}$	0	0
$\{w_b, w_{ab}\}$	$\frac{1}{2}$	0	$\frac{1}{2}$	0
$\{w_a, w_b, w_{ab}\}$	1	0	0	0

ρ_0	$\{w_a\}$	$\{w_b\}$	$\{w_{ab}\}$	$\{w_a, w_b\}$	$\{w_a, w_{ab}\}$	$\{w_b, w_{ab}\}$	$\{w_a, w_b, w_{ab}\}$
OR	0	0	0	1	0	0	0
A	1	0	0	0	0	0	0
B	0	1	0	0	0	0	0
AND	0	0	1	0	0	0	0

σ_1	OR	A	B	AND
$\{w_a\}$	0	1	0	0
$\{w_b\}$	0	0	1	0
$\{w_{ab}\}$	0	0	0	1
$\{w_a, w_b\}$	1	0	0	0
$\{w_a, w_{ab}\}$	0	$\frac{1}{2}$	$\frac{1}{2}$	0
$\{w_b, w_{ab}\}$	0	$\frac{1}{2}$	$\frac{1}{2}$	0
$\{w_a, w_b, w_{ab}\}$	0	$\frac{1}{2}$	$\frac{1}{2}$	0

Table 12: IBR sequence in the epistemically lifted game

weak competence assumption. Strong interpretation games emerge as the special case when the context contains the assumption that the speaker knows the complete answer to the question under discussion.

Definition 3 (Weak Interpretation Game) *Let s be an expression that is uttered in a context ct and that has a set of alternatives $ALT(s)$ such that each element thereof is consistent with ct . The weak interpretation game $\mathcal{G}_s = \langle \mathcal{T}, p^*, \mathcal{S}, \mathcal{A}, u_s, u_r \rangle$ for s is*

constructed as follows, where c_s and c_r are nominal cost functions:

$$\begin{aligned}
\mathcal{W} &= \{w \subseteq \mathcal{S} \mid w \neq \emptyset \wedge w \cap ct \text{ is consistent} \wedge \\
&\quad \forall s' \in \mathcal{S}. w \vdash s \Rightarrow s \in w\}, \\
\mathcal{T} &= POW(W) - \{\emptyset\}, \\
p^*(t) &= \frac{1}{|\mathcal{T}|}, \\
\mathcal{S} &= ALT(s), \\
t \models s' &\Leftrightarrow \forall w \in t. s' \in w, \\
\mathcal{A} &= \mathcal{T}, \\
u_s(t, s', a) &= -c_s(t, s') + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else,} \end{cases} \\
u_r(t, s', a) &= -c_r(s', a) + \begin{cases} 1 & \text{if } t = a, \\ 0 & \text{else,} \end{cases} \\
a_1 \subsetneq a_2 &\Rightarrow \forall s'. c_r(s', a_1) < c_r(s', a_2).
\end{aligned}$$

The difference between the strong and the weak competence assumption may be significant. This applies for instance to *free choice permission* sentences like (10a) (cf. Kamp 1973):

- (10) a. You may take an apple or a banana. (= $\diamond(A \vee B)$)
b. You may take an apple. (= $\diamond A$)
c. You may take a banana. (= $\diamond B$)
d. You may take an apple and a banana. (= $\diamond(A \wedge B)$)

Sentence (a) normally receives a *free choice* interpretation: You may take an apple, and you may take a banana. Furthermore the sentence implicates that the addressee is not allowed to take both an apple and a banana.

The free choice interpretation rests on the background assumption that the speaker is the one who grants the permission. This of course entails that the speaker is maximally competent about the permission state of the addressee. If the speaker simply reports the permission that somebody else granted (enforced for instance by the continuation *but I don't know which*), the free choice inference does not emerge. The disjunction still preferably receives an exhaustive interpretation — the addressee is not allowed to take both kinds of fruit simultaneously.

Let us construct an interpretation game for (10a). It seems reasonable to assume that the alternatives to (10a) are (10b–d), which are taken to have the logical forms which are given in brackets. If we assume the standard inference relation for modal logic (system K), we have four possible worlds here:

- $w_a = \{\diamond A, \diamond(A \vee B)\}$: You may take an apple but not a banana.

- $w_b = \{\diamond B, \diamond(A \vee B)\}$: You may take a banana but not an apple.
- $w_{a;b} = \{\diamond A, \diamond B, \diamond(A \vee B)\}$: You may take an apple and you may take a banana, but not both.
- $w_{ab} = \{\diamond A, \diamond B, \diamond(A \vee B), \diamond(A \wedge B)\}$: You may take both an apple and a banana.

(10a,d) are longer than (b,c), so it is reasonable to assume that they have higher costs. Let us say that $c(t, \diamond(A)) = c(t, \diamond(B)) = 0$ and $c(t, \diamond(A \vee B)) = c(t, \diamond(A \wedge B)) = \frac{1}{1000}$. Under the strong competence assumption, the IBR sequence comes out as in Table 13.

σ_0	$\diamond(A \vee B)$	$\diamond A$	$\diamond B$	$\diamond(A \wedge B)$	ρ_0	w_a	w_b	$w_{a;b}$	w_{ab}
w_a	$\frac{1}{2}$	$\frac{1}{2}$	0	0	$\diamond(A \vee B)$	$\frac{1}{2}$	$\frac{1}{2}$	0	0
w_b	$\frac{1}{2}$	0	$\frac{1}{2}$	0	$\diamond A$	1	0	0	0
$w_{a;b}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	0	$\diamond B$	0	1	0	0
w_{ab}	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\diamond(A \wedge B)$	0	0	0	1

σ_1	$\diamond(A \vee B)$	$\diamond A$	$\diamond B$	$\diamond(A \wedge B)$	ρ_1	w_a	w_b	$w_{a;b}$	w_{ab}
w_a	0	1	0	0	$\diamond(A \vee B)$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
w_b	0	0	1	0	$\diamond A$	1	0	0	0
$w_{a;b}$	0	$\frac{1}{2}$	$\frac{1}{2}$	0	$\diamond B$	0	1	0	0
w_{ab}	0	0	0	1	$\diamond(A \wedge B)$	0	0	0	1

σ_2	$\diamond(A \vee B)$	$\diamond A$	$\diamond B$	$\diamond(A \wedge B)$	ρ_2	w_a	w_b	$w_{a;b}$	w_{ab}
w_a	0	1	0	0	$\diamond(A \vee B)$	0	0	1	0
w_b	0	0	1	0	$\diamond A$	1	0	0	0
$w_{a;b}$	1	0	0	0	$\diamond B$	0	1	0	0
w_{ab}	0	0	0	1	$\diamond(A \wedge B)$	0	0	0	1

Table 13: IBR sequence

In the fixed point (σ_2, ρ_2) , (10a) correctly comes out as having the exhaustive free choice interpretation “You may take an apple and you may take a banana, but not both.”

Now suppose the strong competence assumption is replaced by its weak counterpart. This would be a scenario where the sender is not the person granting the permission, but just somebody who has information about the addressee’s permission state. Then we end up with fifteen information states. For reasons of space, I will only show the the five of them which are relevant (see Table 14). Here a fixed point is already reached at (σ_1, ρ_0) . In this fixed point, (10a) is interpreted as $\diamond A \leftrightarrow \neg \diamond B$, i.e. we get an ordinary exclusive

σ_0	$\diamond(A \vee B)$	$\diamond A$	$\diamond B$	$\diamond(A \wedge B)$		
$\{w_a\}$	$\frac{1}{2}$	$\frac{1}{2}$	0	0		
$\{w_b\}$	$\frac{1}{2}$	0	$\frac{1}{2}$	0		
$\{w_{a;b}\}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	0		
$\{w_{ab}\}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$		
$\{w_a, w_b\}$	1	0	0	0		
\vdots	\vdots	\vdots	\vdots	\vdots		
ρ_0	$\{w_a\}$	$\{w_b\}$	$\{w_{a;b}\}$	$\{w_{ab}\}$	$\{w_a, w_b\}$...
$\diamond(A \vee B)$	0	0	0	0	1	...
$\diamond A$	1	0	0	0	0	...
$\diamond B$	0	1	0	0	0	...
$\diamond(A \wedge B)$	0	0	0	1	0	...
σ_1	$\diamond(A \vee B)$	$\diamond A$	$\diamond B$	$\diamond(A \wedge B)$		
$\{w_a\}$	0	1	0	0		
$\{w_b\}$	0	0	1	0		
$\{w_{a;b}\}$	0	$\frac{1}{2}$	$\frac{1}{2}$	0		
$\{w_{ab}\}$	0	0	0	1		
$\{w_a, w_b\}$	1	0	0	0		
\vdots	\vdots	\vdots	\vdots	\vdots		

Table 14: IBR sequence

wide scope interpretation for the disjunction plus an ignorance implicature (the speaker does not know whether the addressee may take an apple or whether he may take a banana), rather than a free choice reading.

7 Embedded implicatures

Chierchia (2004) makes a case that scalar implicatures are actually not the outcome of a Gricean reasoning procedure that is based on rationality assumptions. Rather, he proposes that these inferences are strictly speaking not implicatures at all, but rather part of the literal meaning of the sentence in question. He motivates this, among other arguments, with examples like (11).

- (11) a. Kai had broccoli or some of the peas. ($B \vee \exists xPx$)
b. Kai had broccoli or all of the peas. ($B \vee \forall xPx$)

The argument runs as follows. Let A be an expression which competes with a logically stronger scalar alternative B . According to the Gricean view, if the speaker would believe that B , he would have said so. Since she didn't say it, she doesn't believe it. If she is also competent, we can conclude that she does not believe that B , and thus B is false. If this is applied to (11a), we have a scalar alternative (11b), which is logically stronger. Hence there should be an implicature to the effect that (11b) is false, i.e. that Kai didn't have broccoli and didn't eat all of the pears.

However, the (a)-sentence does not implicate that Kai didn't have broccoli. What is implicated is much weaker, namely that Kai had broccoli or he had some *but not all* of the pears. The strengthening of *some* to *some but not all* happens within the scope of the disjunction. Chierchia and his followers conclude from this that the information *he did not eat all of the peas* is part of the truth conditions of the second disjunct.

This is not the place to explore the issue whether or not implicatures are computed locally.¹¹ I will confine myself to showing why Chierchia's argument does not apply to the kind of pragmatic reasoning that is formalized by the IBR model.

As pointed out by Sauerland (2004), the alternatives that have to be taken into consideration when evaluating (11a) also include the sentences in (12).

- (12) a. Kai had broccoli. (= B)
 b. Kai had some of the peas. (= $\exists xPx$)
 c. Kai had all of the peas. (= $\forall xPx$)
 d. Kai had broccoli and some of the peas. (= $B \wedge \exists xPx$)
 e. Kai had broccoli and all of the peas. (= $B \wedge \forall xPx$)

These seven alternatives give rise to five different possible worlds:

- $w_{B \rightarrow \exists} = \{B, B \vee \exists xPx, B \vee \forall xPx\}$,
- $w_{\neg B \exists \rightarrow \forall} = \{\exists xPx, B \vee \exists xPx\}$,
- $w_{\neg B \forall} = \{\exists xPx, \forall xPx, B \vee \exists xPx, B \vee \forall xPx\}$,
- $w_{B \exists \rightarrow \forall} = \{B, \exists xPx, B \vee \exists xPx, B \vee \forall xPx, B \wedge \exists xPx\}$,
- $w_{B \forall} = \{B, \exists xPx, B \vee \exists xPx, B \vee \forall xPx, B \wedge \exists xPx, B \wedge \forall xPx\}$.

As the strong competence assumption is not warranted in this example, we have to construct a weak interpretation game, which has 31 different types. For reasons of space I will only present the IBR reasoning for the relevant subset thereof, which are seven types in this case. These are the five maximally informative types, plus two types consisting of two possible worlds each.¹² The first two steps of the IBR sequence are given in Table 15. The

¹¹An extensive discussion from a neo-Gricean point of view can be found in Geurts (2010).

¹²These seven types are the only relevant ones in the sense that all other types have probability 0 for each signal under each ρ_n in the IBR sequence.

σ_0	B	$\exists xPx$	$\forall xPx$	$B \vee \exists xPx$	$B \wedge \exists xPx$	$B \vee \forall xPx$	$B \wedge \forall xPx$
$\{w_{B-\exists}\}$	$\frac{1}{3}$	0	0	$\frac{1}{3}$	0	$\frac{1}{3}$	0
$\{w_{-B\exists-\forall}\}$	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0	0	0
$\{w_{-B\forall}\}$	0	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	0	$\frac{1}{4}$	0
$\{w_{B\exists-\forall}\}$	$\frac{1}{5}$	$\frac{1}{5}$	0	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	0
$\{w_{B\forall}\}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$
$\{w_{B-\exists}, w_{-B\exists-\forall}\}$	0	0	0	1	0	0	0
$\{w_{B-\exists}, w_{-B\forall}\}$	0	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$	0

ρ_0	$\{w_{B-\exists}\}$	$\{w_{-B\exists-\forall}\}$	$\{w_{-B\forall}\}$	$\{w_{B\exists-\forall}\}$	$\{w_{B\forall}\}$	$\{w_{B-\exists}, w_{-B\exists-\forall}\}$	$\{w_{B-\exists}, w_{-B\forall}\}$
B	1	0	0	0	0	0	0
$\exists xPx$	0	1	0	0	0	0	0
$\forall xPx$	0	0	1	0	0	0	0
$B \vee \exists xPx$	0	0	0	0	0	1	0
$B \wedge \exists xPx$	0	0	0	1	0	0	0
$B \vee \forall xPx$	0	0	0	0	0	0	1
$B \wedge \forall xPx$	0	0	0	0	1	0	0

Table 15: IBR sequence

relevant part of σ_1 is simply the transpose of the displayed part of ρ_0 , and (σ_1, ρ_0) form a fixed point.

As can be seen from the fourth row of ρ_0 , (11a) is interpreted as $\{w_{B-\exists}, w_{-B\exists-\forall}\}$. This means that either Kai ate broccoli or some but not all of the peas, but not both, and the speaker does not know which. The IBR model thus predicts a narrow scope exhaustive interpretation of *some*, an exclusive interpretation of *or* and an ignorance implicature, which is empirically correct.

It is illuminating to explore why Chierchia's problem does not apply here. It is true that the sender could have uttered $B \vee \forall xPx$, but she did not do so. There may be more than one reason though why she avoided that signal. In $w_{B-\exists}$ she believes it to be true but prefers to utter B, because it is more informative. In $w_{-B\forall}$ she prefers $\forall xPx$, for the same reason. In $w_{B\exists-\forall}$ she prefers $B \wedge \exists xPx$, again for the same reason. In $w_{B\forall}$, finally, she prefers the more informative $B \wedge \forall xPx$. So if the speaker is maximally competent, she would never utter (11b). The most plausible explanation for her using that signal would be that she is in state $\{w_{B-\exists}, w_{-B\forall}\}$, i.e. she believes that Kai either had broccoli but no peas or all peas but no broccoli, but not both. The fact that she refrained from using this signal thus only entails that she doesn't hold this belief state, which is consistent with the

intuitive interpretation of (11a).

8 Beyond nominal costs

In the model presented so far, the main objective of the players is to achieve correct information transmission. If the receiver guesses the sender's type correctly, both score a utility close to 1; otherwise their utility is close to 0. There are various domains where the degree of communicative success is not binary. If the sender wants to communicate that the temperature is 21 degree and the receiver guesses that the temperature is 20 degree, information transmission was not perfect, but much better than in a scenario where the receiver guesses 10 degree. Quite generally, there may be a graded notion of *similarity* between types, and the goal of communication is to maximize the similarity between the sender's type and receiver's action.¹³ If the degree of similarity between types is continuous, there may be scenarios where the expected gain in communicative success for choosing a more costly signal is lower than the increase in costs. This means that costs are not always nominal anymore. The consequences of this approach are explored in the final example of this article.

Krifka (2002) observes that the pragmatic interpretation of number words follows an interesting pattern:

“RN/RI principle:

- a. Short, simple numbers suggest low precision levels.
- b. Long, complex numbers suggest high precision levels.”

(Krifka 2002:433)

This can be illustrated with the following contrast:

- (13) a. The distance is one hundred meters. (= 100)
b. The distance is one hundred and one meter. (= 101)

The sentence (13b) suggests a rather precise interpretation (with a slack of at most 50 cm), while (13a) can be more vague. It may perhaps mean something between 90 and 110 meter. Actually, (13a) is pragmatically ambiguous; depending on context, it can be rather precise or rather vague. The crucial observation here is: A shorter number term like “one hundred” allows for a larger degree of vagueness than a more complex term like “one hundred and one.”

Krifka also observes that the degree of vagueness of a short term can be reduced by making it more complex—for instance by modifying it with “exactly”:

- (14) The distance is exactly one hundred meter. (= EX-100)

¹³This idea has been explored in detail in Jäger and van Rooij (2007); Jäger (2007a); Jäger et al. (2009).

sim	w_1	w_2	w_3	w_4
w_1	1	0.5	1	0.5
w_2	0.5	1	0.5	1
w_3	1	0.9	1	0.9
w_4	0.9	1	0.9	1

Table 16: Similarity function

Krifka (2002) accounts for these facts in terms of bidirectional OT, assuming a general preference for vague over precise interpretation. Krifka (2007) contains a revised analysis which employs game theoretic pragmatics. Space does not permit a detailed discussion of Krifka’s proposals; in the following I will just briefly sketch how the IBR model accounts for Krifka’s observations.

We use (13a,b) and (14) as expression alternatives. If we assume that the literal denotation of 100 and EXACTLY-100 is identical—100 meter sharp—, we end up with two possible worlds. However, the scenario potentially contains another source of uncertainty of the receiver about the sender’s information state. The sender may or may not consider precision important for the task at hand.¹⁴ This gives us four possible world:

- w_1 : length is 100 meter; precision is important.
- w_2 : length is 101 meter; precision is important.
- w_3 : length is 100 meter; precision is not important.
- w_4 : length is 101 meter; precision is not important.

If $[[\cdot]]$ is the function that maps a signal to the set of types where it is true, we assume

$$\begin{aligned} [[100]] &= [[\text{EXACTLY-100}]] = \{w_1, w_3\} \\ [[101]] &= \{w_2, w_4\} \end{aligned}$$

The utility functions now depend on the similarity between types. Formally, they are defined as:

$$\begin{aligned} u_s(t, s, a) &= \text{sim}(t, a) - c_s(t, s) \\ u_r(t, s, a) &= \text{sim}(t, a) - c_r(s, a) \end{aligned}$$

The similarity function sim is given in Table 16.

Furthermore, I assume that $c_s(t, 100) = 0$ and $c_s(t, 101) = c_s(t, \text{EXACTLY-100}) = 0.15$. The IBR sequence then comes out as shown in Table 17.

¹⁴Like in the examples relating to I- and M-implicatures, this game cannot be constructed according to the recipe given above. For the time being, I simply stipulate the parameters of the game and leave the issue how the game is to be extracted from the linguistic information for further research.

σ_0	100	101	EXACTLY-100
w_1	$\frac{1}{2}$	0	$\frac{1}{2}$
w_2	0	1	0
w_3	$\frac{1}{2}$	0	$\frac{1}{2}$
w_4	0	1	0

ρ_0	w_1	w_2	w_3	w_4
100	$\frac{1}{2}$	0	$\frac{1}{2}$	0
101	0	$\frac{1}{2}$	0	$\frac{1}{2}$
EXACTLY-100	$\frac{1}{2}$	0	$\frac{1}{2}$	0

σ_1	100	101	EXACTLY-100
w_1	1	0	0
w_2	0	1	0
w_3	1	0	0
w_4	1	0	0

ρ_1	w_1	w_2	w_3	w_4
100	$\frac{1}{3}$	0	$\frac{1}{3}$	$\frac{1}{3}$
101	0	1	0	0
EXACTLY-100	$\frac{1}{2}$	0	$\frac{1}{2}$	0

σ_2	100	101	EXACTLY-100
w_1	0	0	1
w_2	0	1	0
w_3	1	0	0
w_4	1	0	0

ρ_2	w_1	w_2	w_3	w_4
100	0	0	$\frac{1}{2}$	$\frac{1}{2}$
101	0	1	0	0
EXACTLY-100	1	0	0	0

Table 17: IBR sequence

In σ_0 and ρ_0 all signals are used/interpreted according to their literal meaning. In σ_1 the sender of type w_4 decides that it is better to use the imprecise but more economical expression 100 rather than 101. Note that in type w_2 , σ_1 still uses 101 because for this type precision is sufficiently important to justify the higher expression costs.

As a consequence, 100 receives a vague interpretation in ρ_1 . To avoid this vagueness, type w_1 in σ_2 prefers the more costly but also more precise expression EXACTLY-100. As a consequence, the expression 100 does not carry any information anymore about the distinction between the two possible lengths in ρ_2 . It still does convey information though, namely the pragmatic information that the sender is in a state where precision is not important. The two complex expressions only receive a precise information in ρ_2 . (σ_2, ρ_2) are a fixed point.

9 Conclusion

Game theoretic methods have received a good deal of attention in the linguistics community in recent years. The aim of this chapter is to illustrate the potential of this approach by showing how the neo-Gricean program of pragmatics can be spelled out in this framework. It is thus representative of one major line of research, which uses rationalistic game theory to model the decision making of language users in specific communicative situations. This is mostly applied to problems of pragmatic interpretation (like the work by Benz, Parikh, van Rooij and others that has already been mentioned above), but there is also some interest in NLP circles to apply these concepts to language generation (see for instance Golland et al. 2010; Klabunde 2009).

Even though game theory has originally been conceived as a tool to make prescriptive claims about the behavior of fully rational agents, since the seminal work of Maynard Smith (1982) it has also become a standard tool in biomathematics. There it is used to model Darwinian evolution in situations where replicative success depends on interactions between different organisms. In this context, evolutionary game theory has been used extensively to study the evolution of animal communication systems, including human language (see for instance Maynard Smith 1991; Nowak and Krakauer 1999).

Since the 1990s it has been recognized by economists that the evolutionary logic can be fruitfully applied to study cultural phenomena in situations where the behavior of humans is governed by imitation and learning. This perspective has been applied to natural language as well to explain the emergence and stability of linguistic conventions in various domains of grammar (see for instance van Rooij 2004; Huttegger 2007; Jäger 2007b).

Acknowledgments

Thanks to Michael Franke and Roger Levy for pointing out several mistakes in a draft version of this article, and to Jason Quinley for correcting my English.

References

- Battigalli, Pierpaolo and Marciano Siniscalchi. 2002. Strong belief and forward induction reasoning. *Journal of Economic Theory* 106, 356–391.
- Benz, Anton, Gerhard Jäger, and Robert van Rooij, eds. 2005. *Game Theory and Pragmatics*. Basingstoke: Palgrave MacMillan.
- Benz, Anton and Robert van Rooij. 2007. Optimal assertions and what they implicate. *Topoi — an International Review of Philosophy* 27, 63–78.
- Blutner, Reinhard. 2001. Some aspects of optimality in natural language interpretation. *Journal of Semantics* 17, 189–216.
- Büring, Daniel. 1995. The 59th street bridge accent. Ph.D. thesis, University of Tübingen.
- Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.
- Camerer, Colin F. and Ernst Fehr. 2006. When does “economic man” dominate social behavior? *Science* 311, 47.
- Chierchia, Gennaro. 2004. Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. *Structures and Beyond*, edited by Adriana Beletti, 39–103, Oxford: Oxford University Press.
- Clark, H. and C. Marshall. 1981. Definite reference and mutual knowledge. *Elements of Discourse Understanding*, edited by A. Joshi, B. Webber, and Ivan Sag, 10–63, Cambridge, UK: Cambridge University Press.
- Crawford, Vincent P. and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50, 1431–1451.
- Dalrymple, Mary, Makoto Kanazawa, Sam Mchombo, and Stanley Peters. 1994. What do reciprocals mean. *Proceedings from Semantics and Linguistic Theory IV*, edited by Mandy Harvey and Lynn Santelmann, 61–78, Cornell University.
- Dekker, Paul and Robert van Rooij. 2000. Bi-directional optimality theory: An application of game theory. *Journal of Semantics* 17, 217–242.
- Fox, Danny. 2007. Free choice and the theory of scalar implicature. *Presupposition and Implicature in Compositional Semantics*, edited by Uli Sauerland and Penka Stateva, 71–120, Hampshire: Palgrave MacMillan.
- Franke, Michael. 2008. Interpretation of optimal signals. *New Perspectives on Games and Interaction*, edited by Krzysztof Apt and Robert van Rooij, 297–310, Amsterdam: Amsterdam University Press.
- Franke, Michael. 2009. Signal to act: Game theory in pragmatics. Ph.D. thesis, University of Amsterdam.
- Franke, Michael and Gerhard Jäger. 2010. Bidirectional optimization from reasoning and learning in games, manuscript, University of Tübingen.
- Gerbrandy, Jelle. 1999. Bisimulations on planet Kripke. Ph.D. thesis, University of Amsterdam.
- Geurts, Bart. 2010. *Quantity Implicatures*. Cambridge University Press.
- Golland, Dave, Percy Liang, and Dan Klein. 2010. A game-theoretic approach to generating spatial descriptions. *Proceedings of EMNLP 2010 (Empirical Methods in Natural*

- Language Processing*), edited by Hang Li and Lluís Màrquez, 410–419, ACL.
- Grice, Herbert Paul. 1975. Logic and conversation. *Syntax and Semantics 3: Speech Acts*, edited by P. Cole and J. Morgan, 41–58, New York: Academic Press.
- Horn, Laurence. 1984. Towards a new taxonomy for pragmatic inference: Q-based and R-based implicatures. *Meaning, Form, and Use in Context*, edited by Deborah Schiffrin, 11–42, Washington: Georgetown University Press.
- Huttegger, Simon H. 2007. Evolution and the explanation of meaning. *Philosophy of Science* 74, 1–27.
- Jäger, Gerhard. 2007a. The evolution of convex categories. *Linguistics and Philosophy* 30, 551–564.
- Jäger, Gerhard. 2007b. Evolutionary Game Theory and typology: a case study. *Language* 83, 74–109.
- Jäger, Gerhard. 2007c. Game dynamics connects semantics and pragmatics. *Game Theory and Linguistic Meaning*, edited by Ahti-Veikko Pietarinen, 89–102, Elsevier.
- Jäger, Gerhard and Christian Ebert. 2009. Pragmatic rationalizability. *Proceedings of Sinn und Bedeutung 13*, edited by Arndt Riester and Torgrim Solstad, number 5 in SinSpeC. Working Papers of the SFB 732, 1–15, University of Stuttgart.
- Jäger, Gerhard, Lars Koch-Metzger, and Frank Riedel. 2009. Voronoi languages, manuscript, University of Bielefeld/University of Tübingen.
- Jäger, Gerhard and Robert van Rooij. 2007. Language structure: Psychological and social constraints. *Synthese* 159, 99–130.
- Jaynes, Edwin Thompson. 2003. *Probability Theory: The Logic of Science*. Cambridge University Press.
- Kamp, Hans. 1973. Free choice permission. *Proceedings of the Aristotelian Society* 74, 57–74.
- Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30, 669–690.
- Klabunde, Ralf. 2009. Towards a game-theoretic approach to content determination. *Proceedings of the 12th European Workshop on Natural Language Generation (ENLG-09)*, 102–105.
- Krifka, Manfred. 2002. Be brief and vague! and how bidirectional optimality theory allows for verbosity and precision. *Sounds and Systems. Studies in Structure and Change. A Festschrift for Theo Vennemann*, edited by David Restle and Dietmar Zaefferer, 439–358, Berlin: Mouton de Gruyter.
- Krifka, Manfred. 2007. Approximate interpretation of number words: A case for strategic communication. *Cognitive foundations of interpretation*, edited by Gerlof Bouma, Irene Krämer, and Joost Zwarts, 111–126, Amsterdam: Koninklijke Nederlandse Akademie van Wetenschappen.
- Levinson, Stephen C. 2000. *Presumptive Meanings*. MIT Press.
- Lewis, David. 1969. *Convention*. Cambridge, Mass.: Harvard University Press.
- Maynard Smith, John. 1982. *Evolution and the Theory of Games*. Cambridge, UK: Cambridge University Press.
- Maynard Smith, John. 1991. Honest signalling: the Philip Sidney game. *Animal Behaviour*

42, 1034–1035.

- Nowak, Martin A. and David C. Krakauer. 1999. The evolution of language. *Proceedings of the National Academy of Sciences* 96, 8028–8033.
- Parikh, Prashant. 2001. *The Use of Language*. Stanford: CSLI Publications.
- Pietarinen, Ahti-Veikko, ed. 2007. *Game Theory and Linguistic Meaning*. Elsevier.
- Rabin, Matthew. 1990. Communication between rational agents. *Journal of Economic Theory* 51, 144–170.
- Rooth, Mats. 1985. Association with focus. Ph.D. thesis, University of Massachusetts, Amherst.
- Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27, 367–391.
- Skyrms, Brian. 2010. *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press.
- Slonim, Robert L. 2005. Competing against experienced and inexperienced players. *Experimental Economics* 8, 55–75.
- Spence, Michael. 1973. Job Market Signaling. *The Quarterly Journal of Economics* 87, 355–374.
- Stalnaker, Robert. 1978. Assertion. *Pragmatics*, edited by Peter Cole, number 9 in Semantics & Pragmatics, 315–32, New York: Academic Press.
- Stone, Matthew. 1998. Modality in dialogue: planning, pragmatics and computation. Ph.D. thesis, Rutgers University.
- van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9, 333–377.
- van Rooij, Robert. 2004. Signalling games select Horn strategies. *Linguistics and Philosophy* 27, 493–527.
- van Rooij, Robert and Katrin Schulz. 2004. Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information* 13, 491–519.

Appendix

Theorem 1 (Fixed Point of IBR sequence) *For each Strong or Weak Interpretation Game, there is a number $n \geq 0$ such that for all $m > n$: $\sigma_m = \sigma_n$ and $\rho_m = \rho_n$.*

Proof: All strategies that are considered in an IBR sequence assign a uniform distribution over a subset of \mathcal{S} to each type (sender) or a uniform distribution over a subset of \mathcal{A} to each signal (receiver). Since the games are finite, there are only finitely many such strategies.

Consider a strategy profile $\langle \sigma, \rho \rangle$. We define five measures over profiles (average sender utility, measure of entropy of σ , average receiver utility, negative of the number of false interpretations assigned to surprise signals, measure of entropy of ρ):

- $m_1(\sigma, \rho) = \sum_{t \in \mathcal{T}} p^*(t) \sum_s \sigma(s|t) \sum_{a \in \mathcal{A}} \rho(a|s) u_s(t, s, a)$,

- $m_2(\sigma, \rho) = \sum_{t \in \mathcal{T}} |\{s | \sigma(s|t) > 0\}|$,
- $m_3(\sigma, \rho) = \sum_{t \in \mathcal{T}} p^*(t) \sum_s \sigma(s|t) \sum_{a \in \mathcal{A}} \rho(a|s) u_r(t, s, a)$,
- $m_4(\sigma, \rho) = - \sum_{s: \forall t \sigma(s|t)=0} |\{a | \rho(a|s) > 0 \wedge a \neq s\}|$,
- $m_5(\sigma, \rho) = \sum_{s \in \mathcal{S}} |\{a | \rho(a|s) > 0\}|$.

Next we define a partial order over profiles via a lexicographic order via their $m_1 - m_5$ -values:

$$\langle \sigma_1, \rho_1 \rangle > \langle \sigma_2, \rho_2 \rangle \text{ iff } \exists i. m_i(\sigma_1, \rho_1) > m_i(\sigma_2, \rho_2) \wedge \forall j < i : m_j(\sigma_1, \rho_1) = m_j(\sigma_2, \rho_2).$$

It is obvious that $>$ must be acyclic.

Let $\langle \sigma_n, \rho_n \rangle$ be an element of an IBR sequence, and let $\sigma_{n+1} \neq \sigma_n$. As $\sigma_{n+1} = BR_s(\rho_n)$, for each t and each s with $\sigma_{n+1}(s|t) > 0$, $EU_s(s|t; \rho_n) = \max_{s'} EU_s(s'|t; \rho_n)$. Hence $m_1(\sigma_{n+1}, \rho_n) = \max_{\sigma} m_1(\sigma, \rho_n)$. The usage of the principle of insufficient reason in the definition of BR_s ensures that within the set of sender strategies having this property, σ_{n+1} is the unique strategy that maximizes $m_2(\cdot, \rho_n)$. Hence $\forall \sigma \neq \sigma_{n+1}. \langle \sigma_{n+1}, \rho_n \rangle > \langle \sigma, \rho_n \rangle$. In particular, $\langle \sigma_{n+1}, \rho_n \rangle > \langle \sigma_n, \rho_n \rangle$.

Now suppose $\rho_{n+1} \neq \rho_n$. Note that for any t, s, a_1, a_2 , if $u_r(t, s, a_1) \geq u_r(t, s, a_2)$, then $u_s(t, s, a_1) \geq u_s(t, s, a_2)$. To see why this is so, suppose $a_1 \neq a_2$. If $u_r(t, s, a_1) \geq u_r(t, s, a_2)$, then either $a_1 = t$ and $a_2 \neq t$, or $a_1, a_2 \neq t$ and $c_r(s, a_1) \leq c_r(s, a_2)$. In both cases, $u_s(t, s, a_1) \geq u_s(t, s, a_2)$.

Since $br_r(s, \sigma)$ gives the set of actions that maximize the expected value of u_r given s and σ , it only contains actions that also maximize the expected value of u_s given s and σ . Hence $BR_r(\sigma)$ is a strategy that maximizes $m_1(\sigma, \cdot)$. m_2 only depends on σ . Hence $BR_r(\sigma)$ maximizes $m_1(\sigma, \cdot)$ and $m_2(\sigma | \cdot)$. By definition, it also maximizes $m_3(\sigma, \cdot)$. By the belief revision policy that is implemented via the best response to surprise signals in the definition of BR_r , $BR_r(\sigma)$ also maximizes $m_4(\sigma, \cdot)$. Finally, the definition of BR_r ensures that $BR_r(\sigma)$ is the unique strategy among all strategies that maximize m_1, \dots, m_4 that maximizes m_5 . Therefore $\langle \sigma_{n+1}, \rho_{n+1} \rangle > \langle \sigma_{n+1}, \rho_n \rangle$.

The relation $>$ is acyclic. All elements of the IBR sequence are strategies that map types (signals) to uniform distributions over subsets of \mathcal{S} (\mathcal{A}). There are only finitely many such strategies. Therefore the IBR sequence must have a fixed point. \dashv