Game dynamics connects semantics and pragmatics

Gerhard Jäger
University of Bielefeld

Abstract

The chapter first gives an overview over the evolutionary interpretation of game theory, and it compares two versions of it, the replicator dynamics and the best response dynamics. The ensuing notions of evolutionary stability are explored. In the second part, it is argued that the best response dynamics lends itself to an epistemic interpretation, and that this provides a suitable game theoretic foundation for pragmatic reasoning in the Gricean tradition.

1 Introduction

Game theory has originally been conceived as a theory of strategic interaction among fully rational agents. Its applicability to real life phenomena like economic or political processes therefore depends on how realistic it is to assume that the acting individuals in these processes are fully rational. Rationality here means, among other things, full awareness of one's own beliefs and preferences and logical omniscience. Even stronger, for classical game theory to be applicable, every agent has to ascribe full rationality to each other agent.

These assumptions are of course unrealistic when applied to humans. This does not devaluate game theoretic models though. An apologist of the classical model might argue that to ride a bike one has to be able to act in accordance with the laws of physics, but one does not need to be able to solve differential equations. Likewise, to successfully embark upon a strategic interaction one does not need to be able to solve game theoretic problems; all that is required is to **act** in accordance with the laws of game theory.[1]

This argument has a certain appeal. If game theory is to be applied as a prescriptive, rather than a normative theory, the question remains open though how not-fully-rational beings achieve the quasi-rationality that is required to apply game theory in the first place.

There are various answers to this problem. One line of research, going back to the work of Herbert Simon [12] (see also Rubinstein's [9]) explores the consequences of giving up the strong rationality assumptions of classical game theory. In other words, agents are assumed to be **boundedly rational**. The **evolutionary** interpretation of game theory

[1] I owe this comparison to Helge Ritter (p.c.).

(see for instance Maynard Smith [7]) completely gives up any rationality assumptions. Instead, game theory is used to describe the dynamics of entire populations of agents. Strategies (in the game theoretic sense) are identified with heritable traits of individuals, and utility with replicative success. Since replicative success of an individual may depend on other individuals of the same population, this involves a strategic component. Game theory can thus be used to model a evolution via natural selection in the Darwinian sense.

It has repeatedly been noticed that Gricean pragmatics has a strong game theoretic flavor (see for instance Stalnaker's discussion in [14]). In particular, it makes the same strong rationality assumptions as classical game theory, and the mentioned objections apply as well. One would thus expect that the notion of bounded rationality has a role to play in laying the foundations of natural language pragmatics. The connection between evolutionary game theory and pragmatics is perhaps not so obvious, but research in economy has shown that evolutionary population dynamics is a useful tool to model cultural processes as well (see for instance [18]). Language use, as a cultural phenomenon, thus falls squarely into the realm of this interpretation of game theory as well.

In the present paper I will explore a particular version of an evolutionary game dynamics called **best response dynamics**. While its standard interpretation applies to learning processes in population, it can also receive a very natural epistemic interpretation involving boundedly rational agents. This model will be applied to the problem of relating (conventionalized) semantic and the (not conventionalized) pragmatic aspects of natural language interpretation.

2 Evolutionary game theory

2.1 The replicator dynamics

Evolutionary game theory (EGT) was developed by theoretical biologists, especially John Maynard Smith (cf. [7]) as a formalization of the neo-Darwinian concept of evolution via natural selection. It builds on the insight that many interactions between living beings can be considered to be games in the sense of game theory — every participant has something to win or to lose in the interaction, and the payoff of each participant can depend on the actions of all other participants. In the context of evolutionary biology, the payoff is an increase in fitness, where fitness is basically the expected number of offspring. According to the neo-Darwinian view on evolution, the units of natural selection are not primarily organisms but heritable traits of organisms. If the behavior of organisms, i.e., interactors, in a game-like situation is genetically determined, the strategies can be identified with gene configurations.

For illustration, consider a simple coordination game (from [5]). The utility matrix is given in Table 1. In the evolutionary setting, this is to be interpreted as follows. There is a large population. Each member of the population belongs to one of two sub-groups, A or B. Group membership is heritable. The individuals in the population reproduce via cloning (i.e. each newborn has exactly one parent). Reproductive success depends on the

	A	B
A	1	0
B	0	1

Tab. 1: Utility matrix for a simple coordination game

interaction with the other members of the population. The average number of offspring of an individual of type A equals the expected utility of a player of strategy A when playing against a random member of the population, and likewise for group B. For the given utility matrix, this means that the average number of offspring of an A-individual equals the proportion of A-players in the population, and the same for B.

If more than half of the population is of type A, A-players will thus on average have more children than B-players, and the proportion of A-players increases over time. The population as a whole will converge towards a state with only A-players. If the B-players have a majority in the initial state, the population converges towards a homogeneous B-state.[2] If the initial state is exactly 50:50, the population will remain in this state because A-players and B-players have exactly the same birth rate.

There are thus three stationary states of the population: 100% A-players, 100% B-players, and precisely fifty-fifty. Note that these are exactly the three Nash equilibria of this game. There is a difference though between the mixed equilibrium on the one hand, and the two pure equilibria on the other hand. Let the population be in the 50:50 state, but let us assume that the population dynamics is slightly noisy. This may be due to sampling effects if the population is finite after all, or replication may be unfaithful with a certain small probability (like mutations in genetic transmission). Then the population may leave the Nash equilibrium and develop a small A-bias or B-bias. However, as soon as the population has a bias, natural selection will enhance that bias until the population converges towards one of the two pure states. If, on the other hand, the population is in one of the two homogeneous states, a small group of invaders from the other strategy will die out soon because the receive a much lower utility against the incumbent population than the incumbents against themselves.

So while Nash equilibria correspond to evolutionarily stationary states, some such states are resistant against mutations, while other states aren't. Maynard Smith dubbed the resistant equilibria **evolutionarily stable states** (ESS).

It turns out that the notion of evolutionary stability is closely related to the rationalistic notion of a Nash equilibrium, but there are subtle differences. It can be shown that the following proper inclusions hold:

$$\textit{Strict Nash Equilibria} \subset \textit{Evolutionarily Stable Strategies} \subset \textit{Nash Equilibria}$$

The mixed equilibrium from the example above demonstrates that there are Nash equilibria that are not an ESS. As an example for an ESS that is not a strict Nash equilibrium,

[2] Standard EGT assumes that populations are so large that they can be considered infinite for all practical purposes. Random variation is disregarded.

	R	P	S
R	1	0	2
P	2	1	0
S	0	2	1

Tab. 2: Utility matrix for Rock-Paper-Scissor

consider the well-known game Rock-Paper-Scissor. The utility matrix is given in Table 2. This game has exactly one Nash equilibrium, the one where each of the three strategies is played with a probability of $\frac{1}{3}$. Suppose a population is in this equilibrium, i.e. each of three sub-populations have exactly the same size. Than each sub-population has the same birth rate and the proportions are stationary. Now suppose Rock gets a small edge over the other two strategies due to unfaithful replication. Then in the next generation, Paper will thrive, one generation later Scissors, than again Rock etc. ad infinitum. Without another unfaithful replication, the population will not return into equilibrium. This illustrates that the single Nash equilibrium of this game is not an ESS.

The population dynamics that ensues if the expected utility is identified with the expected number of offspring is called the **replicator dynamics**. In [7], Maynard Smith gives the necessary and sufficient conditions for a strategy (of a symmetric game) to be evolutionarily stable according to this dynamics:

- s is an Evolutionarily Stable Strategy in the replicator dynamics iff
  1. $u(s, s) \geq u(t, s)$ for all t, and
  2. if $u(s, s) = u(t, s)$ for some $t \neq s$, then $u(s, t) > u(t, t)$.

2.2 The best response dynamics

Many social scientists assume that cultural variables undergo an evolutionary process in a way more or less similar to genes in biology. How close this similarity actually is is a matter of dispute though (see for instance the discussion in [8]). If cultural evolution exists, evolutionary game theory should be applicable in this domain as well. There is in fact a considerable body of literature on the subject from economics and other social sciences (see [13, 17, 18] and the literature cited therein).

A strategy, in the social setting, can be considered a behavioral disposition, very much like in the original, rationalistic interpretation of game theory. To apply an evolutionary model to social phenomena though, it has to be clarified how strategies reproduce. To pose the question in more general terms, which micro-dynamics underlies the macro-dynamics that is modeled by the evolutionary model? Obvious candidates are learning and imitation. If various strategies are differentially successful in being learned and imitated, we expect a process which resembles natural selection.

Such a learning or imitation dynamics disregards the rationality and creativity of human agents. The **best response dynamics** (introduced in [6] and thoroughly investigated in

[3]) takes these aspects into account, but without adopting the extreme assumptions of the classical model.

Suppose there is a population of individuals that play certain strategies of a game, as in the previous setting. In every time step, a new member enters the population. Unlike in the biological setting, the newcomer may freely choose her strategy. If we suppose that newcomers are rational (and well-informed) enough to maximize their expected utility, they will choose a (possibly mixed) strategy that is a best response to the average strategy of the population. Repeating this addition of new members indefinitely, an evolutionary dynamics ensues, but a non-Darwinian one. New strategies may be invented with the purpose of maximizing utility, while in Darwinian evolution, new strategies only emerge due to undirected random mutation.

Despite this massive conceptual difference, the replicator dynamics and the best response dynamics are mathematically similar enough to subsume both under the heading "Evolutionary Game Theory".

The notion of evolutionary stability can be applied to the best response dynamics as well. It is easy to see that all ESSs are Nash equilibria — recall that by definition, a Nash equilibrium is a strategy that is a best response to itself. But what are the sufficient conditions for stability here? Reconsider the Rock-Paper-Scissor game. Suppose the population is close to equilibrium — there is the same number of Paper players and Scissor players, and a tiny excess of Rock players. Then the next newcomer will play Paper, and this will continue until Scissor becomes the best response to the population average, which will be followed by Rock etc. This seems similar to the replicator dynamics scenario. However, here the difference between the state of the population (seen as a vector of fractions) and the Nash equilibrium actually converges to zero. In other words, the Nash equilibrium is actually evolutionarily stable here.

The example illustrates that the best response dynamics induces a notion of ESS that is slightly more inclusive than Maynard Smith's notion. It can be defined in the following way (which is essentially identical to Hofbauer and Sigmund's formulation, see [4], p. 96.

- s is an Evolutionarily Stable Strategy in the best response dynamics iff
  1. $u(s,s) \geq u(t,s)$ for all t, and
  2. if $u(s,s) = u(t,s)$ for some $t \neq s$, then there is a t' with $u(t',s) = u(s,s)$ and $u(t',t) > u(t,t)$

2.3 Epistemic interpretation of the best response dynamic

The best response dynamic can also be given an epistemic interpretation. Let us return to the classical picture of a strategic two-person game, where each player has a preference ordering over the set of profiles (including the mixed ones), that can be represented by a utility function in the standard way. Suppose, however, that the players are entirely irrational. They choose their strategy according to prejudice rather than rational deliberation.

For the sake of simplicity, let us assume that both players have the same prejudices, and this prejudice is common knowledge.

It might occur though that player a is not entirely irrational but makes a rational choice with some probability $\varepsilon > 0$ that may be arbitrarily small. Making a rational choice means to play a best response against the prejudicial strategy of the other player. Depending on the nature of the original strategy, a's choice may now differ from the original strategy by a small amount.

Now suppose that the other player, b, is also rational with probability ε, and furthermore he assumes that a acts as described in the previous paragraph. With probability ε, will thus play the best response to a's strategy, which in turn is the initial strategy with probability $1-\varepsilon$ and a best response to it with probability ε.

This process may be iterated. In this way, we may define an infinite sequence of strategy profiles, starting with the initial prejudice and increasing the strategic depth at every step. If I take the action of my opponent into account in my decision, I have the strategic depth of at least 1. If I also take into account that my partner may take my actions into account, my strategic depth is 2 etc. In general, if I assign strategic depth n to my partner, my own strategic depth is $n+1$. Boundedly rational agents have an upper limit for their strategic depth though. Intuitively, a strategic depth of n is the ability to "think n steps ahead" in a sequential game, or to "think around n corners".

To make this notion formally precise, let a symmetric two-person A be given. x_i, the initial strategy, is a (possibly mixed) strategy for A, and ε is a real number with $0 < \varepsilon \leq 1$. Let $\beta(x)$ be the set of best responses to the strategy x according to A. A **deliberation sequence (based on ε)** $x_0, x_1, x_2, \ldots$ is an infinite sequence of strategies with the property that:

$$x_0 = x_i \quad (1)$$

$$x_{n+1} \in \{\varepsilon y + (1-\varepsilon)x_n | y \in \beta(x_n)\} \quad (2)$$

Let us assume that the agents are conservative, i.e. ε is very small. A deliberation sequence may nevertheless lead to another strategy than x_i. In general we say that **cautious deliberation leads from x_i to x_f** iff there is a positive $\varepsilon_0 \leq 1$ such that for all $\varepsilon < \varepsilon_0$ and for all deliberation sequences $\vec{x}$ based on ε with $x_0 = x_i$ it holds that $\vec{x}$ converges to x_f. Intuitively this means that boundedly rational players with a sufficiently small probability to act rationally and having the prejudice to play x_i will play, with arbitrary approximation, x_f provided they have, and assign to each other, a sufficiently large strategic depth.

Mathematically, cautious deliberation sequences are identical to time series in a discrete best response dynamics. One would thus expect that cautious deliberation may lead to some strategy x_f if and only if x_f is evolutionarily stable according to the best response dynamics. This equivalence (between end points of cautious deliberation and ESSs) does hold if we add the assumption that with an arbitrarily small probability η, players choose their strategy completely at random, without any considerations of prejudice or rationality.

Let us reconsider the two games discussed so far. Suppose Rock-Paper-Scissor players have the prejudice to play the Nash equilibrium $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. A boundedly rational player may decide to play Rock instead with probability ε because this is as good a response to the Nash equilibrium than the equilibrium itself. However, a second round of deliberation will reveal that with a strategic depth of 2, Paper turns out to be the best response. Since ε can become arbitrarily small, the probability of playing Rock rather than the mixed equilibrium is thus arbitrarily small, and the same holds for all other strategies that differ from the equilibrium. Cautious deliberation will thus always remain in the neighborhood of the equilibrium, and this neighborhood may be arbitrarily small.[3]

Compare this to the $(\frac{1}{2}, \frac{1}{2})$ equilibrium of the coordination game in Table 1. Suppose a player decides to play, with probability ε, the pure strategy A instead of the mixed equilibrium. A second round of deliberation reveals that the best response to this mixed strategy is the pure A, etc. ad infinitum. Via iterated deliberation, the probability of A will thus converge to 1. The Nash equilibrium $(\frac{1}{2}, \frac{1}{2})$ is thus not an endpoint of cautious deliberation, while the two pure equilibria A and B are.

2.4 Trajectories

The best response dynamics does not only define a notion of stability, but also sequences of strategies that may start at any point in the strategy space and lead, in most cases, to ESSs. This offers a partial solution to one of the central problems of game theory as such, namely equilibrium selection. If two players can not communicate before the game (either due to the lack of communication channels or the lack of mutual trust) and the game has multiple equilibria, there is no obvious way to predict the action of the other player and thus to make an informed strategy choice oneself, even if it is common knowledge that all players are perfectly rational. The situation may actually improve if the players are boundedly rational in the sense described above, provided they know each other's prejudices, i.e. the strategy that the other will play if he does not apply rational and strategic deliberation.

Schelling's ([10]) observation about the role of **salience** in equilibrium selection is a case in point. Schelling assumes that in a symmetric game with several equilibria, it is advisable to choose one that is more salient than the other(s). The point can nicely be illustrated by an experiment that is reported in Camerer ([1]). The test persons were grouped in couples. Each person was asked to secretly write down a day of the year. If both members of a couple managed to write down the same date (without any previous communication), they both scored a point. It turned out that a majority wrote down salient days like December 25.

This can be seen as a coordination game with 366 different strategies. The utility is 1 at the diagonal and 0 at all other profiles. The expected probability of a test person that is not thinking strategically will be distributed across all 366 dates, with probability peaks at salient days like Christmas eve. The tendency to choose December 25 will be enhanced by each round of strategic thinking, since the best response to a strategy with a probability peak at December 25 is to choose that very date with probability 1. The best

[3] It can actually be shown that cautious deliberation will always lead to the equilibrium in this game, independently of the initial state, but the proof of this fact is left out here for reasons of space. The interested reader is referred to [4].

	S	H	F
S	3	0	-5
H	1	1	-6
F	-2	-3	-10

Tab. 3: The extended stag hunt game

response dynamics is thus bound to converge to this pure strategy, and this is how most test persons behave.

Lewis ([5]) points out that precedent is a good heuristic for equilibrium selection as well. If you played such a coordination game against the same partner before and you managed to meet at an equilibrium, it is prudent to stick to that equilibrium. This is unsurprising if we assume that people are more likely to repeat themselves than to change their strategy without reason.

In the previous examples, each strategy is also a Nash equilibrium if played against itself. Table 3 contains a more complicated example that includes a strictly dominated strategy. If the game is restricted to the strategies S and H, we obtain the well-known stag hunt game. As a reminder: two hunters have the choice to try to catch a stag (S) or a hare (H). A stag is preferable over a hare, but stag hunt requires the two to collaborate, while each hunter can hunt a hare by himself. The worst outcome is to rely on the cooperation of the other hunter and to try to hunt a stag while the other one is in fact not collaborating. The game has two Nash equilibria — both hunting stag or both hunting hare.

In the extended stag hunt game, each player has a third option, namely to coerce the other hunter to go stag hunting by force (F). Such a conflict reduces the utility for both participants, but the one applying force is better off than the one being coerced. Retaliating with force is the worst outcome for both because they will embark upon a costly fight. The best reaction to force is to comply and play S.

The strategy F is strictly dominated and thus should play no role in the considerations of rational players. However, suppose applying force is the first choice that an irrational player would choose if he does not think about the consequences of this action. Then the best response is to play S. S is also the best response to any convex combination of the strategies S and F (i.e. any mixed strategy which assigns some probability p to S and $1-p$ to F). Cautious deliberation thus necessarily leads from F to S.

In real life terms, this illustrates an effective threat to break a deadlock. If one player conveys the impression that he might be willing to apply force — even though this is irrational — then this possibility is reason enough for a sufficiently rational player to comply, as long as compliance is in his own enlightened self interest. (Of course in reality threats also work if the victim is forced to act against his own good interests, but this only works if the threat either does no harm to the bully or else the bully is not arbitrarily rational.)

2.5 Best response dynamics in asymmetric games

So far I have restricted the discussion to symmetric games, i.e. games where both players have the same strategy set and the same utility matrix. A symmetric Nash equilibrium is a strategy in a symmetric game that is a best response to itself. Asymmetric games, on the other hand, are two-person games where the two players play different roles (or, in the population dynamics interpretation, belong to different populations). An asymmetric Nash equilibrium is a pair of strategies $\langle s_i^A, s_j^B \rangle$ (of player A and B respectively), such that s_i^A is the best response to s_j^B and vice versa.

Asymmetric games can be transformed into symmetric ones in a straightforward way. Suppose the players are A and B, their strategy sets are $S^A = s_1^A, \ldots, s_n^A$ and $S^B = s_1^B, \ldots, s_m^B$, and their utility matrices are u_A and u_B respectively. The new symmetric game has the strategy set $S^A \times S^B$. In other words, each strategy of the meta-game is a pair consisting of an A-strategy and a B-strategy. The utility is calculated as

$$u(\langle s_i^A, s_j^B \rangle, \langle s_k^A, s_l^B \rangle) = u_A(s_i^A, s_l^B) + u_B(s_j^B, s_k^A) \qquad (3)$$

The notions defined so far can straightforwardly be applied to asymmetric games, simply by symmetrizing the game first. It turns out that the characterization of evolutionary stability in asymmetric games is actually much simpler to characterize than in the general case:

Theorem 1: A strategy pair $\langle s_i^A, s_j^B \rangle$ of an asymmetric game is evolutionarily stable according to the best response dynamics if and only if it is a strict Nash equilibrium.

Recall that a pair of strategies is a **strict** Nash equilibrium iff each component is the **unique** best response. The proof of the theorem can be found in the appendix; the same result also holds for evolutionary stability in the replicator dynamics, as shown by Selten in [11].

3 From semantics to pragmatics

The main point to be made in this paper is that cautious deliberation, in the sense defined above, leads from conventionalized semantics, i.e. what is said, to the pragmatic content, i.e. what is meant. To illustrate this on an informal level with a standard case of scalar implicature, consider sentence (1).

(1) Some boys came to the party.

The conventionalized semantic strategy is that this sentence serves to convey the proposition that the set of boys coming to the party is non-empty, and this is how a non-rational hearer will interpret it. If we grant that the Gricean maxims ([2]) are somehow part of the utility function of the speaker, she will prefer the statement

(2) All boys came to the party.

if she beliefs that both sentences are true, and confine the usage of (1) to situations where some but not all boys came. A listener with strategic depth of 1 will anticipate this and conclude from (1) exactly this — that some but not all boys came. The scalar implicature of (1) that not all boys came is thus a part of the unique ESS that is reachable from the semantic convention via the best response dynamics. A similar story can be told for other cases of conversational implicatures, provided communication is appropriately modeled as a game, with an utility function that formalizes the Gricean maxims.

Following much work in game theoretic pragmatics (see for instance [15]), I will model communication as a version of a **signaling game** in the sense of Lewis ([5]).

In this setup, a game can be identified with a single utterance situation. Speaker and hearer are the players. Their actions are the production and interpretation of an utterance respectively, and their payoff preferences correspond to speaker economy and hearer economy.

To makes things more precise, let us assume that a fixed set of possible world W is given. The set of meanings M is a set of propositions over W, i.e. $M \subseteq POW(W)$. Furthermore, a set F of forms is given. A speaker strategy is any function s from W to F, i.e. a production grammar. Likewise, a hearer strategy is a comprehension grammar, i.e. a function h from F to M.

Let us thus assume that in each game, some random device, called **nature**, presents the speaker with a possible world w, which is not revealed to the hearer. The speaker than has to choose a form that is shown to the hearer and reveals as much information about w as possible. Nature's choice of w is probabilistic; w is drawn from W according to the probability distribution P_n, which is mutually known by the speaker and the hearer.[4]

I will choose a utility function for this general signaling game setup that formalizes, at least partially, the Gricean maxims. The overarching cooperativity principle translates into the assumption that communication is a game of cooperation. This means that the utility for speaker and hearer are always identical.

Next, let us assume that the hearer has a prior probability function P_H over W (which is also mutually known by speaker and hearer). Since P_n is known by the hearer, for a rational hearer it should hold that $P_n = P_H$. Since communication is pointless if it is mutually known that the hearer would not believe what the speaker is trying to say, I assume that $P_H(m) > 0$ for all $m \in M$. The information state of the hearer is his posterior probability distribution, after incorporating his interpretation of the signal that the speaker emits. This captures part of the maxim of quality — the hearer completely trusts what he is told. For a given speaker strategy s, hearer strategy h, and possible world w, the posterior distribution is $P(\cdot|h(s(w)))$. The number of bits of information that the hearer is still missing to achieve complete information is thus $(-\log_2 P(w|h(s(w))))$. ($\log_2$ is the binary logarithm here. The informativity in this sense can be seen as the number of yes/no-questions that the hearer still has to ask to figure out with certainty what the real

[4] For the sake of simplicity, I assume that W is finite. The model can straightforwardly be extended to an infinite set of possible worlds, if P_n is modeled as a probability density function. I refrain from doing this here for expository reasons.

world is like.)

If the hearer's interpretation $h(s(w))$ is false — $w \notin h(s(w))$ — this is infinite; otherwise it is the lower the more information $h(s(w))$ contains. So according to the quality and the quantity maxim, the speaker should strive to maximize $\log_2 P(w|h(s(w)))$ in each possible world w.[5] (Since utilities have to be finite, we assume that the hearer mistrusts the speaker with some sufficiently small amount η, which ensures that lying leads to an extremely low but still finite utility.)

The maxim of manner refers to the complexity of the form used. I thus assume that there is a cost function c from F to the positive real numbers. The players have an interest in keeping $c(f)$ low.

This game takes the form of a Bayesian game. This means that the utility does not depend just on the strategies of the players, but also on nature's choice. The cooperativity principle together with the maxims of quality, quantity and manner thus lead to the following utility function:

$$u(w, s, h) = \log_2 P_H(w|h(s(w))) - c(s(w)) \tag{4}$$

It is possible to incorporate the maxim or relevance here as well by relativizing utility further to certain decision problem (see for instance van Rooij's [16]). For the sake of simplicity, I will assume in this chapter that any information is relevant to the hearer.

A Bayesian game can be transformed into a strategic game in normal form by averaging over nature's choice:

$$u(s, h) = \sum_{w \in W} P_n(w) \log_2(P_H(w|h(s(w))) - c(s(w))) \tag{5}$$

4 Implicatures

In this section I will explore the behavior of the best response dynamics, given the kind of game that was defined in previous section. Let us first take up example (1) again. To formalize the situation, let us assume that there are just three different possible worlds that can be characterized by first order formulas (where B stands for "boy" and C for "came to the party"). To avoid complications relating to the existential presuppositions of the determiner "all", I assume that there exists boys in all possible worlds.

- $w_1 : \exists x.Bx \wedge \forall y.By \rightarrow Cy$
- $w_2 : \exists x.Bx \wedge Cx \wedge \exists y.By \wedge \neg Cy$
- $w_3 : \exists x.Bx \wedge \neg \exists y.By \wedge Cy$

[5] According to Robert van Rooij (see for instance [16]), this is the utility of a proposition provided every piece of information is equally relevant.

All three worlds are equally likely, both for nature and for the hearer:

$$P_{\{n,H\}}(w_i) = \frac{1}{3} \tag{6}$$

There are four possible forms that the speaker can choose from:

- f_1: "Some boys came to the party."
- f_2: "All boys came to the party."
- f_3: "No boys came to the party."
- f_4: "Some, but not all boys came to the party."

f_4 is more complex than the other forms, which all have about the same complexity. Let us say that

$$c(f_1) = c(f_2) = c(f_3) = 1 \tag{7}$$
$$c(f_4) = 3 \tag{8}$$

The semantic conventions of English relate forms to possible worlds. This corresponds to a pair of strategies: in each possible world, the speaker chooses a conventionally true form at random, and the hearer fully believes the conventional meaning of the form that he perceives. In the current game, this can be depicted as

$(s_0, h_0) =$

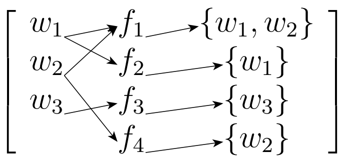

If symmetrized, the utility of this strategy pair against itself is about -3.22.

The best response of the speaker against the conventional hearer strategy is to map w_1 to f_2 (because then the hearer's posterior probability is 1 while the costs of f_1 and f_2 are identical), and to map w_2 to f_1 (which means a loss of .5 in informativity but a gain of 2 in costs). The best response of the hearer to the conventional speaker strategy is still the conventional hearer strategy. So the best response to the conventional strategy pair is

$(s_1, h_1) =$

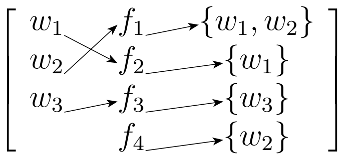

The utility of (s_1, h_1) against (s_0, h_0) is about -2.94, and the utility of (s_1, h_1) against itself is about -2.67. If a cautiously rational player decides to play (s_1, h_1) with a sufficiently small probability ε and otherwise the convention (s_0, h_0), the best response to this mixed strategy is still (s_1, h_1). However, after some finite number n of iterations,[6] the probability of (s_1, h_1) is large enough such that (s_2, h_2) becomes the best response. Nothing changes with regard to the speaker strategy, but in h_2 the hearer has figured out that the speaker utters f_1 if and only if w_2 is true; hence $\{w_2\}$ is the pragmatically informed interepretation of f_2.

$$(s_2, h_2) = \begin{bmatrix} w_1 \to f_2 & f_1 \to \{w_2\} \\ w_2 \to f_1 & f_2 \to \{w_1\} \\ w_3 \to f_3 & f_3 \to \{w_3\} \\ & f_4 \to \{w_2\} \end{bmatrix}$$

The utility of (s_2, h_2) against (s_1, h_1) is -2.333, while the utility of (s_2, h_2) against (s_0, h_0) is -3.0.

While (s_2, h_2) is the best response to a mixed strategy consisting predominantly of (s_1, h_1) and (s_2, h_2), it is still not a stable state. We have to take into account that the speaker beliefs with a small probability η that the speaker picks out a signal at random. If the probability of s_0 and s_1 — the speaker strategies where f_4 may be used to express $\{w_2\}$ — drops below a certain threshold, the best response for the hearer is to ignore f_4 altogether. This leads to

$$(s_4, h_4) = \begin{bmatrix} w_1 \to f_2 & f_1 \to \{w_2\} \\ w_2 \to f_1 & f_2 \to \{w_1\} \\ w_3 \to f_3 & f_3 \to \{w_3\} \\ & f_4 \to \{w_1, w_2, w_3\} \end{bmatrix}$$

This strategy pair is a strict Nash equilibrium and thus evolutionarily stable. Using the terminology from the previous section, we have shown that cautious deliberation leads from the conventional, semantic strategy pair (s_0, h_0) to the pragmatic equilibrium (s_4, h_4).

The transition from the semantic convention (s_0, h_0) to the pragmatically usable strategy pair (s_4, h_4) via the best response dynamics illustrates two important pragmatic phenomena. In s_2, the speaker expresses the fact that all boys come to the party with 100% certainty as f_2, "All boys came to the party", because this is, according to h_0, more specific than the equally true f_1, "Some boys came to the party". This is a consequence of the maxim of quantity. Anticipating this, the hearer pragmatically strengthens the interpretation of "Some boys came to the party" in h_2 to the interpretation $\{w_2\}$, *some but not all boys came to the party.* This is a scalar implicature, and the best response dynamics captures the intuitive reasoning that is used in Gricean pragmatics to explain this effect. Furthermore, the hearer figures out in h_4 that the signal f_4 is pragmatically sub-optimal for the speaker in all conceivable situations, because its conventional meaning can be conveyed via f_1 in a more economical way. Therefore this signal ceases to carry a pragmatic meaning and is ignored in the stable state. In the literature, this phenomenon is called **total blocking** (as opposed to **partial blocking**, where only part of the conventional

[6] where $n > \frac{\log(u(1,0)-u(1,1))-\log(u(1,0)-u(1,1)+u(2,1)-u(1,1))}{\log(1-\varepsilon)} \approx \frac{-2.783}{\log(1-\varepsilon)}$ for the particular numbers chosen here; $u(i, j)$ being the symmetrized utility of (s_i, h_i) against (s_j, h_j).

meaning of an expression is pragmatically blocked by a competing expression). A good example for this phenomenon are regular derivations that compete with underived words, like "pig" that cannot be used to refer to meat from pigs (while "chicken" or "lamb" can be used to refer to meat from the respective animals) because there is a special term, "pork" with precisely this meaning.

In the example at hand ("Some but not all boys came to the party"), this effect does actually not occur. This wrong prediction is related to the fact that I assumed it to be common knowledge that the speaker has perfect knowledge. In most situations, the hearer cannot be sure about this. If the speaker has possibly only partial knowledge, the scalar implicature only arises with a certain probability, not with certainty. This is sufficient to preempt total blocking.

It should be noted that not much hinges on the particular numbers chosen here. What is relevant is just the following inequality:

$$-\log_2(P_H(w1|\{w_1, w_2\})) > c(f_4) - c(f_1) \tag{9}$$

In some sense, this inequality compares incommensurable quantities, namely the differential informativity of two strategies versus that differential complexity of two expressions. The relative weight of informativity and complexity depends on various situational factors, and thus the inequality may actually be true or false depending on the context. A more complete model could model the relative importance of these two factors by some random factor that is itself a component of strategic reasoning.

5 Conclusion

Space does not permit to spell out the consequences of this approach to other pragmatic phenomena in detail. I will thus conclude by briefly setting the approach that was developed in this chapter into a broader context.

The basic idea to connect semantics and pragmatics via an iterated process of computing the best response in a signaling game setting is due to Robert Stalnaker ([14]). In the dynamics that Stalnaker uses, a strategy is entirely replaced by the best response to it. (In terms of my formalization, $\varepsilon = 1$ in Stalnaker's model.) The two models are not equivalent, but they coincide in many applications (including the example discussed above). Best response dynamics (in either version) shares the notion of a Nash equilibrium as a stable configuration in a strategic interaction with most alternative game theoretic solution concepts. However, the attractive feature of iterated best response (again in either version) is that equilibria can be grounded in strategy profiles that need not be in equilibrium, and that may even be strictly dominated (like the F-strategy in the extended stag hunt game above). This seems to be an important asset in many situations of strategic interaction, including communication, where salience and precedent single out a profile which may or may not be rationally justifiable. I hasten to add that I do believe that natural languages are in an evolutionarily stable state, at least with a very high probability. This applies to languages (in the sense of populations of utterances) as a whole though. A strategy profile which may be optimal on average may be non-rationalizable in a particular utterance situation. Best response dynamics thus serves to establish a link between the macro-structure and the micro-structure of linguistic communication.

References

[1] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, 2003.

[2] H. P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics 3: Speech Acts*, pages 41–58. Academic Press, New York, 1975.

[3] Josef Hofbauer. Stability for the best response dynamics. preprint, 1995.

[4] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, 1998.

[5] David Lewis. *Convention*. Harvard University Press, Cambridge, 1969.

[6] Akihiko Matsui. Best response dynamics and socially stable strategies. *Journal of Economic Theory*, 57:343–362, 1992.

[7] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, 1982.

[8] Peter J. Richerson and Robert Boyd. *Not by Genes Alone. How Culture Transformed Human Evolution*. The University of Chicago Press, Chicago and London, 2005.

[9] Ariel Rubinstein. *Modeling Bounded Rationality*. MIT Press, Cambridge (Mass.), 1998.

[10] Thomas C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge (Mass.), 1960.

[11] Reinhard Selten. A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology*, 84:93–101, 1980.

[12] Herbert Simon. *Models of Bounded Rationality*, volume 2. MIT Press, Cambridge (Mass.), 1982.

[13] Brian Skyrms. *Evolution of the Social Contract*. Cambridge University Press, Cambridge, UK, 1996.

[14] Robert Stalnaker. Saying and meaning, cheap talk and credibility. In Anton Benz, Gerhard Jäger, and Robert van Rooij, editors, *Game Theory and Pragmatics*, pages 83–100. Palgrave MacMillan, 2005.

[15] Robert van Rooij. Signalling games select Horn strategies. *Linguistics and Philosophy*, 27:493–527, 2004.

[16] Robert van Rooij. Utility, informativity, and protocols. *Journal of philosophical logic*, 33:389–419, 2004.

[17] Fernando Vega-Redondo. *Evolution, Games, and Economic Behaviour.* Oxford University Press, Oxford, 1996.

[18] H. Peyton Young. *Individual Strategy and Social Structure. An Evolutionary Theory of Institutions.* Princeton University Press, Princeton, 1998.

Appendix

Proof of Theorem 1: It is immediate from the definition that each strict Nash equilibrium is an ESS in the best response dynamics. So suppose $\langle s_i^A, s_j^B\rangle$ is a non-strict Nash equilibrium that is an ESS in the best response dynamics. This means that there is a strategy pair $\langle s_k^A, s_l^B\rangle \neq \langle s_i^A, s_j^B\rangle$ such that $u_A(s_k^A, s_j^B) = u_A(s_i^A, s_j^B)$ and $u_B(s_l^B, s_i^A) = u_B(s_j^B, s_i^A)$. Therefore either $i \neq k$ or $j \neq l$. Let us assume, without loss of generality, that $i \neq k$. Since $\langle s_i^A, s_j^B\rangle$ is an ESS, there must be a pair $\langle s_l^A, s_m^B\rangle$ with $u_A(s_l^A, s_j^B) = u_A(s_i^A, s_j^B)$ and $u_B(s_m^B, s_i^A) = u_B(s_j^B, s_i^A)$ such that $u_A(s_l^A, s_j^B) + u_B(s_m^B, s_k^A) > u_A(s_k^A, s_j^B) + u_B(s_j^B, s_k^A)$. Since both $\langle s_l^A, s_m^B\rangle$ and $\langle s_k^A, s_j^B\rangle$ are best responses to $\langle s_i^A, s_j^B\rangle$, $u_A(s_l^A, s_j^B) = u_A(s_k^A, s_j^B) = u_A(s_i^A, s_j^B)$. Hence $u_B(s_m^B, s_k^A) > u_B(s_j^B, s_k^A)$, and thus $m \neq j$.

We restrict the game to the sub-game G that results if all strategies are eliminated that are not best responses to $\langle s_i^A, s_j^B\rangle$. It follows from the previous paragraph that in the resulting sub-game, there are at least two A-strategies (including s_i^A) and at least two B-strategies (including s_j^B). The definition of ESS entails that $\langle s_i^A, s_j^B\rangle$ is the only Nash equilibrium in this sub-game.

We define an accessibility relation R between profiles in the following way: $R(\langle s_a^A, s_b^B\rangle, \langle s_c^A, s_d^B\rangle)$ iff either $s_a^A = s_c^A$ and and s_d^B is a best response to s_b^B, or $s_b^B = s_d^B$ and s_c^A is a best response to s_a^A. Let R^* be the reflexive and transitive closure of R. Now there are two options:

1. There is a profile x such that not $xR^*\langle s_i^A, s_j^B\rangle$. Then we can form the sub-game G' consisting of all the strategies in G that are components of profiles reachable from x via R^*. Neither s_i^A nor s_B^j belong to this G', because either each strategy in G is a best response either to s_i^A or to s_j^B by construction. According to Nash's theorem, G' has a Nash equilibrium. This equilibrium must simultaneously be a Nash equilibrium of G, since none of the excluded strategies is a best response to any of the strategies in G'. So G has a second Nash equilibrium besides $\langle s_i^A, s_j^B\rangle$, which is in contradiction with the assumption that $\langle s_i^A, s_j^B\rangle$ is an ESS.

2. $\langle s_i^A, s_j^B\rangle$ is reachable from any profile in G via R^*. This entails that there are strategies $s_o^A \neq s_i^A$ and $s_p^B \neq s_j^B$ such that either $R(\langle s_o^A, s_p^B\rangle, \langle s_i^A, s_p^B\rangle)$ and $R(\langle s_i^A, s_p^B\rangle,$

$\langle s_i^A, s_j^B\rangle)$, or $R(\langle s_o^A, s_p^B\rangle, \langle s_o^A, s_j^B\rangle)$ and $R(\langle s_o^A, s_j^B\rangle, \langle s_i^A, s_j^B\rangle)$. In the former case, $\langle s_i^A, s_p^B\rangle$ must be a Nash equilibrium of G, and likewise in the latter case $\langle s_o^A, s_j^B\rangle$.

So both scenarios lead to the conclusion that G has a Nash equilibrium besides $\langle s_i^A, s_j^B\rangle$, which is in contradiction with the assumtion that $\langle s_i^A, s_j^B\rangle$ is an ESS. We have thus proved that any asymmetric non-strict Nash equilibrium is not evolutionarily stable, or equivalently, that every asymmetric ESS is a strict Nash equilibrium. $\dashv$