Regret Equilibria in Games

Droste, E.J.R.; Kosfeld, M.; Voorneveld, M.

Publication date:
1998

Regret Equilibria in Games[1]

Edward Droste

Michael Kosfeld

Mark Voorneveld[2]

Department of Econometrics and CentER, Tilburg University, P.O.Box 90153, 5000 LE Tilburg, The Netherlands.

Abstract: We study boundedly rational players in an interactive situation. Each player follows a simple choice procedure in which he reacts optimally against a combination of actions of his opponents drawn at random from the distribution generated by a player's beliefs. By imposing a consistency requirement we obtain an equilibrium notion which we call *regret equilibrium*. An existence proof is provided and it is shown that the concept survives the iterated elimination of never-best responses. Additional properties are studied and the regret equilibrium concept is compared with other game theoretic solution concepts. The regret equilibrium concept is illustrated by means of interesting examples. It is shown that in the centipede game, players will continue to play with large probability.

Journal of Economic Literature Classification Number: C72

Keywords: equilibrium, regret, procedural rationality, bounded rationality, game theory.

[1]Preliminary version

[2]Corresponding author. E-mail: M.Voorneveld@kub.nl. This author's research is financially supported by the Dutch Foundation for Mathematical Research (SWON) through project 613-04-051.

1 Introduction

Traditional economic analysis rests on two fundamental assumptions. First, economic agents have a particular goal, e.g. utility or profit maximization. Second, economic agents behave in such a way that these goals are achieved within the limits imposed by given constraints. Simon (1976) refers to this kind of behavior as being substantively rational. Whenever the assumptions with respect to utility or profit maximization and substantive rationality hold, incorporating behavioral considerations in economic analysis is superfluous.

Economics became concerned with human cognitive processes originated in psychology when uncertainty and expectations became an explicit part of the models. Cournot, for example, showed that the notion of profit-maximization is ill-defined in an oligopoly market. The quantity that would be substantively rational for each firm depends on the choices made by the other firms: no firm can choose without making assumptions about how others will choose. Thus, the presence of uncertainty prevents them from acting substantively rational. However, as stated by Simon (1976), being boundedly rational does not prevent the firms from being procedurally rational. Procedural rationality means that behavior is the outcome of some process of appropriate reasoning given available knowledge and computational capabilities. See also Rubinstein (1997).

In games players make choices based on an association between actions and consequences. We consider a game theoretic model in which the players behave according to a very simple procedure of deliberation. Facing a certain situation, i.e. given a pure strategy profile of the opponents this procedure prescribes that each best response is played with equal probability. If a player faces a mixed strategy profile of his opponents, he is confronted with uncertainty. In those situations regret considerations may become important. A player ends up feeling regret in case the chosen action was not a best response ex post. We focus on static, one-shot games where a mixed strategy means that each player does choose a pure strategy, but this pure strategy is drawn from the probability distribution over his strategy set induced by the player's mixed strategy. We assume that a player has beliefs about the mixed strategy profile of his opponents. Thus, a player can determine the probability distribution over the pure strategy profiles of his opponents from these beliefs. Each player assumes that the realized pure strategy profile of his opponents is randomly drawn from this distribution. Given the procedure described above, he would react to such a realization by playing each of his best responses with equal probability. We now assume that he uses his beliefs to weigh each of these reactions accordingly. Thus, the procedure together with the player's beliefs about the mixed strategies of his opponents determines the mixed strategy of the player. The equilibrium requirement of our concept is that the beliefs of each player about the other players' behavior are consistent with the other players' behavior. Consequently, in equilibrium each player plays a pure strategy with a probability equal to the probability that this strategy does not give rise to regret.

The notion of Nash equilibrium and the equilibrium concept introduced by Osborne and Rubinstein (1997) also require the beliefs to be consistent with the other players' behavior. However, they differ with respect to the way in which players construct beliefs about the consequences of their actions. The notion of Nash equilibrium prescribes that each player constructs probabilistic beliefs about the other players' behavior. Furthermore, given these beliefs each player associates a distribution of consequences to each action. In the equilibrium concept of Osborne and Rubinstein each player constructs beliefs about the consequences of his actions by sampling each action and associating a consequence with each action independently.

The concept of Osborne and Rubinstein (1997) aims at constructing a procedure in case the players have limited knowledge of the economic environment. Contrary to both the Nash equilibrium and our concept, they do not require players to know the relationship between their own action, the other players' actions, and the outcome. Their players only need to know their own set of possible actions. However, Osborne and Rubinstein do assume players to be able to compare outcomes resulting from different action profiles of the other players. This is not necessary in our model where a very simple preference relation for each player suffices. In particular, our players do not have to be able to compare lotteries over outcomes or even outcomes resulting from different action profiles of the opponents. Furthermore, given the actions of the opponents, we allow for the possibility that a player cannot compare some of his actions; he only must be able to determine all best responses.

As mentioned before, regret considerations become important whenever people face uncertain situations. In a decision theoretic set-up such situations are modelled by lotteries, see Loomes and Sudgen (1982). Their aim was to give a simple explanation of the observations in the experiments of Kahneman and Tversky (1979). Regret considerations in a dynamic game theoretic setting are studied by Hart and Mas-Colell (1997).

The structure of the paper is as follows. In Section 2 we define terminology that is used throughout the paper. The regret equilibrium notion is derived from the corresponding procedure in Section 3. This section also includes the existence result and some results concerning weakly and strictly dominated strategies. Section 4 concerns the structure and size of the set of regret equilibria. In Section 5 we compare our concept with other solution concepts in noncooperative game theory. Section 6 applies the equilibrium concept to several classes of games, including a Rock-Scissors-Paper game, a class of Hawk-Dove games, and two-person coordination games. We proceed by applying the regret equilibrium concept to extensive form games, in particular the centipede game. It is shown that the equilibrium concept provides a potential resolution of the paradoxical flavor of this game: The equilibrium conditions indicate that the players understand the structure of the centipede game, but still want to continue at the early stages of the game with a relatively large probability. Section 7 concludes.

2 Preliminaries

This section defines some standard game theoretic notions which are used hereafter. A *(strategic) game* is a tuple $G = \langle N, (S_i)_{i\in N}, (\succeq_i)_{i\in N}\rangle$, where $N = \{1, \ldots, n\}$ is a finite set of players; each player $i \in N$ has a finite set S_i of pure strategies and a binary relation $\succeq_i$ over $\prod_{i\in N} S_i$, reflecting his preferences over the outcomes. The binary relation $\succeq_i$ is assumed to be reflexive and its asymmetric part $\succ_i$, defined for all $s, t \in \prod_{i\in N} S_i$ by

$$s \succ_i t \Leftrightarrow s \succeq_i t \text{ and not } t \succeq_i s,$$

is assumed to be acyclic. In the following we also consider cases in which the preference relations $\succeq_i$ induce von Neumann-Morgenstern utility functions $\pi_i : \prod_{i\in N} S_i \to \mathbb{R}$ and denote the corresponding game by $G = \langle N, (S_i)_{i\in N}, (\pi_i)_{i\in N}\rangle$. For notational convenience we write $S = \prod_{i\in N} S_i$, $S_{-i} = \prod_{j\in N\setminus\{i\}} S_j$; for a strategy tuple $s = (s_1, \ldots, s_n) \in S$ we denote $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ and, with a slight abuse of notation, $s = (s_i, s_{-i})$. We denote by

$$\Delta_i := \{\sigma_i : S_i \to \mathbb{R} \mid \forall s_i \in S_i : \sigma_i(s_i) \geq 0, \sum_{s_i\in S_i} \sigma_i(s_i) = 1\}$$

the set of mixed strategies for player i; analogous to the pure strategy case, we use notations Δ, Δ_{-i}, $\sigma = (\sigma_i, \sigma_{-i})$. For a mixed strategy profile σ_{-i}, we write $\sigma_{-i}(s_{-i}) := \prod_{j\in N\setminus\{i\}} \sigma_j(s_j)$, the probability that the opponents of player i play the strategy profile $s_{-i} \in S_{-i}$.

Consider a game $\langle N, (S_i)_{i\in N}, (\succeq_i)_{i\in N}\rangle$. Denote for each player $i \in N$ and each profile $s_{-i} \in S_{-i}$ of pure strategies of his opponents the set of pure best replies, i.e., the pure strategies that player i cannot improve upon, by $B_i(s_{-i})$:

$$B_i(s_{-i}) := \{s_i \in S_i \mid \nexists \tilde{s}_i \in S_i : (\tilde{s}_i, s_{-i}) \succ_i (s_i, s_{-i})\}.$$

Of course, for games $\langle N, (S_i)_{i\in N}, (\pi_i)_{i\in N}\rangle$ with utility functions we have:

$$B_i(s_{-i}) := \{s_i \in S_i \mid \forall \tilde{s}_i \in S_i : \pi_i(s_i, s_{-i}) \geq \pi_i(\tilde{s}_i, s_{-i})\}.$$

Since S_i is finite and $\succ_i$ is acyclic, $B_i(s_{-i})$ is nonempty. In a game $\langle N, (S_i)_{i\in N}, (\succeq_i)_{i\in N}\rangle$, we call a pure strategy $s_i \in S_i$ a *never-best response* if

$$\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i})\} = \emptyset.$$

For a game $\langle N, (S_i)_{i\in N}, (\pi_i)_{i\in N}\rangle$ we have that $s_i \in S_i$ is a never-best response if

$$\pi_i(s_i, s_{-i}) < \max_{\tilde{s}_i\in S_i} \pi_i(\tilde{s}_i, s_{-i})$$

for each $s_{-i} \in S_{-i}$. A pure strategy $s_i \in S_i$ is *weakly dominated* by a mixed strategy $\sigma_i \in \Delta_i$ if

$$\forall s_{-i} \in S_{-i} : \pi_i(\sigma_i, s_{-i}) \geq \pi_i(s_i, s_{-i})$$

with strict inequality for at least one s_{-i}, and *strictly dominated* if all inequalities are strict. A strictly dominated strategy is clearly a never-best response.

3 Model

The players in our model base their mixed strategy on their beliefs about the strategy profile of the opponents and on a very simple procedure. Each player $i \in N$ has beliefs about the strategy profile of his opponents, in the sense that he believes his opponents adopt a strategy profile $\sigma_{-i} \in \Delta_{-i}$. We focus on one-shot games, where a mixed strategy means that each player *does* play a pure strategy, but this pure strategy is drawn from the probability distribution over his strategy set induced by a player's mixed strategy. Hence, the player is uncertain about the exact pure strategy profile $s_{-i} \in S_{-i}$ that will be chosen by his opponents. In the Nash equilibrium context, players use their beliefs to calculate expected payoffs. In our model, players reason differently. Since they are playing a one-shot game, they are sure that in the end they will face a pure strategy profile of the opponents. If they do not play a best response to this pure strategy profile they will feel regret. This consideration is due to the uncertainty of the situation that is coming from the beliefs that the players have. We assume that each player $i \in N$ adopts a very simple procedure. Given a pure strategy profile s_{-i} of his opponents, he will play each best response to s_{-i} with equal probability. Each player assumes that the realized pure strategy profile of his opponents is drawn from the probability distribution over the pure strategy profiles that is induced by his beliefs. Thus, using this procedure, a player would be able to react to each realization ex post. However, ex ante the realized pure strategy profile is stochastic. Players now use their beliefs about the profiles of the opponents to weigh their reactions accordingly. Consequently, a player plays each of his pure strategies with a certain probability, that is an outcome of the procedure together with his beliefs. The equilibrium requirement of our concept is that the beliefs of each player about the other players' behavior are consistent with the other players' behavior: we require correctness of beliefs.

We first give an example and then derive the equilibrium conditions in the general case. Consider the two-player game in Figure 1, where $\epsilon \in (0, 1)$. Player 1 has two pure strategies, A and B. Player 2 has three pure strategies, a, b, and c. We ignore the payoffs to player 2

	a	b	c
A	$1, 0$	$2, 0$	$3, 0$
B	$3 + \epsilon, 0$	$1 + \epsilon, 0$	$2 + \epsilon, 0$

Figure 1: Regret considerations

in this example and focus on the situation for player 1. Assume that player 1 believes that his opponent will play each of his pure strategies with equal probability, i.e., that his opponent plays $\sigma_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Hence, player 1 considers each of the realizations a, b, and c to be equally likely. If the observation is a, his unique best response is to play B, in the other cases he should play his unique best response A. According to the procedure and the beliefs, the probability that player 1 will play A is then $\frac{2}{3}$, and that he will play B is $\frac{1}{3}$, which determines his mixed strategy.

In general, what equilibrium conditions does the above imply? Consider a strategy profile $\sigma \in \Delta$ and a player $i \in N$. What is the probability that player i will choose to play a strategy $s_i \in S_i$ if he follows the above procedure and believes that his opponents play according to σ_{-i}? Player i will not play s_i in response to a pure strategy profile s_{-i} to which s_i is not a pure best reply. Now suppose that s_i is indeed a pure best reply to s_{-i}; the probability of event s_{-i}, i.e., the probability that i's opponents play s_{-i}, equals $\sigma_{-i}(s_{-i})$. Given this event, he will play all pure best replies to s_{-i} with equal probability and consequently will play s_i with probability $\frac{1}{|B_i(s_{-i})|}$. Summarizing, given a profile σ_{-i} of player i's opponents, the probability that player i decides to play s_i if he follows the procedure and believes that his opponents play according to σ_{-i} equals

$$\sum_{\{s_{-i} \in S_{-i} | s_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) =: r_i(s_i, \sigma_{-i}), \tag{1}$$

where the empty sum is zero by definition. In equilibrium, every player's beliefs must be correct, so expression (1) should equal player i's probability of playing s_i, $\sigma_i(s_i)$.

The expression for $r_i(s_i, \sigma_{-i})$ in (1) can be interpreted as player i's probability of not feeling regret after playing s_i. Necessary for not feeling regret of s_i is that s_i is a best reply to the played strategy. If there were more than one pure best reply, he feels regret for not being able to play the other pure best replies. This last type of regret might seem strange, but it can be compared to the regret of not choosing an equally pleasing destination for your vacation: you are aware that some alternative destinations are just as good, but it is still unpleasant that you cannot be in more places at the same time.

Notice that in Figure 1 the probability that player 1 feels no regret if he plays A equals $\frac{2}{3}$; if he plays B it equals $\frac{1}{3}$. Given beliefs $\sigma_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the expected payoff of playing A is 2 and of playing B, $2+\epsilon$. Hence an expected utility maximizing player 1 would select B with probability one.

Summarizing, we define our equilibrium notion as follows:

Definition 3.1 *Let $G = \langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$ be a game. A mixed strategy profile $\sigma \in \Delta$ is a* regret equilibrium *if for every player $i \in N$ and for every $s_i \in S_i$:*

$$\sigma_i(s_i) = r_i(s_i, \sigma_{-i}). \tag{2}$$

The set of regret equilibria of a game G is denoted by $RE(G)$.

The following proposition shows that regret equilibria exist for every game.

Proposition 3.2 *Let $G = \langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$ be a game. Then $RE(G) \neq \emptyset$.*

Proof. Let $i \in N, \sigma \in \Delta$. Notice that

$$\begin{aligned} \sum_{s_i \in S_i} r_i(s_i, \sigma_{-i}) &= \sum_{s_i \in S_i} \sum_{\{s_{-i} \in S_{-i} | s_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\ &= \sum_{s_{-i} \in S_{-i}} \sum_{s_i \in B_i(s_{-i})} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\ &= \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i}) \\ &= 1. \end{aligned}$$

Hence the mapping

$$\begin{aligned} r: \quad & \Delta \to \Delta \\ & \sigma \mapsto r(\sigma) \end{aligned}$$

with $r(\sigma)_i(s_i) = r_i(s_i, \sigma_{-i})$ is well-defined. In the definition (1) of the function r_i neither the index set in the summation sign nor the number $|B_i(s_{-i})|$ of pure best responses depends on the strategy combination σ. Hence, this mapping is obviously continuous. Application of the Brouwer fix point theorem yields the existence of a strategy profile $\sigma \in \Delta$ such that $\sigma = r(\sigma)$, which is a regret equilibrium. □

Remark 3.3 It follows from the proof that $\sum_{s_i \in S_i} \sigma_i(s_i) = 1 = \sum_{s_i \in S_i} r_i(s_i, \sigma_{-i})$ for each $\sigma \in \Delta, i \in N$. As a consequence, one of the conditions $\sigma_i(s_i) = r_i(s_i, \sigma_{-i})$ of player i is redundant.

A game H is said to be *obtained by iterated elimination of never-best responses* from a game $G = \langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$ if there exists a number $k \in \mathbb{N}$ of elimination rounds and for each player $i \in N$ a collection of sets $S_i^0, S_i^1, \ldots, S_i^k$ and a sequence $\succeq_i^0, \succeq_i^1, \ldots, \succeq_i^k$ of relations such that:

1. For each player $i \in N : S_i = S_i^0 \supseteq S_i^1 \supseteq \cdots \supseteq S_i^k$;
2. For each player $i \in N$ and each $l = 0, 1, \ldots, k$: $\succeq_i^l$ is the preference relation $\succeq_i$ from the game G restricted to $\prod_{j \in N} S_j^l$;
3. For each $l = 0, 1, \ldots, k-1$ there exists a player $i \in N$ such that $S_i^l \setminus S_i^{l+1}$ is nonempty and contains only never-best responses of player i in the game $\langle N, (S_i^l)_{i \in N}, (\succeq_i^l)_{i \in N} \rangle$;
4. H is the game $\langle N, (S_i^k)_{i \in N}, (\succeq_i^k)_{i \in N} \rangle$;
5. In the game H, no player $i \in N$ has never-best responses.

The behavior of the regret equilibrium concept with respect to dominated strategies and elimination thereof is summarized in the next result.

Proposition 3.4 *The following results hold:*

(i) In a regret equilibrium σ^* of a game $\langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$ never-best responses are played with zero probability.

Moreover,

(ii) the set of regret equilibria of a game $G = \langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$ equals — up to zero probability assigned to eliminated pure strategies — the set of regret equilibria of a game that is obtained by iterated elimination of never-best responses.

Finally,

(iii) let $G = \langle N, (S_i)_{i \in N}, (\pi_i)_{i \in N} \rangle$ be a game with von Neumann-Morgenstern utility functions and let σ^* be a regret equilibrium of G. If player i's pure strategy s_i is weakly dominated by the mixed strategy σ_i, then:

$$\text{for all } \bar{s}_i \in S_i: \text{ if } \sigma_i(\bar{s}_i) > 0, \text{ then } \sigma_i^*(\bar{s}_i) \geq \sigma_i^*(s_i).$$

Proof. The proof of (i) is easy: if $s_i \in S_i$ is a never-best response, then the set $\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i})\}$ is empty and hence according to (2): $\sigma_i^*(s_i) = r_i(s_i, \sigma_{-i}^*) = 0$.

To prove (ii), it suffices to prove that the first round of eliminations does not change the equilibrium set, since the proof can then be repeated for the additional rounds. Assume for simplicity that in this first elimination round we eliminate all the never-best responses

$$NB_i := \{s_i \in S_i \mid s_i \text{ is a never-best response of player } i \text{ in } G\}$$

of each player $i \in N$, thus obtaining a smaller game G'. The equilibrium conditions in the game G are that for each $i \in N$ and each $s_i \in S_i$:

$$\begin{aligned}
\sigma_i(s_i) &= r_i(s_i, \sigma_{-i}) \\
&= \sum_{\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\
&= \sum_{\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i}) \text{ and } \forall j \in N \setminus \{i\}: s_j \notin NB_j\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\
&+ \sum_{\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i}) \text{ and } \exists j \in N \setminus \{i\}: s_j \in NB_j\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\
&= \sum_{\{s_{-i} \in \prod_{j \in N \setminus \{i\}} S_j \setminus NB_j \mid s_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma_{-i}(s_{-i}) \\
&+ \sum_{\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i}) \text{ and } \exists j \in N \setminus \{i\}: s_j \in NB_j\}} \frac{1}{|B_i(s_{-i})|}
\end{aligned}$$

By (i), strategies $s_j \in NB_j$ are played with zero probability in a regret equilibrium. Hence the second sum in the last equality above equals zero. What remains, for each player $i \in N$ and each strategy $s_i \in S_i \setminus NB_i$, are exactly the equilibrium conditions for the game G'.

To prove (iii), assume w.l.o.g. that $s_i \in B_i(s_{-i})$. Since σ_i weakly dominates s_i and $s_i \in B_i(s_{-i})$, for every $\bar{s}_i \in S_i$ such that $\sigma_i(\bar{s}_i) > 0$ we must have that $\bar{s}_i \in B_i(s_{-i})$:

$$\{s_{-i} \in S_{-i} \mid s_i \in B_i(s_{-i})\} \subseteq \{s_{-i} \in S_{-i} \mid \bar{s}_i \in B_i(s_{-i})\},$$

which together with the definition of $r_i(\cdot, \sigma^*_{-i})$ implies the result:

$$\begin{aligned}
\sigma_i(\bar{s}_i) > 0 \Rightarrow \sigma^*_i(\bar{s}_i) &= r_i(\bar{s}_i, \sigma^*_{-i}) \\
&= \textstyle\sum_{\{s_{-i} \in S_{-i} | \bar{s}_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma^*_{-i}(s_{-i}) \\
&\geq \textstyle\sum_{\{s_{-i} \in S_{-i} | s_i \in B_i(s_{-i})\}} \frac{1}{|B_i(s_{-i})|} \sigma^*_{-i}(s_{-i}) \\
&= r_i(s_i, \sigma^*_{-i}) \\
&= \sigma^*_i(s_i).
\end{aligned}$$

□

Notice that the result above does not rule out that weakly dominated strategies are played with positive, even quite large probability.

Example 3.5 Consider the game in Figure 2. T weakly dominates B and L strictly dominates R. T is a best response against L, but player 1 will regret not playing the other best response B against L, and T is a unique best response against R. Hence in equilibrium we have the condition that

$$\sigma_1(T) = \frac{1}{2}\sigma_2(L) + \sigma_2(R).$$

The condition for $\sigma_1(B)$ is redundant, since the probabilities have to add up to one. Similarly, for player 2 we see that L is a unique best response to both T and B, so that his equilibrium condition becomes

$$\sigma_2(L) = \sigma_1(T) + \sigma_1(B).$$

Solving these equations and taking into acount that $(\sigma_1, \sigma_2) \in \Delta_1 \times \Delta_2$ we find that the unique regret equilibrium equals $((\frac{1}{2}, \frac{1}{2}), (1, 0))$. Observe that the weakly dominated strategy is not only played with positive probability, but that there is not even an alternative strategy with a higher probability.

	L	R
T	$1,1$	$1,0$
B	$1,1$	$0,0$

Figure 2: The game from Example 3.5

4 The size and structure of the set of regret equilibria

The size of an equilibrium set can be seen as a measure of the cutting power of an equilibrium concept. With respect to the size of the set of regret equilibria of a game $\langle N, (S_i)_{i \in N}, (\succeq_i)_{i \in N} \rangle$, remark that it is always a relatively small subset of Δ. A strategy tuple $\sigma_{-i} \in \Delta_{-i}$ completely determines $r_i(\cdot, \sigma_{-i})$ and hence in an n-player game it suffices to know only $n-1$ components of a regret equilibrium to compute the equilibrium strategy for the remaining n-th player. This implies that $RE(G)$ is always of lower dimension than Δ. In particular, it is impossible that $RE(G) = \Delta$.

The structure of the set of Nash equilibria has been studied by several authors, including Winkels (1979) and Jansen (1981), who show that in two-person games the set of Nash equilibria has a nice decomposition into a finite number of polytopes. Concerning the structure of the set of regret equilibria, we see that if the game G has only two players, then $RE(G)$ is a polytope, since the set of regret equilibria is then determined by finitely many linear equations and linear weak inequalities in the variables $(\sigma_i(s_i))_{i \in N, s_i \in S_i}$. If the game has at least three players, its set of regret equilibria is determined by a set of polynomial equations over a Cartesian product of simplices. This leads to the observations that — analogous to the set of Nash equilibria — the set of regret equilibria may be curved or disconnected. The following two examples indicate that both possibilities indeed occur.

Example 4.1 Consider the three player game in Figure 3. Here we denote by $p, q, r \in [0, 1]$ the probability with which player 1 chooses his first row, player 2 chooses his first column, and player 3 chooses his first matrix, respectively.

	q	$1-q$
p	$1,1,1$	$1,0,0$
$1-p$	$0,1,1$	$0,0,0$

r

	q	$1-q$
p	$1,0,0$	$0,1,1$
$1-p$	$0,0,0$	$1,1,1$

$1-r$

Figure 3: A game with a curved set of regret equilibria

Considering Remark 3.3, it suffices to determine an equilibrium constraint only for p, q, and r, since those for $1-p, 1-q, 1-r$ will follow immediately. The first strategy (the top row) of player 1 is a unique best response to three combinations of pure strategies of his opponents, namely to those in which player 2 chooses either his first or his second column and player 3 chooses the first matrix, which occurs with probability $qr + (1-q)r$, and to the strategy in which player 2 chooses his first column and player 3 chooses the second matrix, which occurs with probability $q(1-r)$. Together with the constraints for the other two players, we find that the conditions for a regret equilibrium are

$$\begin{cases} p &=& qr+(1-q)r+q(1-r) &=& q+(1-q)r \\ q &=& pr+(1-p)r &=& r \\ r &=& pq+(1-p)q &=& q \\ p,q,r &\in& [0,1] \end{cases}$$

Consequently, the set of regret equilibria equals

$$\{((p,1-p),(q,1-q),(r,1-r)) \mid p = q(2-q), r = q, q \in [0,1]\},$$

which is a curved equilibrium set.

Example 4.2 Consider the three player game in Figure 4.

	q	$1-q$
p	$1,1,1$	$0,0,0$
$1-p$	$0,0,1$	$1,1,0$

r

	q	$1-q$
p	$0,0,0$	$1,1,1$
$1-p$	$1,1,0$	$0,0,1$

$1-r$

Figure 4: A game with a disconnected set of regret equilibria

The conditions for a regret equilibrium are

$$\begin{cases} p &=& qr+(1-q)(1-r) &=& 2qr-q-r+1 \\ q &=& pr+(1-p)(1-r) &=& 2pr-p-r+1 \\ r &=& pq+(1-p)q &=& q \\ p,q,r &\in& [0,1] \end{cases}$$

This is equivalent with (after substitution of $r = q$):

$$\begin{cases} p &=& 2q^2-2q+1 \\ q &=& 2pq-p-q+1 \\ r &=& q \\ p,q,r &\in& [0,1] \end{cases}$$

Subtracting the first equality from the second, we find:

$$\begin{cases} p &=& 2q^2-2q+1 \\ pq &=& q^2 \\ r &=& q \\ p,q,r &\in& [0,1] \end{cases}$$

Hence

$$\begin{cases} p &=& 2q^2-2q+1 \\ q &=& 0 \\ r &=& q \\ p,q,r &\in& [0,1] \end{cases} \quad \text{or} \quad \begin{cases} p &=& 2q^2-2q+1 \\ p &=& q \\ r &=& q \\ p,q,r &\in& [0,1] \end{cases}$$

Consequently, the set of regret equilibria equals

$$\{((1,0),(0,1),(0,1))\} \cup \{((\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}))\} \cup \{((1,0),(1,0),(1,0))\},$$

consisting of three components.

5 Comparison with other solutions

Osborne and Rubinstein (1997) introduce a procedure in which a player samples all his pure strategies once and chooses the one which yields — in that particular sample — the highest payoff (breaking ties equi-probably). Comparing this to our procedure, we observe that in the Osborne and Rubinstein procedure a player only needs to know his set of pure strategies, but he must be able to compare all outcomes in S. In our procedure, *the players must know more, but need to be able to compare less* than in their procedure. In particular,

- the players must know the possible outcomes $S = \prod_{i \in N} S_i$,
- they are required only to be able to perform a very simple kind of comparative statics: given a profile $s_{-i} \in S_{-i}$ of pure strategies for his opponents, player i only needs to determine whether or not a strategy $s_i \in S_i$ can be improved upon.

This very limited type of comparisons is similar to those made by a player in the notion of pure Nash equilibrium. Notice that a very simple preference relation for a player suffices: It is not necessary to compare lotteries over outcomes, not even outcomes that arise facing different strategies of the opponents. This last point is particularly important: The players are not required to compare outcomes arising from very 'exotic' what-if situations. Only responses to the same environment, i.e., to the same combination of pure strategies of their opponents, are compared. This allows the nature of the payoff (money, days of vacation, ...) to be different for each combination of strategies of the opponents.

Moreover, even if the profile of pure strategies of the opponents is fixed, a player is not required to be able to compare all his strategies, only to be able to determine all best responses. Hence even the order on $\{(s_i, s_{-i}) \mid s_i \in S_i\}$ need not be complete.

Chen, Friedman, and Thisse (1997), using a model in which players have subconscious utility functions, consider socalled boundedly rational Nash equilibria with equilibrium requirements similar to those in our model.

Despite the relatively prudent behavior with respect to (weakly) dominated strategies as expressed in Proposition 3.4, the set of regret equilibria and Nash equilibria have no obvious relation. In the game of Figure 5, for instance, the unique Nash equilibrium equals $((\frac{1}{3},\frac{2}{3}),(\frac{2}{3},\frac{1}{3}))$, the unique RE equals $((\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}))$. We can, however, indicate a relation with the notion of *strict equilibria* introduced by Harsanyi (1973) as those strategy profiles σ satisfying the

	L	R
T	$0,2$	$2,0$
B	$1,0$	$0,1$

Figure 5: The Nash and regret equilibrium concept differ

condition that each player plays his unique best response to the strategies of the opponent:

$$\forall i \in N : \{\sigma_i\} = \{\tau_i \in \Delta_i \mid \nexists \tilde{\tau}_i : \pi_i(\tilde{\tau}_i, \sigma_{-i}) > \pi_i(\tau_i, \sigma_{-i})\}.$$

It is clear that a strict Nash equilibrium is always a pure strategy Nash equilibrium and (consequently) that strict Nash equilibria not always exist. However, if they exist, they are exactly the pure strategy regret equilibria of the game.

Proposition 5.1 *The set of strict Nash equilibria of a game $\langle N, (S_i)_{i \in N}, (\pi_i)_{i \in N}\rangle$ coincides with the set of pure strategy regret equilibria.*

The proof is straightforward and left to the reader.

The results with respect to the iterated elimination of never-best responses in Proposition 3.4 call to mind the notion of rationalizability introduced in Bernheim (1984) and Pearce (1984). Without going into the formal definitions, it follows immediately from Proposition 3.4 and Bernheim (1984, pp. 1015-1016) that every action that is played with positive probability in a regret equilibrium is rationalizable. However, in a regret equilibrium σ, the mixed strategies σ_i themselves need not be rationalizable.

6 Examples

In this section we apply the concept of a regret equilibrium to several classes of games, including two-person coordination games and a class of Hawk-Dove games. Moreover, one can apply the concept of a regret equilibrium to the reduced strategic form of extensive games. We present one brief example and one more elaborate case, in which we solve a T-choice centipede game.

	R, q_1	S, q_2	$P, 1-q_1-q_2$
R, p_1	$0,0$	$1,-1$	$-1,1$
S, p_2	$-1,1$	$0,0$	$1,-1$
$P, 1-p_1-p_2$	$1,-1$	$-1,1$	$0,0$

Figure 6: Rock, Scissors, Paper

Example 6.1 Consider the Rock, Scissors, Paper game in Figure 6, where R, S, P, have the obvious meaning and the corresponding probabilities with which these strategies are played are denoted by p_i, q_i. The conditions for a regret equilibrium are

$$\begin{cases} p_1 & = & q_2 \\ p_2 & = & 1 - q_1 - q_2 \\ q_1 & = & p_2 \\ q_2 & = & 1 - p_1 - p_2 \\ p_1, p_2, q_1, q_2 & \in & [0, 1] \\ p_1 + p_2 & \leq & 1 \\ q_1 + q_2 & \leq & 1 \end{cases}$$

Simple calculus leads to the conclusion that the unique regret equilibrium equals the unique Nash equilibrium in which both players choose each of their pure strategies with probability $\frac{1}{3}$.

Example 6.2 A two-player game is a coordination game if both players have the same set of pure strategies and the unique best response to a pure strategy of the opponent is to play the same strategy. An example of a coordination game is the Battle of the Sexes game given in Figure 7. From the definition of a coordination game it is clear that a player does not regret

	boxing	ballet
boxing	$3, 2$	$0, 0$
ballet	$0, 0$	$2, 3$

Figure 7: Battle of the Sexes; a coordination game

playing a pure strategy if and only if his opponent plays the same strategy. As a consequence, a profile of mixed strategies is a regret equilibrium if and only if both players play the same mixed strategy. This illustrates an important difference with the Nash equilibrium concept: The pure Nash equilibria of a coordination game are the combinations of pure strategies in which the players indeed coordinate (choose the same pure strategy). Since these Nash equilibria are strict, they are also regret equilibria. However, there may be mixed strategy Nash equilibria in which the players do not coordinate exactly. In the example above, the mixed strategy Nash equilibrium is $((\frac{2}{5}, \frac{3}{5}), (\frac{3}{5}, \frac{2}{5}))$. The regret equilibrium concept provides — in our opinion — a much more intuitive solution.

Example 6.3 In this example we consider a class of Hawk-Dove games with the structure of the payoff matrix given in Figure 8. Here V and W are real numbers satisfying the condition $W < V$. We consider several cases.

	q	$1-q$
p	V, V	$0, 2V$
$1-p$	$2V, 0$	W, W

Figure 8: A class of Hawk-Dove games

1. If $W > 0$, we have a Prisoner's dilemma; both players have a strictly dominating strategy. The unique regret equilibrium equals $((0,1),(0,1))$.

2. If $W = 0$, the conditions for a regret equilibrium are

$$\begin{cases} p &= \frac{1}{2}(1-q) \\ q &= \frac{1}{2}(1-p) \\ p,q &\in [0,1] \end{cases}$$

   The regret equilibrium is $((\frac{1}{3},\frac{2}{3}),(\frac{1}{3},\frac{2}{3}))$.

3. If $W < 0 < V$, we have a Chicken game. The conditions for a regret equilibrium are

$$\begin{cases} p &= 1-q \\ q &= 1-p \\ p,q &\in [0,1] \end{cases}$$

   The set of regret equilibria is $\{((p,1-p),(1-p,p)) \mid p \in [0,1]\}$.

4. If $V = 0$, the conditions for a regret equilibrium are

$$\begin{cases} p &= \frac{1}{2}q + (1-q) \\ q &= \frac{1}{2}p + (1-p) \\ p,q &\in [0,1] \end{cases}$$

   The regret equilibrium is $((\frac{2}{3},\frac{1}{3}),(\frac{2}{3},\frac{1}{3}))$.

5. If $V < 0$, the first strategy of both players strictly dominates the second, so the unique regret equilibrium is $((1,0),(1,0))$.

Example 6.4 Consider the extensive form game in Figure 9. In this game, player 1 is given the choice to stop (S) or continue (C). If he continues, player 2 is given the same choice. The game ends if either player decides to stop or both decide to continue.

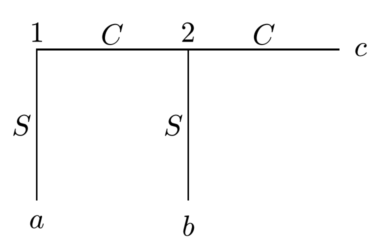

Figure 9: An extensive form game

Assume that $c \succ_1 a$ and $c \succ_2 b$. Consequently, we have that the outcome c is the unique subgame perfect equilibrium of the game. Denote by p the probability that player 1 chooses to stop and by q the probability that player 2 chooses to stop. By Remark 3.3, it suffices to determine the equilibrium conditions for p and q. Player 2's choice to stop is not a best response to player 1's strategy to continue. If player 1 stops, it is of little concern what player 2 chooses: either strategy is a best response. Hence the equilibrium condition for player 2 is:

$$q = \frac{1}{2}p, q \in [0, 1].$$

The equilibrium condition for player 1 is either $p = q$ or $p = \frac{1}{2}q$ or $p = 0$, depending on whether he finds a better than, equivalent to, or worse than outcome b. In the first two cases, i.e., if $a \succeq_1 b$, there is a Nash equilibrium yielding outcome a which is never played in a regret equilibrium. No matter what preferences player 1 has over a and b, the unique regret equilibrium in all cases is that both players decide to continue with probability one.

Example 6.5 In the T-choice centipede game, introduced by Rosenthal (1981), players 1 and 2 alternately move. In any of the $2T$ periods, the player whose turn it is to move can decide to stop the game (S) or to continue (C). Consequently, both players have $T+1$ pure strategies: stopping at any one of the T opportunities, or continue all the time. The game ends if one of the players decides to stop or if neither player has decided to do so after each of them has had T opportunities. For each player, the outcome when he stops the game in period t is better than that in which the other player stops the game in period $t+1$ (or the game ends), but worse than any outcome that is reached if in period $t+1$ the other player passes the move to him. In terms of regret:
Player 2 feels no regret of stopping at his k-th opportunity exactly in the following cases:

- *player 1 stops immediately; then all of player 2's $T+1$ pure strategies are a pure best response;*
- *$k = T$; the unique best response to player 1's choice to continue always is to stop at the final stage;*
- *if player 1 decides to stop at opportunity $k+1$.*

Player 1 feels no regret of stopping at his k-th opportunity in exactly one case:

- *if player 2 decides to stop in the next period, at his k-th opportunity.*

An example of a 3-choice centipede game is given below.

Denote by $p_i(q_i)$ the probability of player 1 (2) to stop at his i-th opportunity, once this opportunity is reached ($i = 1, \ldots, T$). Thus, our computations are in behavioral, rather than in mixed strategies. We show that for each number $T \in \mathbb{N}$ of choices and each $k \in \{0, \ldots, T-1\}$:

$$p_{T-k} = q_{T-k} = \frac{2}{k+3}. \tag{3}$$

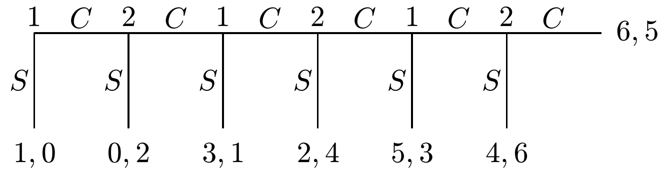

Figure 10: A 3-choice centipede game

In particular, if the number of choices T approaches infinity, the probability for each player to stop at the first (and by the same argument at any finite) opportunity, converges to zero. This provides a potential resolution of the paradoxical aspect of the centipede game. In the unique Nash equilibrium both players stop immediately with probability one, in the regret equilibrium concept they stop immediately with positive probability, but the solution in (3) indicates that there is a strong urge to continue playing, thus providing the possibility to achieve more preferable outcomes.

The description of the T-choice centipede game in terms of regret (emphasized above) immediately gives rise to the following conditions for player 1:

$$p_1 = q_1 \tag{I.1}$$

$$(1-p_1)p_2 = (1-q_1)q_2 \tag{I.2}$$

$$\cdots$$

$$(1-p_1)(1-p_2)\cdots(1-p_{T-1})p_T = (1-q_1)(1-q_2)\cdots(1-q_{T-1})q_T, \tag{I.T}$$

and for player 2:

$$q_1 = \frac{p_1}{T+1} + (1-p_1)p_2 \tag{II.1}$$

$$(1-q_1)q_2 = \frac{p_1}{T+1} + (1-p_1)(1-p_2)p_3 \tag{II.2}$$

$$\cdots$$

$$(1-q_1)(1-q_2)\cdots(1-q_{T-1})q_T = \frac{p_1}{T+1} + (1-p_1)(1-p_2)\cdots(1-p_T). \tag{II.T}$$

The conditions that arise from always continuing are redundant (see Remark 3.3).

We prove first of all, that in the T-choice centipede game we have for each $i = 1, \ldots, T-1$:

$$p_i = q_i, p_i \notin \{0, 1\}. \tag{4}$$

This is necessary to avoid division by zero when we solve the game. We know from condition (I.1) that $p_1 = q_1$. Suppose $p_1 = 1$. Substitution in (II.1) yields $1 = \frac{1}{T+1}$, a contradiction. Suppose $p_1 = 0$. Then $p_2 = q_2$ by (I.2) and $p_2 = 0$ by (II.1). Hence $p_3 = q_3$ by (I.3) and $p_3 = 0$ by (II.2).

Proceeding in this fashion yields that $p_k = q_k = 0$ for all $k = 1, \ldots, T$, which contradicts (II.T). Hence $p_1 = q_1, p_1 \notin \{0, 1\}$. Now assume that we have shown for some $k \in \{1, \ldots, T-2\}$

$$\forall n \leq k : p_n = q_n, p_n \notin \{0, 1\}.$$

We proceed to show that the same holds for $k+1$. First of all, we have from (I.k+1) that $p_{k+1} = q_{k+1}$. Consider condition (II.k+1):

$$(1-q_1)(1-q_2)\cdots(1-q_k)q_{k+1} = \frac{p_1}{T+1} + \underbrace{(1-p_1)(1-p_2)\cdots(1-p_{k+1})p_{k+2}}_{\geq 0}.$$

If $q_{k+1} = 0$, then its left hand side equals zero, which would imply that $p_1 \leq 0$, whereas we know from the above that $p_1 > 0$. If $p_{k+1} = 1$, condition (II.k+2) reduces to

$$0 = \frac{p_1}{T+1},$$

a contradiction. This finishes the proof. This part was necessary to avoid division by zero in the following solution of the game.

Substitute the left-hand side of player 1's conditions in the left-hand side of player 2's conditions. This yields

$$\begin{aligned}
p_1 &= \frac{p_1}{T+1} + (1-p_1)p_2 \\
(1-p_1)p_2 &= \frac{p_1}{T+1} + (1-p_1)(1-p_2)p_3 \\
&\cdots \\
(1-p_1)(1-p_2)\cdots(1-p_{T-1})p_T &= \frac{p_1}{T+1} + (1-p_1)(1-p_2)\cdots(1-p_T)
\end{aligned}$$

Obviously, the first equation is equivalent to

$$\frac{T}{T+1}p_1 = (1-p_1)p_2.$$

Using this equality to replace the left-hand side of the second equation leads to

$$\frac{T}{T+1}p_1 = \frac{p_1}{T+1} + (1-p_1)(1-p_2)p_3,$$

which is equivalent to

$$\frac{T-1}{T+1}p_1 = (1-p_1)(1-p_2)p_3.$$

Use this equation, again, in order to replace the left-hand side of the third equation. This leads to

$$\frac{T-2}{T+1}p_1 = (1-p_1)(1-p_2)(1-p_3)p_4.$$

Continuing in this way, we get the following equivalent system of T equations:

$$\begin{aligned}
\frac{T}{T+1}p_1 &= (1-p_1)p_2 \\
\frac{T-1}{T+1}p_1 &= (1-p_1)(1-p_2)p_3 \\
&\cdots \\
\frac{2}{T+1}p_1 &= (1-p_1)(1-p_2)\cdots(1-p_{T-1})p_T \\
\frac{1}{T+1}p_1 &= (1-p_1)(1-p_2)\cdots(1-p_{T-1})(1-p_T).
\end{aligned}$$

The final step is to roll it up backwards again. Add the last and the second last equation to get

$$\frac{3}{T+1}p_1 = (1-p_1)(1-p_2)\cdots(1-p_{T-1}). \tag{5}$$

In combination with the second last equation and (4), which assures that we do not divide by zero, this immediately leads to

$$p_T = \frac{2}{3}.$$

Now start with equation (5) and first, add the third last equation, and second, divide in a similar way. This yields first,

$$\frac{6}{T+1}p_1 = (1-p_1)(1-p_2)\cdots(1-p_{T-2}), \tag{6}$$

and second,

$$p_{T-1} = \frac{3}{6} = \frac{1}{2}.$$

Now do this again with equation (6) in combination with the fourth last equation and get

$$\frac{10}{T+1}p_1 = (1-p_1)(1-p_2)\cdots(1-p_{T-3}),$$

and so

$$p_{T-2} = \frac{4}{10} = \frac{2}{5}.$$

This procedure stops when reaching the first equation, thereby generating the following sequence of probabilities:

$$\frac{2}{3}, \frac{1}{2}, \frac{2}{5}, \frac{1}{3}, \frac{2}{7}, \frac{1}{4}, \ldots$$

It is easy to see that

$$\forall T \in \mathbb{N}, \forall k \in \{0, \ldots, T-1\} : p_{T-k} = q_{T-k} = \frac{2}{k+3}.$$

In particular, if the number of choices T approaches infinity, the probability for each player to stop at the first (and by the same argument at any finite) opportunity, converges to zero. Osborne and Rubinstein (1997), in computing their equilibrium notion in the centipede game, conclude that their equilibrium notion makes sense only if both players fail to understand the structure of the game. In our equilibrium notion, the equilibrium conditions form an almost immediate translation of the structure of the game, where it is a unique best response to stop exactly one period ahead of your opponent's intent to do so. Still, we find a potential resolution of the paradox posed by the centipede game: The players play the unique Nash equilibrium of stopping immediately with positive probability, but there is a strong urge to continue playing.

7 Conclusion

In this paper we have investigated a solution concept based on a simple procedure. Each player plays optimally against a profile of pure strategies of the opponents that is randomly drawn from the distribution over the set of their pure strategies that is induced by his beliefs about the mixed strategies of his opponents. The notion of regret equilibrium follows by imposing a consistency requirement. Most of the analysis was done for a very general class of games. Only very simple types of comparisons are required of the players, since a player does not have to compare lotteries or even outcomes that arise facing different strategies of the opponents. It is sufficient that for each profile of pure strategies of the opponent, a player is able to decide whether or not his strategies are pure best responses. This makes the concept widely applicable.

References

Bernheim B.D. (1984): "Rationalizable Strategic Behavior", *Econometrica*, 52, 1007-1028.

Chen H-C., Friedman J.W., and Thisse J-F. (1997): "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach", *Games and Economic Behavior*, 18, 32-54.

Harsanyi J.C. (1973): "Games with Randomly Distributed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points", *International Journal of Game Theory*, 2, 1-23.

Hart S. and Mas-Colell A. (1997): "A Simple Adaptive Procedure leading to Correlated Equilibrium", DP 126, Hebrew University, Jerusalem.

Jansen M.J.M. (1981): "Maximal Nash Subsets for Bimatrix Games", *Naval Research Logistics Quarterly*, 28, 147-152.

Kahneman D. and Tversky A. (1979): "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, 47, 263-291.

Loomes G. and Sudgen R. (1982): "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty", *Economic Journal*, 92, 805-824.

Osborne M. and Rubinstein A. (1997): "Games with Procedurally Rational Players", to appear in *American Economic Review*.

Pearce D.G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection", *Econometrica*, 52, 1029-1050.

Rosenthal R.W. (1981): "Games of Perfect Information, Predatory Pricing, and the Chain-Store Paradox", *Journal of Economic Theory*, 25, 92-100.

Rubinstein A. (1997): *Modeling Bounded Rationality*, Cambridge MA: MIT Press.

Simon H.A. (1976): "From Substantive to Procedural Rationality", in *Method and Appraisal in Economics*, S.J. Latsis (ed.), Cambridge: Cambridge University Press, pp. 129-148.

Winkels H.M. (1979): "An Algorithm to Determine all Equilibrium Points of a Bimatrix Game", in *Game Theory and Related Topics*, O. Moeschlin and D. Pallaschke (eds.), Amsterdam: North Holland, pp. 137-148.