Economics Letters 36 (1991) 141–146
North-Holland

Strategy subsets closed under rational behavior *

Kaushik Basu

Delhi School of Economics, Delhi 110007, India
Princeton University, Princeton, NJ 08544, USA

Jörgen W. Weibull

Stockholm University, S-10691 Stockholm, Sweden
Princeton University, Princeton, NJ 08544, USA

Received 19 October 1990
Accepted 29 January 1991

A set of strategy profiles is here said to be closed under rational behavior (*curb*) if it contains all its best replies. Each *curb* set contains the support of at least one Nash equilibrium in mixed strategies, but there are perfect Nash equilibria that are not contained in any minimal *cuirb* set. It is shown that every game with compact strategy sets and continuous payoff functions possesses at least one minimal *curb* set, that every minimal *curb* set is identical with its best replies and that it is contained in the set of rationalizable strategy profiles.

1. Introduction

While a Nash equilibrium is a point in the Cartesian product of the players' strategy spaces such that no player can increase his payoff by a unilateral deviation, a strict Nash equilibrium requires that any unilateral deviation actually incurs a loss. Any strict Nash equilibrium has all the strategic stability properties that the refinement literature asks for, but many games lack such equilibria. In any non-strict Nash equilibrium, at least one player is indifferent between some of his pure strategies even under his Nash equilibrium beliefs. As is well known, such indifference can make the equilibrium highly 'unstable'.

The point is illustrated in fig. 1 below. The game in diagram (a) has a unique Nash equilibrium, and in this equilibrium both players randomize between their first two strategies, player 1 choosing T with probability $2/3$ and M with probability $1/3$, and player 2 choosing L with probability $1/4$. However, under these Nash equilibrium beliefs, player 1 is indifferent between T and M, and player 2 is indifferent between L and R. If player 1 would assign probability $p > 2/3$ to the uncertain event that 2 will choose L, then 1's unique optimal strategy is B. And if player 2 would assign probability $q > 1.3$ to the event that 1 will choose M, then 2's unique strategy is L etc. Strategy T, which is assigned probability $2/3$ in Nash equilibrium, is optimal for player 1 only if he assesses

* We are grateful for helpful comments from Dilip Abreu and Brian Kanaga. The research of Weibull was supported by the Swedish Council for Research in the Humanities and Social Sciences (HSFR).

	L	R
T	4,1	1,2
M	1,3	2,1
B	6,2	0,3

(a)

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

(b)

Fig. 1

$1/4 < p < 1.3$. It thus appears that the unique Nash equilibrium of this game (hence its only perfect equilibrium) is 'unstable'. In particular, it does not seem justified to exclude, as does the Nash equilibrium criterion, the possibility that player 1 will play B, even if pre-play communication is presumed.

While the example in fig. 1a suggests a coarsening of the Nash equilibrium concept, the same type of argument applied to the game in fig. 1b suggests a refinement of the Nash equilibrium requirement. For although the Nash equilibrium (M, R) is undominated and hence perfect, a pre-play agreement to play (M, R) does not appear to be 'self-enforcing', since also L is a best reply to M, and in view of this indifference on behalf of player 2, player 1 may contemplate playing T. If player 2 assigns positive probability to the possibility that 1 will play T, and a smaller probability to the event that 1 will play B, then 2's unique optimal strategy is L, in which case player 1 should certainly play T instead of M. It hence seems that we are led to a rejection of the perfect but non-strict Nash equilibrium (M, R) in favor of the strict equilibrium (T, L).

We will argue that even if one agrees to treat pre-play communication as possible but implicit, it is not clear why one should presume that 'agreements' take the form of a single strategy (pure or mixed) profile, rather than a set of strategy profiles. Indeed, a set-valued solution concept has been developed by Kohlberg and Mertens (1986). However, while they select sets of Nash equilibria, i.e. sets contained in their own best replies, we here select sets containing all their own best replies, a 'dual' approach which can be viewed as a set-theoretic coarsening of the notion of strict Nash equilibrium while their approach is a set-theoretic refinement of the Nash equilibrium concept.

One set-valued solution concept developed in this paper is the 'tight curb' notion, i.e. sets which are identical with their own best replies. In the light of the work of Bernheim (1984) and Pearce (1984) one sees that, in a game with continuous payoff functions and compact strategy sets, a maximal tight curb set coincides with the set of rationalizable strategy profiles. Hence, while their work may be seen as an investigation into the properties of the maximal tight curb set in a game, the present paper may be viewed as an exploration of the whole spectrum of tight curb sets, with particular attention paid to the opposite end of this spectrum, viz. the minimal tight curb sets. One reason for highlighting minimal tight curb sets is that these are the 'nearest' set-valued generalization of strict equilibria. Moreover, minimality has the advantage of reducing strategic ambiguity. However, in certain games such an advantage turns out to be too costly in terms of other game theoretic desiderate, so in this first exploration we consider minimal tight curb sets only as useful benchmarks. In addition to studying curb and tight curb sets, defined in terms of best replies, we also investigate a weaker notion, *curb** sets, based on undominated best replies.

2. Notation and preliminaries

Attention in this paper is focused on normal-form games $G = (N, S, U)$, where $N = \{1, 2, \ldots, n\}$ is the set of players, S is the Cartesian product of the players' strategy sets S_i, and U is a mapping of

S into $\mathbb{R}^n$, such that $U_i(s) \in \mathbb{R}$ is the i'th player's von Neumann–Morgenstern utility level when strategy profile s is played. We will let $\mathscr{G}$ denote the class of all games G in which each strategy set S_i is a compact set in some Euclidean space and each payoff function U_i: $S \to \mathbb{R}$ is continuous. For any game G, let P be the collection of all products of non-empty and compact subsets of the players strategy sets, i.e. $X \in P$ if and only if X is the Cartesian product of nonempty compact sets $X_i \subset S_i$ $[i = ,1, 2, \ldots, n]$. (In particular, $S \in P$ if $G \in \mathscr{G}$.) As is usual, we will write $s \backslash s_i'$ when the ith component of a profile s is replaced by s_i'. A strategy profile $s \in S$ is a Nash equilibrium if $U_i(s) \geqslant U_i(s \backslash s_i') \forall i \in \mathbb{N}$, $\forall s_i' \in S_i$, and it is a strict (Nash) equilibrium if $U_i(s) > U_i(s \backslash s_i')$ $\forall i \in \mathbb{N}$, $\forall s_i' \in S_i\text{-}\{s_i\}$.

In the present paper, a player's belief about others' strategies takes the form of a product probability measure on all players' strategy sets (we include his own strategy set in the domain only for notational convenience). Hence, a player's belief is formally identical with a mixed strategy profile $m \in M(S)$, where $M(S)$ is the Cartesian product of the sets $M(S_i)$ of Borel probability measures over each strategy set S_i. For any Borel subset $X_i \subset S_i$, let $M(X_i)$ be the (Borel) probability measures with support in X_i, i.e. $M(X_i) = \{m \in M(S_i)\colon m(X_i) = 1\}$, and for any $X \in P$ let $M(X)$ be the Cartesian product of the sets $M(X_i)$.

For each player $i \in \mathbb{N}$, streategy $s_i \in S_i$ and belief $m \in M(s)$, let $u_i(s_i \mid m)$ be the player's expected utility under belief m when he plays s_i [1]. Let $\beta_i(m)$ be the i'th player's set of optimal strategies in S_i under belief $m \in M(S)$, i.e., strategies $s_i \in S_i$ such that $u_i(s_i \mid m) > u_i(s_i' \mid m)$ $\forall s_i' \in S_i$. If $G \in \mathscr{G}$, that is, payoff functions are continuous and strategy sets compact, then each set $\beta_i(m) \subset S_i$ is non-empty and compact. For any set $X \in P$, let $\beta_i(X)$ denote the i'th player's set of optimal strategies under beliefs in $M(X)$:

$$\beta_i(X) = \bigcup_{m \in M(X)} \beta_i(m), \tag{1}$$

and write $\beta(X)$ for the Cartesian product of the sets $\beta_i(X)$. This study is restricted to games G in $\mathscr{G}$, a class of games for which $X \in P$ implies $\beta(X) \in P$.

3. Sets closed under rational behavior

A set X of strategy profiles will be said to be *closed under rational behavior* (*curb*) if $X \in P$ and $\beta(X) \subset X$. In words: a set X in P is curb if the belief that strategies outside X will not be played implies that such strategies will indeed not be played by players who are rational in the sense of never playing strategies that are suboptimal.

In spirit, the present criterion is related to the notion of strict equilibrium, and, indeed, every such equilibrium, viewed as a singleton set, meets the curb condition: if $s \in S$ is a strict equilibrium then $\{s^*\} \subset \beta(\{s^*\})$. However, the curb criterion is also met by the set $X = S$ of all strategy profiles in the game, the set $X = S$ thus being the maximal curb set. Conversely, one may ask whether there exist *minimal* curb sets, i.e., curb sets which do not contain any proper subset which is a curb set:

[1] More exactly, we define $u_i\colon S_i \times M(X) \to \mathbb{R}$ by $u_i(s_i \mid m) = \int U_i(z \backslash s_i) \mathrm{d}m(z)$. Note that $u_i(s_i \mid m)$ is functionally independent of the component $m_i \in M(S_i)$.

Proposition 1. *Every game $G \in \mathscr{G}$ possesses at least one minimal curb set.*

Proof. Let Q be the (non-empty) collection of curb sets in S, partially ordered by (weak) set inclusion. By Hausdorff's Maximality Principle, Q contains a maximal nested sub-collection. Let $Q' \subset Q$ be such a sub-collection, and, for each $i \in \mathrm{N}$, let $\tilde{X}_i$ be the intersection of all sets X_i' for which $X' \in Q'$. Since each set X_i' is non-empty and compact, so is $\tilde{X}_i$, by the Cantor Intersection Theorem. Hence, $X \in P$. Suppose $s_i \in \beta_i(\tilde{X})$. Since $M(\tilde{X}) \subset M(X') \forall X' \in Q'$, we have $s_i \in \beta_i(X') \forall X' \in Q'$, and thus $s_i \in X_i' \forall X' \in Q'$ (since all $X' \in Q'$ are curb). Hence, $s_i \in \tilde{X}_i$, so $\beta_i(\tilde{X}) \subset \tilde{X}_i \forall i \in \mathrm{N}$, i.e., $\tilde{X}$ is curb. □

Generalizing the definition of strict equilibrium from singleton sets to arbitrary product sets, we call a curb set X tight if $\beta(X) = X$ [2]. In particular, a profile $s \in S$ is a strict equilibrium if and only if $\{s\}$ is a tight curb set. Note that a tight curb set is 'immune' to iterated elimination of suboptimal strategies under beliefs in $M(X)$. For if it is common knowledge that no player will use a strategy outside X, then each player knows that other (rational) players will play in $\beta(X)$ and hence each player should play in $\beta(\beta(X))$ etc. If X is tight curb, then such iteration has no effect: $\beta^n(X) = X$, for all n.

While many games lack strict equilibria, every game G in $\mathscr{G}$ possesses at least one tight curb set. In fact, in such games, every minimal curb set is tight, indeed a minimal tight curb set (i.e., it contains no proper subset which is a tight curb set). Conversely, every minimal tight curb set is a minimal curb set. Formally:

Proposition 2. *A set in a game $G \in \mathscr{G}$ is a miminal curb set iff it is a minimal tight curb set.*

Proof. First, suppose X is minimal curb but not tight. Then there exists some player $j \in \mathrm{N}$ for which $\beta_j(X) \subset X_j$ and $\beta_j(X) \neq X_j$. Let $X_j' = \beta_j(X)$ and $X_i' = X_i \forall_i \neq j$. Then $X' \subset X$, $X' \neq X$, $\beta_j(X') = X_j'$ and $M(X') \subset M(X)$, so $\beta_i(X') \subset \beta_i(X) \subset X_i' = X_i' \forall i \neq j$. Hence, $\beta(X') \subset X'$. The payoff function u_j being continuous, the correspondence β_j from $M(X)$ to S_j is non-empty – and compact – valued, and, by Berge's Maximum Theorem upper hemi-continuous [3]. Moreover, $M(X)$ is compact, so its image $\beta_j(X) = \beta_j(X') = X_j'$ is non-empty and compact, i.e., $X' \in P$. In sum: X' is curb, contradicting the hypothesis that X is minimal. Thus, any minimal curb set is tight, and, being minimal among curb sets, it is minimal among tight curb sets. Secondly, suppose X is a minimal tight curb set. Applying the proof of Proposition 1 to the curb set X in the role of S, one establishes the existence of a minimal curb set $X' \subset X$. By the first part of the present proof, such a set X' is tight, and, since by hypothesis X is a minimal tight curb set, $X' = X$. Hence, X is a minimal curb set. □

In order to relate the concept of a minimal curb set to established solution concepts, two further observations are useful, both being valid for all games G in $\mathscr{G}$. First, every curb set in such a game contains the support of at least one Nash equilibrium in mixed strategies. To see this, suppose $X \subset S$ is a curb set in $G = (N, S, U)$ and consider the game $G' = (N, X, U)$ obtained when players are restricted to the (non-empty and compact) strategy subsets X_i. Like G, the 'subgame' G' meets the

[2] In the language of Pearce (1984), a set $X \in P$ has the 'best response property' if $X \subset \beta(X)$, so a curb set X is tight iff it has the best response property.

[3] We endow $M(S)$ with the *topology of weak convergence*. A sequence (m^t) from $M(S)$ is said to *converge weakly* to $m \in M(S)$ if $\int f \mathrm{d}m^t \to \int f \mathrm{d}m$ for all continuous functions $f: \mathrm{S} \to \mathbb{R}$. Since S is a compact metric space, so is $M(S)$, see e.g. Theorem 6.4 in Parthasarathy (1967).

	L	R
T	3,1	0,0
B	0,0	1,3

(a)

	L	M	R
T	1,1	0,0	0,0
M	0,0	2,3	3,2
B	0,0	3,2	2,3

(b)

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

(c)

Fig. 2.

conditions of the Glicksberg Theorem concerning the existence of Nash equilibrium in mixed strategies. Now, if $m \in M(X)$ is such an equilibrium of G', then it is also a Nash equilibrium of G, since by hypothesis each restricted strategy set X_i contains all best replies in S_i to strategies in X. Secondly, the set $R \subset S$ of rationalize strategy profiles is non-empty and is the largest product set $X \subset S$ satisfying the equality $X = \beta(X)$. Hence, if the set R is compact, then $R \in P$ and R is the maximal tight curb set. Indeed, one can prove that R is compact (and non-empty) in games $G \in \mathscr{G}$ [Basu and Weibull (1990)]. In the light of this observation and Proposition 2, it should be clear why a minimal curb set and the set of rationalizable strategy profiles can be thought of as the two ends of a spectrum.

To illustrate these general findings, let us briefly return to fig. 1. One notes that the only curb set in 1.a is the full strategy space $\{T, M, B\} \times \{L, R\}$ *itself, and it contains as a proper subset the support* $(\{T, M\} \times \{L, R\}$ itself, and it contains as a proper subset the support $(\{T, M\} \times \{L, R\})$ of the unique (mixed strategy) Nash equilibrium, which also happens to be quasi-strict. The only minimal curb set in 1b is $\{T\} \times \{L\}$, the support of the unique strict equilibrium. Note that the support of the undominated and hence perfect Nash equilibrium (M, R) is not contained in any minimal curb set.

Figure 2a below shows the 'battle of the sexes' with three curb sets, $\{T\} \times \{L\}$, $\{B\} \times \{R\}$ and $\{T, B\} \times \{L, R\}$, all of which are tight, but only two of which are minimal. The non-minimal curb set appears to be a plausible pre-play 'agreement', viz. if the two players cannot agree on any of the two strict equilibria. The game in 2b has two minimal curb sets, $\{T\} \times \{L\}$ and $\{M, B\} \times \{M, R\}$, the first containing a strict equilibrium and the second a non-strict equilibrium. Intuitively, the second set seems the more likely pre-play agreement. The game in 2c suggests that the curb requirement may in some games be too restrictive. Consider the set $\{T\} \times \{L\}$ in that game. This set is evidently not curb, since player 1 can costlessly deviate from T. Yet one could argue that it is 'closed' if not under 'rational' play, at least under 'rational and cautious' play, since the only other best reply for player 1 to 2's strategy L is his weakly dominated strategy B.

4. Sets closed under rational and cautious behavior

The above observation about optimal but weakly dominated strategies suggests the following weakening of the curb criterion. For each player $i \in$ N, let $S_i^* \subset S_i$ be his subset of strategies that are not weakly dominated, i.e., for which there exists no mixed strategy $m_i \in M(S_i)$ which weakly dominates s_i [4]. For any belief $m \in M(S)$, let $\beta_i^*(m) = \beta_i(m) \cap S_i^*$, the i'th player's undominated optimal strategies under m, and for any $X \in P$, let $\beta_i^*(X) = \beta_i(X) \cap S_i^*$. A set $X \subset S$ will be called

[4] More exactly, but with a slight abuse of notation: there exists no $m_i \in M(S_i)$ such that $U_i(s' \backslash m_i) \geqslant U_i(s' \backslash s_i)$ for all $s' \in S$, with strict inequality for at least one $s' \in S$.

closed under rational and cautious behavior (curb*) if $X \in P$ and $\beta^*(X) \subset X$. In words: A set X of (pure) strategies is curb* if it is a non-empty and compact product set, and if the belief that strategies outside X will not be played implies that such strategies will indeed not be played by any player who is (a) rational in the sense of never playing strategies that are suboptimal, and (b) cautious in the sense of never playing a weakly dominated strategy [5]. This weaker criterion is evidently met by the set $X = \{T\} \times \{L\}$ in fig. 2.c.

It is not difficult to verify that minimal *curb** sets exist in all games with continuous payoff functions and compact strategy sets. The proof of Proposition 2 does not apply, though. For while (in games $G \in \mathcal{G}$) the set $\beta(X)$ is non-empty and compact if X is, the set $\beta^*(X)$ need not be compact. However, one can show that it is non-empty for every $X \in P$ [6]. Hence, if we use the slightly weaker tightness condition that X be contained in the closure of $\beta^*(X)$ – such an X may be called almost tight – then a minor elaboration of the proof of Proposition 2 applies to *curb**, mutatis mutandis, leading to the following parallel conclusion: a set is a minimal *curb** set iff it is a minimal almost tight *curb** set [7].

References

Basu K. and J.W. Weibull, 1990, Strategy subsets closed under rational behavior, Discussion paper no. 62 (Princeton University, Princeton, NJ).

Bernheim D., 1984, Rationalizable strategic behavior, Econometrica 52, 1007–1029.

Kohlberg E. and J.-F. Mertens, 1986, On the strategic stability of equilibria, Econometrica 54, 1003–1037.

Parthasarathy, K., 1967, Probability measures on metric spaces (Academic Press, New York).

Pearce D., 1984, Rationalizable strategic behavior and the problem of perfection, Econometrica 52, 1029–1050.

Weibull J.W., 1990, Self-enforcement in extensive-form games, Seminar paper no. 470 (Institute for International Economic Studies, Stockholm University, Stockholm).

[5] As argued in Weibull (1990), play of weakly dominated strategies can actually be both 'rational' and 'cautious' in games in which some player moves more than once in some play of the game. Hence, we take the present criterion to be generally valid only for games in which each player moves at most once in every play of the game.

[6] For a proof of this claim, and for an example of a non-compact set $\beta^*(X)$, see Basu and Weibull (1990).

[7] More precisely, we define a set $X \subset S$ to be an almost tight curb* set if $X \in P$ and $X = \overline{\beta^*(X)}$. Note that if $X \in P$ and $\beta^*(X) \subset X$, then $\overline{\beta^*(X)} \subset X$.