

MULTISTAGE SITUATIONS
ECONOMETRICA (1996) 64, 1415–1437

JOSEPH GREENBERG*, DOV MONDERER**, AND BENYAMIN SHITOVITZ***

* Department of Economics, McGill University,
and C.R.D.E., Université de Montréal.

** Faculty of Industrial Engineering and Management, The Technion, Israel.

*** Department of Economics, Haifa University.

We would like to thank Geir B. Asheim, Martin Hellwig, Steven Tadelis, Licun Xue, and the anonymous referees for very helpful comments. Financial support from the Research Council of Canada [both the Natural Sciences and Engineering (NSERC), and the Social Sciences and Humanities (SCHRC)], from Quebec's Fonds (FCAR), from the Fund for the Promotion of Research at the Technion, and from the Joint Technion–Haifa University Research Fund is gratefully acknowledged.

Typeset by $\mathcal{A}\mathcal{M}\mathcal{S}$ -TEX

MULTISTAGE SITUATIONS.

Joseph Greenberg*, Dov Monderer**, and Benyamin Shitovitz***

* Department of Economics, McGill University, 855 Sherbrooke st.w., Montreal, Quebec H3A2T7, Canada.

** Faculty of Industrial Engineering and Management, The Technion, Haifa, Israel.

*** Department of Economics, Haifa University. Haifa, Israel.

Abstract. We introduce and analyze “multistage situations”, which generalize “multistage games” (which, in turn, generalize “repeated games”). One reason for this generalization is to avoid the perhaps unrealistic constraint - inherent to noncooperative games - that the set of strategy tuples must be a Cartesian product of the strategy sets of the players. Another reason is that in most economic and social activities (e.g. in sequential bargaining without a rigid protocol) the “rules of the game” are rather amorphous; the procedures are rarely pinned down. Such social environments can, however, be represented as multistage situations and be effectively analyzed through the theory of social situations.

The paper contributes to the theory of social situations within the framework of multistage situations (e.g., the existence of a largest conservative stable standard of behavior which yields a definition that extends “subgame perfect equilibrium paths”), and to the theory of multistage games (e.g., the existence of ϵ -generalized perfect equilibrium, for all $\epsilon > 0$). It also provides an equivalence theorem between subgame perfection and the largest conservative stable standard of behavior for multistage games. The usefulness of our approach is further illustrated by our notion of “ k -rationality” whereby (at each subsituation) players look only k steps ahead.

Keywords: Multistage games, Multistage situations, Social situations, Incomplete specifications.

1. Introduction¹

Multistage games have been extensively studied by, among others, Gillette (1957), Fudenberg and Levine (1983,1986), Harris (1985 a,b), Hellwig and Leininger (1987), Hellwig, Leininger, Reny, and Robson (1990), Börgers (1989,1991), and Harris, Reny, and Robson (1995). A multistage game (with observed actions), like an extensive form game, is defined by a tree but unlike the latter, several players may act at the same decision node. It turns out that the description and analysis of such games may be considerably simplified if they are represented equivalently by paths. Specifically a multistage game is a set of paths satisfying two main properties:

- (1) *The continuation property*: If all finite truncations of a path belong to the set, so does the path.
- (2) *The Cartesian product property*: At each node there is a one to one correspondence between a Cartesian product of action spaces and the nodes that can be directly reached from this node.

Our paper is motivated by the observation that many real life situations fail to satisfy one (or both) of these conditions. We shall presently present several examples of a variety of dynamic social environments which support our assertion. Such situations cannot be analyzed by the classical game theoretic solution concepts because, when focusing on subgame perfect equilibrium or its refinements, as we do in this paper, assigning a low payoff to “inadmissible courses of actions” would, in general, extend the set of “punishments” and thereby considerably and unjustifiably enlarge the solution (equilibrium) set.

In this paper we analyze social environments in which the defining set of paths may fail to satisfy either Condition (1) or Condition (2). We call such environments

¹This paper is an extensive revision of: “Conservative stability and subgame perfection” by Greenberg and Shitovitz, D.P. 5/90, McGill University.

“multistage situations” because they are not games but they are social situations as defined in the theory of social situations (TOSS) (Greenberg (1990)).

We analyze multistage situations through the solution concept of a conservative stable standard of behavior (CSSB). In TOSS a standard of behavior is a function that assigns a set of outcomes to each situation; here it is a function that assigns a set of paths to each node (subsituation). The paths that are thus singled out can be thought of as those paths that players believe may be followed when the node in question is reached. A standard of behavior is called stable if (i) no path in the standard admits a “profitable” deviation to a new subsituation (internal stability) and (ii) every path outside the standard does admit a “profitable” deviation to a new subsituation (external stability). This notion of stability is similar in spirit, but not equivalent to that of von Neumann and Morgenstern. It also contrasts with the usual notion of noncooperative equilibrium which focuses on the particular equilibrium without reference to the set of equilibria as a whole.

In considering the “profitability” of a deviation to a new subsituation, agents refer to the prescriptions of the given standard of behavior for this subsituation. Given that the standard of behavior is set-valued, there are many ways of doing so, and hence many stability concepts. A conservative player considers a deviation to be profitable if and only if *all* paths singled out by the standard of behavior for the new subsituation make him better off. So a conservative stable standard of behavior is a stable standard of behavior for conservative players. In contrast, an optimistic stable standard of behavior (OSSB) would be a stable standard for players who consider a deviation to be profitable if and only if *some* path prescribed for the new subsituation would make them better off².

²Optimistic stable standards of behavior are discussed (among others) in Asheim 1988, Greenberg 1990, Shitovitz 1994, and Tadelis 1994. They are strongly related to von Neumann and Morgenstern abstract stable sets.

An example of a standard of behavior for multistage games is the mapping that assigns to each node of the game the set of subgame perfect equilibrium paths in the corresponding subtree. Denote this standard of behavior by PEP. Greenberg (1989) proved that for repeated games the PEP is the largest conservative standard of behavior. In this paper we extend this result to multistage situations, and establish conservative stability as a generalization of subgame perfection³.

It should be noted that the notion of CSSB (like the notion of subgame perfection) imposes no “subsituation consistency”⁴ restrictions. In particular, the CSSB may assign different sets of paths to two subsituations that are identical except for the histories that led to them. We now provide the promised examples.

Example 1: *Normative restrictions:* There are social environments in which players may be forced (because of, e.g., legal, historical, social, or ethical considerations) to restrict their actions so that, for example, the resulting outcomes be Pareto optimal. That is, only Pareto optimal outcomes can be considered and similarly, objections to a proposed outcome cannot be based on deviations to non-Pareto outcomes. But the set of Pareto optimal strategy profiles is, in general, a strict subset of the set of all a-priori possible profiles, and, moreover, it cannot be represented as a Cartesian product of the individual strategy sets⁵.

Example 2: *General Competitive Equilibrium:* Debreu (1952) proved the existence of a competitive equilibrium by representing the market economy as a (one–

³Subgame perfection is closely related to the concept of self-generating sets that were introduced by Abreu, Pearce, and Staccetti (1986) (see van Damme 1987, pp. 183–184 for details). This concept was generalized by Bergin and Macleod (1993), and this generalization was further discussed by Asheim (1992). Geir Asheim pointed out to us the potential interesting relationships (that are still to be further explored) between (generalized) self-generating sets and conservative stable standards of behaviors. Note, however, that self-generating sets are defined in the payoff space while we consider the set of paths.

⁴In the sense of Harsanyi and Selten (1988) (see also Guth, Leininger, and Stephan (1991)).

⁵In a previous version of this paper we fully analyzed the multistage situations derived from the repeated Prisoner’s Dilemma game in which only Pareto optimal paths are admissible. The unique CSSB in this situation includes only the two extreme non-cooperative paths.

shot) strategic form game in which a consumer's strategy is a consumption bundle and a strategy of the auctioneer is a vector of prices. However, the bundle chosen by a consumer must belong to his budget set, which is determined by the strategy chosen by the auctioneer, himself a player in that game. Thus, the ("price-feasible") strategy choices cannot be made independently.⁶ It is precisely for this reason that Debreu introduced the notion of an "abstract economy" which extends that of normal form games. The more general framework of multistage situations allows us to further extend the analysis and consider dynamic versions of Debreu's social equilibrium model.

Example 3: *Lack of the continuation property:* Consider Rubinstein's (1982) sequential bargaining model except that payoffs to the infinite paths (of never reaching an agreement) are not specified. That is, the tree that we wish to study consists of all finite paths of Rubinstein's original game. While this tree can readily be analyzed within our framework of multistage situations, it cannot be analyzed within classical game theory, since it is not a game tree. Indeed, if we were to use strategies (rather than paths) then we must include the payoff to the path that results from the strategy profile in which every player always refuses the offer made to him. That is, we must (because of the continuation property of multistage games) be able to assign a payoff to the resulting path. But this path is not in the tree under consideration since the possibility that agreement will never be reached is not anticipated.

Another example where the same phenomenon arises is "the continuous dollar

⁶As one of the referees notes, one could replace Nash equilibrium in the above model with a Stackelberg equilibrium, the auctioneer being the natural (and only) leader. (This destroys, however, the appealing property that the equilibria of the resulting game are precisely the Walrasian equilibria in the market.) But, in general, there may be no "leaders". Such is the case, for example, in a model with price rigidities, where consumers' budget sets are constrained by the rationing scheme, and at the same time rationing can be imposed only on the short side of the market.

auction game”, where two bidders bid for a prize of, say, 1 dollar. Players make alternating bids under the requirement that every bidder who does not chose to drop out has to strictly overbid the other bidder’s bid. Every player who drops out loses his highest bid. While these rules are “clear” in real life, in order to analyze this social environment as a game, “it is therefore necessary to amend the rules of the game by assigning payoffs to the cases in which bidding does not stop after finitely many bids” (Leininger, 1988, p.240). As we have already mentioned, assigning payoffs to infinite bidding paths may unjustifiably enlarge the solution set by adding feasible punishments.

Example 4: *Realistic specifications:* In many (if not most) “real life” situations, payoffs are assigned only to a (rather limited) subset of all potentially possible courses of action. This fact is reflected by expressions such as: “inconceivable”, “not done”, “inappropriate”, “is not supposed to take place”, etc. For example, no legal system covers all possible eventualities, not even in principle⁷. As theorists, especially over the last two decades, we are accustomed to analyzing models where the payoff matrix (or game tree) is complete, i.e., every entry (or path) is assigned a payoff (or a distribution of payoffs in the case of incomplete information). In practice, however, there are circumstances in which players do not entertain the possibility of finding themselves, (possibly because they did not conceive of these, or they were judged to be, prior to deriving the resulting equilibrium, “rare and unlikely”), even though such circumstances are (“theoretically”) a-priori feasible.

To be more concrete, consider, for example, a variant of Rubinstein’s (1982) sequential bargaining model where at each period one and only one of the two players can “speak” - either make an offer or accept an offer that was previously

⁷Such observations motivated the analysis of incomplete contracts (see e.g., Grossman and Hart (1986)).

made to him. These guidelines are, of course, incomplete. Not only do they not specify who can speak when, but they do not even specify what might happen if, say, both players make offers at the same time. This possibility is simply ruled out⁸; once one player speaks - the other must be silent. Clearly, this “civilized negotiation process” cannot be represented as a game, since the set of admissible actions (strategies) is not a Cartesian product. Yet, we suggest, such specifications are often encountered in real life bargaining and they can easily be represented as a multistage situation.

“Coalition formation” where the negotiation process is not rigidly specified is another source of examples to social environments where either Condition (1) or Condition (2) fail to hold.⁹

Example 5: *Infinite number of players:* There is an extensive and important literature that involves an infinite number of players. Such models typically impose some restrictions on the set of admissible strategy tuples. For example, Alesina and Rosenthal (1992) consider a model with a continuum of voters, each of whom must vote for one of two candidates. The candidate who attains the majority of votes is the winner. For this rule to apply, the set of voters who support a candidate must be measurable. Again, as in the previous examples, the set of “measurable strategies” cannot be represented as a Cartesian product of the individual strategy sets. Therefore, Alesina and Rosenthal’s model is a multistage situation rather than a multistage game¹⁰.

⁸It is *not* that, in equilibrium, players will not choose to speak at the same time. Rather, this option is not feasible.

⁹See e.g., Asheim (1988 and 1991) who studied renegotiation proofness, Mariotti (1993) who considered a variant of the “contingent threat situation” to predict the agreements individuals are likely to reach in normal form games, Muto and Nakayama (1992), Muto and Okada (1992a, 1992b) who studied various applications of “open negotiation processes” (Greenberg 1990, Chapter 7), and Xue (1995) who used conservative stability to analyze coalition formation by both partially and fully farsighted players.

¹⁰In fact, the currently employed *definition* of “a nonatomic game in normal form”, due to Schmeidler (1973), imposes a joint measurability constraint on the strategy sets. Hence it is no

Example 6: *Bounded rationality*: As was suggested to us by one of the referees, TOSS can incorporate a form of bounded rationality in which the players “look ahead” only a finite number of steps. This form of bounded rationality cannot be captured using strategies. It can, however, be analyzed within our framework because a CSSB consists of paths. In Section 6 we provide the formal definition of this notion and show that for every positive discount factor the cooperative outcome in the *finitely* repeated Prisoner’s Dilemma game can (almost) be reached. Moreover, increasing the (finite) number of repetitions yields more cooperation.

The contributions of this paper fall into three main categories: Contributions to the theory of social situations within the framework of multistage situations, contributions to the theory of multistage games, and an equivalence theorem between subgame perfection and the largest conservative stable standard of behavior for multistage games.

The paper is organized as follows. In Section 2 we define multistage situations which generalize the notion of multistage games and in Section 3 we define the concept of “ ϵ -conservative stable standard of behavior” (ϵ -CSSB; see Greenberg 1990, Section 2.6), capturing the fact that deviations from the “status quo” are often costly. We then prove the existence of a largest nonempty valued ϵ -CSSB, (a result that holds for every social situation and not only for multistage situations). In Section 4 we study existence, uniqueness, and nonempty valuedness of ϵ -CSSB for multistage situations that have some particular properties, such as: finite horizon, finite outdegree, perfect information, continuity at infinity, and continuity. Section 5 is devoted to an analysis of the implications of the previous results to multistage games. We extend and generalize the principle, established by Greenberg (1989) for infinitely repeated games with discounting, that for “continuous” games the

longer a game in normal form; it is, however, a (one stage) multistage situation.

largest CSSB coincides with the PEP. For general games (that are not necessarily continuous) we have the weaker equivalence between the largest CSSB and the generalized subgame perfect equilibrium paths (GPEP; a concept due to Börgers, 1989). This result generalizes Abreu (1988). Among the other results in this section is a considerable generalization of Harris (1985a): We establish the existence, for every $\epsilon > 0$, of ϵ -GPEP in perfect information multistage games which are bounded from above and continuous at infinity (but not necessarily continuous, e.g., with non-compact action spaces).

In Section 6 we define the notion of “ k -rationality” capturing a new aspect of bounded rationality. This notion demonstrates that TOSS not only extends game theory, but it also opens the way to new solution concepts for games. In particular, “ k -rationality” differs from playing automata with a fixed size (as in Neyman, 1985), as well as from the concept of k -perfection (see e.g., Chakrabarti (1991)). In particular, unlike these notions, “ k -rationality” may yield a *full* cooperation in the finitely repeated Prisoner’s Dilemma game.

To facilitate reading and allow concentration on concepts rather than techniques all proofs are relegated to Section 7 (the Appendix).

2. Multistage games and multistage situations.

We first describe multistage games (with observed actions¹¹). Our representation of multistage games follows Fudenberg and Levine (1983), Harris (1985a,b), and Börgers (1989, 1991). There is a finite¹² set of players, $N = \{1, 2, \dots, n\}$, and an infinite number of stages (periods), $T = \{1, 2, \dots\}$. At each stage $t \in T$ every player $i \in N$ chooses an action $x_t^i \in S_t^i$. These choices are made simultaneously. Denote $S_t = \times_{i \in N} S_t^i$ and $S = \times_{t=1}^{\infty} S_t$. For $x \in S$ and $t \geq 1$ we denote the history

¹¹Thus, we are (almost) assuming pure strategies. In subsequent papers, we intend to extend our analysis to mixed and correlated strategies.

¹²Most of our results do not make use of this assumption.

$(x_1, x_2, \dots, x_{t-1})$ by x^t . (x^1 denotes the empty history.) At period t , players know x^t . The set of actions which are actually available to player i after observing x^t may depend on the observed history, i.e., it may be a strict subset of S_t^i . The set of feasible paths is given by $P \subseteq S$ and is part of the description of the game. For $x \in S$ and a stage $t \in T$ set

$$P(x, t) = \{y \in P : y_r = x_r \text{ for all } 1 \leq r < t\}.$$

Thus, $P(x, 1) = P$ for every $x \in S$. Clearly, for all $x \in P$ and $t \geq 1$, $x \in P(x, t)$.

Multistage games satisfy also the converse implication, that is,

$$(2.1) \quad x \in S \text{ and } P(x, t) \neq \emptyset \text{ for all } t \implies x \in P.$$

For $x \in P$ and a stage $t \in T$ denote by $A_t(x)$ the set of possible actions at x^t .

That is,

$$(2.2) \quad A_t(x) = \{y_t \in S_t : \text{there exists } z \in P(x, t) \text{ such that } z_t = y_t\}.$$

For $i \in N$ let

$$A_t^i(x) = \{y_t^i \in S_t^i : \text{there exists } z_t \in A_t(x) \text{ such that } z_t^i = y_t^i\}.$$

A *multistage game* is a pair (P, u) where $u = (u^i)_{i \in N}$, $u^i : P \rightarrow R$ is player i 's utility (or payoff) function, the set of feasible paths $P \subseteq S$ satisfies (2.1), and the following Cartesian product condition holds:

$$(2.3) \quad A_t(x) = \times_{i \in N} A_t^i(x) \text{ for every } x \in P \text{ and for every } t \geq 1.$$

Both conditions (2.1) and (2.3) are necessary in all classical solution concepts that make use of the notion of strategy. As mentioned in the Introduction, there are, however, many interesting cases in which condition (2.1) and/or condition (2.3) are

most restrictive. The theory of social situations does not impose these conditions. Define, therefore, a *multistage situation* to be a pair (Π, u) where $u = (u^i)_{i \in N}$, $\Pi \subseteq S$ and $u^i : \Pi \rightarrow R$ is the payoff function of player i .

Clearly, every multistage game is also a multistage situation (and therefore every result that holds for multistage situations holds for multistage games). The additional generality enables the analysis of many interesting multistage situations that are not multistage games.

3. Conservative stability.

Let $G = (\Pi, u)$ be a multistage situation. For $x \in \Pi$ and a stage t , $G(x, t)$ denotes the subsituation that results when the path $x \in \Pi$ was offered and followed up to stage t .¹³ A *standard of behavior (SB)* for G is a mapping σ that assigns to every subsituation $G(x, t)$ a (possibly empty) subset of $\Pi(x, t)$, denoted by $\sigma(x, t)$. Paths in $\sigma(x, t)$ are interpreted to be those paths which can be “reasonably” expected to result once the subsituation $G(x, t)$ is reached. Paths in $\Pi(x, t) \setminus \sigma(x, t)$ will not be followed if at least one of the players will choose to deviate. The formalization of this property of the SB σ is captured by conservative stability: For all x and t , $\sigma(x, t)$ consists of all those paths that will not be rejected by any player, when the players are aware of, and believe in the specifications of the mapping σ . Given the mapping σ , a path y will not be accepted by player i if and only if there exists some stage τ in which i will benefit from not following y , no matter which of the paths¹⁴ that σ prescribes after i 's deviation will be followed.

In many cases deviating from the “status quo” is costly. It is for this reason that the above stability notion is extended to ϵ – conservative stability. This extension will also prove to be technically very useful. For $\epsilon \geq 0$, let $\Delta^\epsilon(\sigma, (x, t))$ denote the

¹³Formally, $G(x, t) = (\tilde{\Pi}, w)$ where $\tilde{\Pi} = \Pi(x, t) \times \{(x, t)\}$ and for $y \in \Pi(x, t)$ and $i \in N$, $w^i(y, (x, t)) = u^i(y)$. Clearly, subsituations extend the notion of subgames.

¹⁴Thus, players behave conservatively.

“ ϵ – conservative dominion of σ at (x, t) ”, consisting of all paths in $\Pi(x, t)$ such that there exists at least one player who can unilaterally deviate at some future stage, and every path that is thereafter recommended by σ , makes him at least ϵ better off. That is, a path $y \in \Pi(x, t)$ belongs to $\Delta^\epsilon(\sigma, (x, t))$ if and only if there exist $i \in N$, $\tau \geq t$, and $z \in \Pi(y, \tau)$ with $z_\tau^j = y_\tau^j$ for every $j \neq i$, such that

$$(3.1) \quad \sigma(z, \tau + 1) \neq \emptyset,$$

and

$$(3.2) \quad u^i(\eta) > u^i(y) + \epsilon \quad \text{for all } \eta \in \sigma(z, \tau + 1).$$

Let σ be a standard of behavior for the multistage situation G and let $\epsilon \geq 0$. Following Greenberg (1990, Definition 2.6.2) we say that σ is ϵ – conservative internally stable if for every (x, t) ,

$$\sigma(x, t) \subseteq \Pi(x, t) \setminus \Delta^\epsilon(\sigma, (x, t)).$$

σ is ϵ – conservative externally stable if for every (x, t) ,

$$\sigma(x, t) \supseteq \Pi(x, t) \setminus \Delta^\epsilon(\sigma, (x, t)).$$

σ is ϵ – conservative stable (ϵ –CSSB) if for every (x, t) ,

$$\sigma(x, t) = \Pi(x, t) \setminus \Delta^\epsilon(\sigma, (x, t)).$$

For $\epsilon = 0$, we denote 0–CSSB by CSSB.

Conservative stability is particularly appealing in environments where players are “conservative”, i.e., they exhibit extreme aversion to “Knightian” uncertainty. That is, when the “utility” a player derives from a set of outcomes (over which no probability distribution is given) equals the utility the player derives from the worst

outcome in this set. An SB σ is a CSSB if a path $y \in \Pi(x, t)$ would be followed if and only if no single player would choose to deviate from this path, realizing that the set of paths that would be followed in induced positions are those specified by σ , and all players are behaving “conservatively”. There are other types of stability (that differ from the conservative one) that are discussed in TOSS. In this paper, however, we consider only conservative stability.

The following claim establishes that conservative stable standards of behavior have some desirable properties. First, a conservative stable standard of behavior cannot be identically empty valued.¹⁵ Second, observe that the definition of a SB allows $\sigma(x, t)$ to depend on the entire path x ; in particular it may depend on actions that are prescribed for the future, i.e., on the values x_τ assumes for $\tau \geq t$. Claim 3.3 asserts that for a CSSB σ , $\sigma(x, t)$ depends only on the history $x^t = (x_1, x_2, \dots, x_{t-1})$. Third, a conservative stable standard of behavior satisfies the *truncation property*: if a path is recommended at some stage t then its continuation at any later stage must also be recommended for the (sub)situation reached at that stage.

Claim 3.3. ¹⁶

Let $\epsilon \geq 0$. Let σ be an ϵ -conservative stable standard of behavior for the multi-stage situation $G = (\Pi, u)$. Then

(1) σ is not identically empty valued. That is, there exists $(x, t) \in \Pi \times T$ such that $\sigma(x, t) \neq \emptyset$.

(2) $\sigma(x, t)$ depends only on x^t . That is, $\sigma(y, t) = \sigma(z, t)$ for all $y, z \in \Pi(x, t)$.

(3) σ satisfies the truncation property. That is,

$$y \in \sigma(x, t) \implies y \in \sigma(y, \tau) \text{ for every } \tau \geq t.$$

¹⁵This property holds for every situation. See, Greenberg (1990, Claim 2.5.1).

¹⁶As mentioned in the Introduction, all proofs are relegated to the Appendix.

Another desirable property of conservative stability is that it admits a largest element.¹⁷ Let σ_1 and σ_2 be two standards of behavior. Then $\sigma_2 \subseteq \sigma_1$ if $\sigma_2(x, t) \subseteq \sigma_1(x, t)$ for every (x, t) . The SB σ is *nonempty valued* if $\sigma(x, t) \neq \emptyset$ for every (x, t) .

Theorem 3.4. *Let $\epsilon \geq 0$ and let $G = (\Pi, u)$ be a multistage situation. If the set Σ of all ϵ -conservative internally stable nonempty valued standards of behavior is nonempty, then it admits a largest element with respect to the set inclusion order. That is, there exists $\sigma^L \in \Sigma$ such that $\sigma \subseteq \sigma^L$ for every $\sigma \in \Sigma$. Moreover, σ^L is the largest nonempty valued ϵ -CSSB.*

Let $\epsilon_1 < \epsilon_2$. Note that if σ is ϵ_1 -conservative internally stable then it is ϵ_2 -conservative internally stable, and if σ is ϵ_2 -conservative externally stable then it is ϵ_1 -conservative externally stable. Therefore, ϵ_1 -conservative stability neither implies nor is implied by ϵ_2 -conservative stability. However, for the largest nonempty valued ϵ -CSSB we have the following useful result:

Claim 3.5. *Let $G = (\Pi, u)$ be a multistage situation that admits a nonempty valued ϵ -CSSB for some $\epsilon \geq 0$. Then G admits a nonempty valued ϵ_1 -CSSB for all $\epsilon_1 > \epsilon$. Moreover, $\sigma^{L_1} \subseteq \sigma^{L_2}$, where σ^{L_1} and σ^{L_2} are, respectively, the largest nonempty valued ϵ_1 -CSSB and ϵ_2 -CSSB for G , and $\epsilon \leq \epsilon_1 \leq \epsilon_2$.*

4. Further results on multistage situations.

In this section we analyze multistage situations that have some special properties. Our definitions of these properties extend their analogs in multistage games (see Fudenberg and Levine 1983 and Harris 1985a,b).

A multistage situation $G = (\Pi, u)$ has a *finite horizon* if there exists a stage τ such that $\Pi(x, t) = \{x\}$ for all $x \in \Pi$ and $t \geq \tau$. It has a *finite outdegree* if for all

¹⁷This property holds for every situation.

$i \in N$ and for all (x, t) ,

$$B_t^i(x) = \{z_t^i : z \in \Pi(x, t) \text{ and for all } j \neq i \ z_t^j = x_t^j\}$$

is a finite set. (Recall that “finite action multistage games” require the stronger condition that $A_t(x)$ is a finite set for every $x \in \Pi$ and for every $t \geq 1$.) Player i is *non-active at (x, t)* if i cannot make a unilateral deviation, at stage t , given the history x^{t+1} . That is, if $y \in \Pi(x, t)$ is such that $y_t^j = x_t^j$ for every $j \neq i$, then $y_t^i = x_t^i$ (hence $y_t = x_t$). A multistage situation has *perfect information* if at each (x, t) , all players with the possible exception of one, are non-active. A multistage situation is *bounded (from above)* if for every $i \in N$, u^i is bounded (from above). A multistage situation $G = (\Pi, u)$ is *continuous at infinity* if for every $i \in N$,

$$\lim_{t \rightarrow \infty} \sup_{y, z \in \Pi(x, t)} |u^i(y) - u^i(z)| = 0, \quad \text{uniformly in } x \in \Pi.$$

“Continuity at infinity” allows for non-compact action and paths spaces, as well as for discontinuous and unbounded payoff functions.

A multistage situation $G = (\Pi, u)$ is *continuous* if it satisfies the following conditions: the spaces S_t^i , S_t , and S , where $i \in N$ and $t \in T$, are compact metric spaces and Π is a compact subset of S . Moreover, the payoff functions and the correspondences $x \rightarrow C_t^i(x)$, $x \in \Pi$ are continuous on Π , where for every $x \in \Pi$, $i \in N$, and $t \geq 1$,

$$C_t^i(x) = \{z \in \Pi : z^t = x^t \text{ and for all } j \neq i \ z_t^j = x_t^j\}.$$

Theorem 4.1. *Let $G = (\Pi, u)$ be a continuous at infinity multistage situation. Then, for every $\epsilon > 0$, G admits a unique ϵ -CSSB.*

Since every finite horizon multistage situation is continuous at infinity, Theorem 4.1 yields that every finite horizon multistage situation admits a unique ϵ -CSSB, for every $\epsilon > 0$. By Greenberg (1990, Theorem 5.4.2) this is valid also for $\epsilon = 0$.

The existence theorem above does not exclude CSSBs that assume some empty values (see Example 4.8 below). We now provide conditions that guarantee that every ϵ -CSSB is nonempty valued. This is an important property of a standard of behavior since it ensures that in every subsituation there is at least one path that will be accepted by the players.

Theorem 4.2. *Let $G = (\Pi, u)$ be a bounded from above multistage situation with perfect information and let $\epsilon > 0$. Then every ϵ -CSSB for G is nonempty valued.*

Theorem 4.3. *Let $G = (\Pi, u)$ be a multistage situation with perfect information and finite outdegree, and let $\epsilon \geq 0$. Then every ϵ -CSSB for G is nonempty valued.*

Example 4.8 below demonstrates that boundedness from above is essential for Theorem 4.2 and Example 4.9 below shows that Theorem 4.2 (unlike Theorem 4.3) cannot be generalized for the case where $\epsilon = 0$.

It is easily verified that continuity implies continuity at infinity and boundedness of the payoff functions. Therefore, Theorems 4.1, 4.2 and 4.3 imply

Corollary 4.4. *For all $\epsilon > 0$, every continuous multistage situation with perfect information $G = (\Pi, u)$ admits a unique ϵ -CSSB, σ^ϵ , which is nonempty valued.*

The nonempty valuedness property of the unique ϵ -CSSB in continuous multistage situations (even if they are not of perfect information) has the following two important implications.

Lemma 4.5. *Let G be a continuous multistage situation and let $\epsilon > 0$. If the unique ϵ -CSSB, σ^ϵ , is nonempty valued, then it is compact valued.*

Theorem 4.6. *Let $G = (\Pi, u)$ be a continuous multistage situation. Suppose that for every $\epsilon > 0$, the unique ϵ -CSSB of G , σ^ϵ , is nonempty valued. Then $\sigma = \bigcap_{\epsilon > 0} \sigma^\epsilon$ is the largest nonempty valued CSSB for G .*

Theorems 4.1, 4.2 and 4.6 imply the following existence theorem for continuous multistage situations with perfect information:

Theorem 4.7. *Let $G = (\Pi, u)$ be a continuous multistage situation with perfect information, and for every $\epsilon > 0$ let σ^ϵ be the unique (nonempty valued) ϵ -CSSB of G . Then G admits a nonempty valued CSSB and $\sigma = \bigcap_{\epsilon > 0} \sigma^\epsilon$ is the largest nonempty valued CSSB for G .*

We end this section by providing a few examples that demonstrate that the assumptions we employed to derive our results are not redundant. The first example concerns Theorems 4.2 and 4.3. We show that there are multistage situations (in fact games) that are continuous at infinity whose unique ϵ -CSSB assumes some empty values.

Example 4.8. *Consider the one-person game in extensive form where at the first stage the player chooses between U , in which case he gets a payoff of 1, and D , which brings him to the second stage where he can choose any real number $\alpha \in \mathbb{R}$. If he chooses α , his payoff is α . It is easily verified (see Appendix) that for every $\epsilon > 0$ this game admits a unique ϵ -CSSB, σ , which assigns the empty set to the subsituation representing the second stage.*

If we replace the unbounded choice in the second stage (i.e. the range of α) by the open and bounded interval $(0, 1)$, then the unique CSSB assigns the empty set to the subsituation representing the second stage. Note that by Theorem 4.2, since the utilities here are bounded, for every $\epsilon > 0$, the ϵ -CSSB is nonempty valued.

The next example concerns Theorems 4.3 and 4.7. We show that multistage situations with perfect information and finite outdegree that are not continuous need not admit a CSSB.

Example 4.9. *Consider the one-person game in extensive form where at each*

stage the player has to choose between C - to continue to play or E - to exit. If he exits at stage t , his payoff is $1 - \frac{1}{t}$. If he never exits, his payoff is 0. (See Appendix for a formal proof that this game admits no CSSB.)

5. Applications to subgame perfection in multistage games.

We refer the reader to Fudenberg and Levine (1983), Harris (1985a,b), and Börgers (1989, 1991) for the precise definitions of strategies and related notions in multistage games. We recall now briefly the definitions concerning subgame perfection.

Let $G = (P, u)$ be a multistage game and let $\epsilon \geq 0$. A *perfect ϵ -equilibrium* is a strategy tuple f that satisfies: There exists no subgame in which a player can gain more than ϵ by unilaterally deviating from the restriction of f to this subgame. A *generalized perfect ϵ -equilibrium* is a strategy tuple f that satisfies: There exists no subgame in which a player can gain more than ϵ through a *single-stage* unilateral deviation from the restriction of f to this subgame. This latter definition is due to Börgers (1989). For an n -tuple of strategies f we denote by $p(f)$ the path (which, by (2.1), necessarily belongs to P) defined by f and we denote by $f_{(x,t)}$ the restriction of f to the subgame (x, t) , (i.e., the subgame beginning after the history x^t).

Since multistage situations are represented by paths rather than strategies, in order to apply our results to games we need the following notation. For every (x, t) , let ϵ -PEP(x, t) be the set of all paths $y \in P(x, t)$ for which there exists a strategy tuple f with $p(f) = y$ such that $f_{(x,t)}$ is a perfect ϵ -equilibrium in the subgame (x, t) . Similarly, let ϵ -GPEP(x, t) be the set of all paths $y \in P(x, t)$ for which there exists a strategy tuple f with $p(f) = y$ such that $f_{(x,t)}$ is a generalized perfect ϵ -equilibrium in the subgame (x, t) . Note that the two mappings ϵ -PEP and ϵ -GPEP are standards of behavior, that each is nested in ϵ , and that for a fixed ϵ ,

ϵ -PEP \subseteq ϵ -GPEP. We shall write PEP for 0-PEP, and GPEP for 0-GPEP.¹⁸ In general, GPEP and PEP are distinct SBs. Example 5.8 below is a typical example that demonstrates this fact.

We now state an equivalence theorem between the existence of generalized perfect ϵ -equilibrium paths and the existence of ϵ -conservative stable standards of behavior.

Theorem 5.1. *Let $\epsilon \geq 0$ and let $G = (P, u)$ be a multistage game. Then*

(1) *G admits a generalized perfect ϵ -equilibrium if and only if G admits a nonempty valued ϵ -CSSB.*

(2) *Every nonempty valued ϵ -CSSB σ recommends only generalized perfect ϵ -equilibrium paths. That is, $\sigma \subseteq \epsilon$ -GPEP. Moreover, if G admits a generalized perfect ϵ -equilibrium then ϵ -GPEP is the largest nonempty valued ϵ -CSSB.*

Remark 5.2. *The proof of Theorem 5.1 establishes a stronger result. The theorem remains valid if CSSB is replaced by “conservative internally stable standard of behavior”.*

Consider a multistage game $G = (P, u)$ that admits a subgame perfect equilibrium. Since every subgame perfect equilibrium is also a generalized subgame perfect equilibrium, Theorem 5.1 implies that the GPEP is the largest nonempty valued CSSB. As Example 5.7 below demonstrates, GPEP does not, in general, coincide with the PEP. Thus, Theorem 5.1 does not tell us whether the PEP itself is a CSSB. That such is always the case is asserted in the following theorem.

Theorem 5.3. *Let $G = (P, u)$ be a multistage game that admits a perfect equilibrium. Then PEP is a CSSB.*

Theorem 5.1 allows us to obtain results for multistage games by applying our

¹⁸PEP [GPEP] stand for [Generalized] Perfect Equilibrium Paths.

results on conservative stability in multistage situations. For example, Theorem 4.2 yields the following extension¹⁹ of Harris (1985a):

Corollary 5.4. *Let $G = (P, u)$ be a bounded from above and continuous at infinity multistage game with perfect information. Then G admits generalized perfect ϵ -equilibrium for every $\epsilon > 0$.*

Remark 5.5. *For games that are continuous at infinity, by the one stage deviation principle $GPEP=PEP$ (see e.g., Harris 1985a and Fudenberg and Tirole 1991). Thus, for such games Theorem 5.1 is valid when ϵ - $GPEP$ and ϵ - $CSSB$ are replaced by PEP and $CSSB$, respectively, thereby generalizing Greenberg (1989,1990).*

It is noteworthy and verified in Example 5.8 below, that the one stage deviation principle does not apply to ϵ -one stage deviation. In particular, therefore, for $\epsilon > 0$, ϵ - $GPEP$ does not coincide with ϵ - PEP , not even in continuous multistage games (in fact, not even in finite game trees with perfect information).

Remark 5.5 yields the following generalization of Abreu's (1988) characterization, using "optimal penal codes", of perfect equilibrium in repeated games with discounting.²⁰

Theorem 5.6. *Let $G = (P, u)$ be a continuous multistage game. Then there exists a nonempty valued standard of behavior σ such that for every (x, t) , $\sigma(x, t)$ contains at most n paths and for every $y \in P$: y is a perfect equilibrium path if and only if the standard of behavior obtained from σ by adding y to $\sigma(x, 1)$ for all $x \in P$ is conservative internally stable.*

We end this section by showing that the ϵ - PEP and the ϵ - $GPEP$ are different concepts. The first example demonstrates that PEP may differ from $GPEP$ if the

¹⁹Harris (1985a) imposed stronger conditions (such as continuity) and considered only the case of $\epsilon = 0$.

²⁰Note that such games are continuous.

game is not continuous at infinity. The second example shows that for $\epsilon > 0$, ϵ -GPEP and ϵ -PEP may differ even in a finite game tree with perfect information and finite outdegree.

Example 5.7.

Consider an infinitely repeated game, where the utility functions $(u^i)_{i \in N}$ depend only on the tail of the path. That is, for every player i ,

$$u^i(x_1, x_2, x_3, \dots) = u^i(x_2, x_3, \dots) \quad \text{for every } x \in S.$$

For example,

$$u^i(x) = \lim_{T \rightarrow \infty} \frac{g^i(x_1) + g^i(x_2) + \dots + g^i(x_T)}{T},$$

where g^i denotes the payoff function of player i in the one-shot game and \lim stands for \liminf , \limsup , or any Banach limit. Clearly, this property does not, in general, imply that every path is a perfect equilibrium path. In contrast, as we show in the Appendix,

$$\text{GPEP}(x, t) = P(x, t) \quad \text{for all } (x, t).$$

Example 5.8.

Consider the following one player game tree: The player can choose L or R at stage 1. If he chooses L , he gets a payoff of 1. If he chooses R , he reaches a node where he can choose U or D . U yields a payoff of 2 and D yields a payoff of 3.

Obviously for all ϵ , $2 > \epsilon > 1$, ϵ -GPEP includes the three paths – $\{L, (R, U), (R, D)\}$, where ϵ -PEP contains only the two paths – $\{(R, U), (R, D)\}$.

6. Bounded Rationality.

The fact that multistage situations employ paths allows us to introduce a new variant of bounded rationality where, at every subsituation, players consider deviations from a proposed path only at one of the k successive periods, for some fixed

$k \geq 1$. This leads to the definition of k -CSSB, given below. Players are *k-rational* if they follow a k -CSSB. We show that “ k -rationality” converges to “full rationality” when $k \rightarrow \infty$, which, of course, is a desirable property.

Applying “ k -rationality” to multistage games yields a new solution concept, demonstrating that TOSS not only extends game theory, but it also opens the way to new solution concepts for games. In particular, our new notion is neither related to playing automatons with a fixed size (as in Neyman, 1985), nor is it related to the concept of k -perfection (see e.g., Chakrabarti 1991). It turns out, however, as shown at the end of this section, that in the finitely repeated Prisoner’s Dilemma game we get similar (but not identical) results to those in Neyman (1985). In particular, if players are “ k -rational” then the unique k -CSSB of this game contains the path in which the cooperative outcome is played after the first k steps (and the noncooperative outcome in the first k steps).

Let σ be a standard of behavior for the multistage situation $G = (\Pi, u)$, let $\epsilon \geq 0$, and let $k \geq 1$. For $x \in \Pi$ and stage t define $\Delta_k^\epsilon(\sigma, (x, t))$ as the set of all $y \in \Pi(x, t)$ for which there exist $i \in N$, $t \leq \tau < t + k$, and $z \in \Pi(y, \tau)$ satisfying:

- (i) for every $j \neq i : z_\tau^j = y_\tau^j$,
- (ii) $\sigma(z, \tau + 1) \neq \emptyset$, and
- (iii) $u^i(\eta) > u^i(y) + \epsilon$ for every $\eta \in \sigma(z, \tau + 1)$.

We say that σ is (ϵ, k) -conservative internally stable if for every (x, t) ,

$$\sigma(x, t) \subseteq \Pi(x, t) \setminus \Delta_k^\epsilon(\sigma, (x, t)),$$

σ is (ϵ, k) -conservative externally stable if for every (x, t) ,

$$\sigma(x, t) \supseteq \Pi(x, t) \setminus \Delta_k^\epsilon(\sigma, (x, t)),$$

σ is (ϵ, k) -conservative stable ((ϵ, k) -CSSB) if for every (x, t) ,

$$\sigma(x, t) = \Pi(x, t) \setminus \Delta_k^\epsilon(\sigma, (x, t)).$$

We write “ k -conservative” for “ $(0, k)$ -conservative”.

Analogous to Theorem 4.1, it can be proved that for $\epsilon > 0$ and $k \geq 1$, every continuous²¹ multistage situation admits a unique (ϵ, k) -CSSB, denoted by σ_k^ϵ . Moreover, following the proof of Lemma 4.5, it can be shown that for $\epsilon > 0$ and $k \geq 1$, σ_k^ϵ is compact valued whenever it is nonempty valued, in which case the analog of Theorem 3.4 guarantees the existence of a largest nonempty valued (ϵ, k) -CSSB. Whenever the multistage situation admits a largest nonempty valued CSSB (k -CSSB) we denote it by σ^L (σ_k^L). Recall that in this case, by Remark 5.5, $\sigma^L = PEP$. Obviously, for all $k \geq 1$, $\sigma^L \subseteq \sigma_k^L$.

Theorem 6.1. *Let $G = (\Pi, u)$ be a continuous multistage situation. If σ_k^ϵ is nonempty valued for every $\epsilon > 0$ and $k \geq 1$, then G admits a nonempty valued k -CSSB for every $k \geq 1$. Moreover, the largest nonempty valued CSSB satisfies:*

$$\sigma^L = \bigcap_{k=1}^{\infty} \sigma_k^L,$$

We now study the notion of “ k -rationality” in the following finitely repeated Prisoner’s Dilemma game with a fixed discount factor $\delta \in (0, 1)$:

	a	b
a	1,1	5,0
b	0,5	4,4

That is, the action set for both players is $S_t^1 = S_t^2 = A = \{a, b\}$ for all $t \geq 1$.

Thus

$$S = \{(x_t)_{t=1}^{\infty} : x_t \in A \times A \text{ for every } t \geq 1\}.$$

Let $g^1, g^2 : A \times A \rightarrow R$ denote the payoff functions in the one-shot game, where for $i = 1, 2$, $g^i(a, a) = 1$, $g^i(b, b) = 4$, $g^1(a, b) = 5$, $g^2(a, b) = 0$, $g^1(b, a) = 0$, and

²¹In fact, continuity at infinity is sufficient.

$g^2(b, a) = 5$. The (discounted) payoff functions in the repeated game, $u^1, u^2 : S \rightarrow R$, are given by

$$u^i(x) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} g^i(x_t) \quad \text{for every } x \in S, i = 1, 2.$$

The fact that this game admits a (unique) subgame perfect equilibrium and that for all $k \geq 1$, $\sigma^L \subseteq \sigma_k^L$ imply that this game admits a nonempty valued k -CSSB, which, by Greenberg (1990, Theorem 5.4.1) is unique. The following claim asserts that if the players are “ k -rational” then this unique k -CSSB contains the path in which the cooperative outcome is played after the first k steps and the noncooperative outcome in the first k steps. Recall that the cooperative path is denoted by $c = ((b, b), (b, b), (b, b), \dots)$ and denote $e = ((a, a), (a, a), (a, a), \dots)$.

Claim 6.2. *For the T -fold repeated Prisoner’s Dilemma game, if $T > k$, then $x(k) = (e^{k+1}, c)$ belongs to $\sigma_k^L(x, 1)$ for every $x \in \Pi$.*

Our analysis of the finitely repeated Prisoner Dilemma game and Claim 6.2 can be easily generalized as follows:

Theorem 6.3. *Consider an n -person one-shot game in which $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ is a Nash equilibrium and $\beta = (\beta_1, \beta_2, \dots, \beta_n)$ weakly Pareto dominates α . Then, for the T -fold repeated discounted game ($1 \leq T \leq \infty$) and for $k < T$,*

$$x(k) \in \sigma_k^L(x, 1) \quad \text{for all } x \in \Pi,$$

where $x(k) = (A^{k+1}, B)$ with $A = (\alpha, \alpha, \alpha \dots)$ and $B = (\beta, \beta, \beta \dots)$.

It can be easily seen that Theorem 6.3 can be further generalized if we assume that β strongly Pareto dominates α . In this case the cooperative path B itself belongs to $\sigma_k^L(x, 1)$ for T and δ sufficiently large.

7. Appendix (Proofs).

Proof of Claim 3.3.

(1) Follows from ϵ -conservative external stability. For a more detailed proof see Greenberg (1990, Claim 2.5.1).

(2) Obviously $\Pi(x, t)$, and hence $\Delta^\epsilon(\sigma, (x, t))$, depend only on x^t . Therefore the assertion follows from the definition of ϵ -conservative stability.

(3) Assume in negation that the truncation property does not hold. Then there exists $y \in \sigma(x, t)$ such that $y \notin \sigma(y, \tau)$ for some $\tau \geq t$. By ϵ -conservative external stability, $y \in \Delta^\epsilon(\sigma, (y, \tau)) \subseteq \Delta^\epsilon(\sigma, (y, t))$. As shown in part (2) of this claim

$$\Delta^\epsilon(\sigma, (y, t)) = \Delta^\epsilon(\sigma, (x, t)).$$

Hence $y \in \sigma(x, t) \cap \Delta^\epsilon(\sigma, (x, t))$, contradicting the ϵ -conservative internal stability of σ . \square

Proof of Theorem 3.4. Define the SB σ^L as follows:

$$(3.4.1) \quad \sigma^L(x, t) = \bigcup_{\sigma \in \Sigma} \sigma(x, t).$$

Obviously σ^L is nonempty valued. We proceed to show that it is ϵ -conservative internally stable.

Indeed, let $y^L \in \sigma^L(x, t)$. We have to show that $y^L \notin \Delta^\epsilon(\sigma^L, (x, t))$. By (3.4.1) there exists $\sigma \in \Sigma$ such that $y^L \in \sigma(x, t)$. By the ϵ -conservative internal stability of σ , $y^L \notin \Delta^\epsilon(\sigma, (x, t))$. Hence, it suffices to show that for every (x, t) ,

$$(3.4.2) \quad \Delta^\epsilon(\sigma, (x, t)) \supseteq \Delta^\epsilon(\sigma^L, (x, t)).$$

Let $y \in \Delta^\epsilon(\sigma^L, (x, t))$. Then, there exists a player $i \in N$ who can deviate from y at some future stage $\tau \geq t$, and induce a path $z \in \Pi(y, \tau)$ such that (3.1) and (3.2)

are satisfied for σ^L . As σ is nonempty valued, (3.1) is satisfied for σ . As $\sigma \subseteq \sigma^L$, (3.2) is also satisfied by σ . Hence $y \in \Delta^\epsilon(\sigma, (x, t))$, which proves (3.4.2).

It remains to show that σ^L (defined in 3.4.1) is ϵ -conservative externally stable. Suppose, in negation, that it is not. Then there exist (x_0, t_0) and $y_0 \in \Pi(x_0, t_0)$ but $y_0 \notin \sigma^L(x_0, t_0) \cup \Delta^\epsilon(\sigma^L, (x_0, t_0))$. Define the SB σ as follows:

$$\sigma(x_0, t_0) = \sigma^L(x_0, t_0) \cup \{y_0\},$$

and

$$\sigma(x, t) = \sigma^L(x, t) \quad \text{for all } (x, t) \neq (x_0, t_0).$$

It can be easily verified that σ is ϵ -conservative internally stable and nonempty valued, contradicting the maximality of σ^L . \square

Proof of Claim 3.5. Let σ^L be the largest ϵ -CSSB. As $\epsilon_i \geq \epsilon$, $i = 1, 2$, σ^L is ϵ_i -conservative internally stable. Therefore by Theorem 3.4 the largest ϵ_i -CSSB, σ^{L_i} , exists. As σ^{L_1} is ϵ_2 -conservative internally stable, $\sigma^{L_1} \subseteq \sigma^{L_2}$. \square

Proof of Theorem 4.1. Let $\epsilon > 0$. By continuity at infinity there exists an integer J such that for every $i \in N$, for every $x \in \Pi$, and for every $t \geq J$,

$$\sup_{y, z \in \Pi(x, t)} |u^i(z) - u^i(y)| \leq \epsilon.$$

For every $x \in \Pi$ and for every $\tau \geq J$ define

$$\sigma(x, \tau) = \Pi(x, \tau).$$

For $t < J$ define σ by the following recursive formula:

$$\sigma(x, t) = \Pi(x, t) \setminus \Delta^\epsilon(\sigma, (x, t)).$$

Obviously σ is the unique ϵ -CSSB. \square

Proof of Theorem 4.2. Let σ be an ϵ -CSSB for G , and assume, in negation, that there exist $x \in \Pi$ and $t \geq 1$ such that $\sigma(x, t) = \emptyset$. By ϵ -conservative external stability, there exist $\tau \geq t$ and $y \in \Pi(x, t)$ such that $\sigma(y, \tau + 1) \neq \emptyset$. Set

$$(4.2.1) \quad k = \min\{\tau \geq t : \text{there exists } z \in \Pi(x, t) \text{ with } \sigma(z, \tau + 1) \neq \emptyset\}.$$

Let $\bar{y} \in \Pi(x, t)$ satisfy $\sigma(\bar{y}, k + 1) \neq \emptyset$. Since the multistage situation is with perfect information there exists at most one active player, j , at (\bar{y}, k) .

Denote

$$M^j(y) = \{z \in C_k^j(y) : \sigma(z, k + 1) \neq \emptyset\}.$$

As $\bar{y} \in M^j(\bar{y})$, $M^j(\bar{y}) \neq \emptyset$. Denote

$$(4.2.2) \quad \alpha = \sup\{u^j(\eta) : \exists z \in M^j(\bar{y}) \text{ with } \eta \in \sigma(z, k + 1)\}$$

Since the multistage situation is bounded from above, $\alpha < \infty$. Let $z \in M^j(\bar{y})$ and $\eta \in \sigma(z, k + 1)$ satisfy

$$(4.2.3) \quad u^j(\eta) > \alpha - \frac{\epsilon}{2}.$$

We shall now show that $\eta \in \sigma(\bar{y}, k)$. Indeed, since σ is ϵ -conservative stable, it suffices to show that $\eta \notin \Delta^\epsilon(\sigma, (\bar{y}, k))$. Indeed, if $i \neq j$, then i is non-active at (\bar{y}, k) and therefore i cannot ϵ -benefit by deviating from η at k . And, by (4.2.3), neither can player j . No player can ϵ -benefit by a deviation from η at a stage larger than k , since $\eta \in \sigma(z, k + 1)$. It follows that $k = t$, as otherwise we get a contradiction to 4.2.1. This implies that $\eta \in \sigma(\bar{y}, t)$. By Claim 3.3, $\sigma(x, t) = \sigma(\bar{y}, t)$. Hence, $\sigma(x, t) \neq \emptyset$. A contradiction. \square

Proof of Theorem 4.3. The proof proceeds as the proof of Theorem 4.2 up to the definition of α in (4.2.2). As G is not necessarily bounded from above we cannot exclude $\alpha = \infty$. We therefore replace α by β which is defined as follows: Let A

be the set of all $a \in A_k(\bar{y})$ for which there exists $z \in M^j(\bar{y})$ with $z_k = a$. As $\bar{y}_k \in A$, $A \neq \emptyset$. For every $a \in A$, choose $z(a) \in M^j(\bar{y})$ with $z(a)_k = a$, and choose $\eta(z(a)) \in \sigma(z(a), k+1)$. Because the multistage situation has a finite outdegree, A is a finite set, and therefore we can define β in the following way:

$$\beta = \max\{u^j(\eta(z(a))) : a \in A \text{ and } z(a) \in M^j(\bar{y})\}.$$

We now choose $b \in A$ such that $u^j(\eta(z(b))) = \beta$, denote $\eta(z(b))$ by η , and proceed as in the proof of Theorem 4.2 \square

Proof of Lemma 4.5. Let $\epsilon > 0$. We prove the stronger result that $\sigma^\epsilon(\cdot, t)$ has a closed graph for every $t \geq 1$. By continuity at infinity there exists an integer J such that for every $i \in N$, for every $x \in \Pi$, and for every $t \geq J$,

$$\sup_{y, z \in \Pi(x, t)} |u^i(z) - u^i(y)| \leq \epsilon.$$

For $\tau \geq J$ we have

$$(4.5.1) \quad \sigma^\epsilon(x, \tau) = \Pi(x, \tau) \quad \text{for every } x \in \Pi.$$

As Π is compact, (4.5.1) implies that $\sigma^\epsilon(\cdot, \tau)$ has a closed graph. We prove the claim for $t < J$ by backward induction. Suppose the claim is valid for every $\tau > t$.

Let $x_k \rightarrow x$ and $y_k \rightarrow y$ with $y_k \in \sigma^\epsilon(x_k, t)$ and assume in negation that $y \notin \sigma^\epsilon(x, t)$. By ϵ -conservative external stability of σ^ϵ , there exist $i \in N$, $t \leq \tau$ and $z \in \Pi(y, \tau)$ with

- (i) for every $j \neq i$, $z_\tau^j = y_\tau^j$;
- (ii) $\sigma^\epsilon(z, \tau + 1) \neq \emptyset$, and
- (iii) $u^i(\eta) > u^i(y) + \epsilon$ for every $\eta \in \sigma^\epsilon(z, \tau + 1)$.

Since $C_\tau^i(\cdot)$ is a lower semicontinuous correspondence, $y_k \rightarrow y$ and $z \in C_\tau^i(y)$, there exists $z_k \rightarrow z$ with $z_k \in C_\tau^i(y_k)$. As σ^ϵ is nonempty valued, $\sigma^\epsilon(z_k, \tau + 1) \neq \emptyset$.

Because $y_k \in \sigma^\epsilon(x_k, t)$ and because σ^ϵ is ϵ -conservative internally stable, there exists

$$(4.5.2) \quad \eta_k \in \sigma^\epsilon(z_k, \tau + 1)$$

with

$$u^i(\eta_k) \leq u^i(y_k) + \epsilon.$$

As Π is compact, $\eta_{k_r} \rightarrow \bar{\eta}$ for a subsequence. Therefore, by continuity of u^i ,

$$u^i(\bar{\eta}) \leq u^i(y) + \epsilon.$$

By (4.5.2) and by our inductive hypothesis, $\bar{\eta} \in \sigma^\epsilon(z, \tau + 1)$ contradicting (iii) above. \square

Proof of Theorem 4.6. By Lemma 4.5, σ^ϵ is compact valued for all $\epsilon > 0$. By Claim 3.5, $\sigma^{\epsilon_1} \subseteq \sigma^{\epsilon_2}$ whenever $\epsilon_1 < \epsilon_2$. Therefore σ is nonempty valued. We proceed to prove that σ is conservative stable.

Conservative internal stability: Let $y \in \sigma(x, t)$, let $i \in N$, and let $\tau \geq t$. We have to show that player i cannot benefit by unilaterally deviating from y at stage τ . Let then $z \in \Pi(y, \tau)$ satisfy $z_\tau^j = y_\tau^j$ for all $j \neq i$. As $y \in \sigma(x, t)$, we have that for every $\epsilon > 0$, $y \in \sigma^\epsilon(x, t)$. By ϵ -conservative internal stability of σ^ϵ , there exists $\eta^\epsilon \in \sigma^\epsilon(z, \tau + 1)$ such that

$$(4.6.1) \quad u^i(\eta^\epsilon) \leq u^i(y) + \epsilon \quad \text{for every } \epsilon > 0.$$

Since Π is compact, $\Pi(z, \tau + 1)$ is compact and therefore there exists²² a vanishing decreasing sequence $(\epsilon_n)_{n=1}^\infty$ and $\eta \in \Pi(z, \tau + 1)$ such that $\lim_{n \rightarrow \infty} \eta^{\epsilon_n} = \eta$. We now show that $\eta \in \sigma(z, \tau + 1)$. Let $\epsilon > 0$. Choose $m \geq 1$ with $\epsilon_m < \epsilon$. By Claim 3.5, $\eta^{\epsilon_n} \in \sigma^{\epsilon_m}(z, \tau + 1)$ for all $n \geq m$. By Lemma 4.5 $\eta \in \sigma^{\epsilon_m}(z, \tau + 1)$, because

²² $\Pi(z, \tau + 1)$ is a compact metric space and hence it is sequentially compact.

$\sigma^{\epsilon_m}(z, \tau + 1)$ is closed. Hence, $\eta \in \sigma^\epsilon(z, \tau + 1)$. Therefore $\eta \in \bigcap_{\epsilon > 0} \sigma^\epsilon(z, \tau + 1) = \sigma(z, \tau + 1)$.

The continuity of u^i together with (4.6.1) imply that $u^i(\eta) \leq u^i(y)$. Hence, σ is conservative internally stable.

Conservative external stability: Let $y \in \Pi(x, t) \setminus \sigma(x, t)$. Then, there exists $\epsilon > 0$ such that $y \notin \sigma^\epsilon(x, t)$. By the ϵ -conservative stability of σ^ϵ , there exists a player and an ϵ -profitable deviation for that player with respect to σ^ϵ . This deviation is also profitable for the player with respect to σ because $\sigma \subseteq \sigma^\epsilon$, σ is nonempty valued, and $\epsilon > 0$.

We have to show that σ is the largest nonempty valued CSSB. Denote the largest nonempty valued CSSB by $\tilde{\sigma}$. Obviously, $\sigma \subseteq \tilde{\sigma}$. To prove the reversed inclusion, note that by Claim 3.5, $\tilde{\sigma} \subseteq \sigma^\epsilon$ for every $\epsilon > 0$. Therefore $\tilde{\sigma} \subseteq \bigcap_{\epsilon > 0} \sigma^\epsilon = \sigma$. \square

Verification of Example 4.8. We describe this game as an (infinite) multistage game (P, u) as follows: Denote $\bar{U} = (U, U, \dots)$, and for $\alpha \in R$, $\bar{\alpha} = (\alpha, \alpha, \dots)$. Then,

$$P = \{\bar{U}\} \cup \{(D, \bar{\alpha}) : \alpha \in R\},$$

$$P(\bar{U}, t) = \{\bar{U}\} \quad \text{for all } t > 1,$$

$$P((D, \bar{\alpha}), 2) = \{(D, \bar{\beta}) : \beta \in R\} \quad \text{for all } \alpha \in R,$$

and

$$P((D, \bar{\alpha}), t) = \{(D, \bar{\alpha})\} \quad \text{for all } \alpha \in R \text{ and } t \geq 3.$$

The payoff function is given by $u(\bar{U}) = 1$, and $u((D, \bar{\alpha})) = \alpha$ for every $\alpha \in R$.

Define the standard of behavior σ as follows: $\sigma(\bar{U}, t) = \{\bar{U}\}$ for all $t \geq 1$, $\sigma((D, \bar{\alpha}), 2) = \emptyset$ for all $\alpha \in R$, and $\sigma((D, \bar{\alpha}), t) = \{(D, \bar{\alpha})\}$ for all $\alpha \in R$ and $t \geq 3$.

It can be easily seen that σ is the unique ϵ -CSSB for every $\epsilon \geq 0$.

Verification of Example 4.9. Assume, in negation, that G admits a CSSB σ . Note, first, that the path $C^* = (C, \dots, C, \dots)$ of never exiting, does not belong to $\sigma(C^*, t)$ for all t . Indeed, the path yields the payoff 0, while the path (C, \dots, C, E) of exiting at period $(t + 1)$ yields a positive payoff. Moreover, by Theorem 4.3, $\sigma(C^*, t)$ is nonempty, for all t . Therefore, consider $y \in \sigma(x, 1)$. Since y is of the form $y = (C, C, \dots, C, E)$ the player will deviate from E (at stage τ) to C , inducing the subtree $(x, \tau + 1)$, yielding him the payoff which is at least strictly greater than his payoff under y . Thus, $y \notin \sigma(x, 1)$. A contradiction.

Proof of Theorem 5.1. In view of Theorem 3.4 it suffices to prove the following two claims:

Claim 5.1.1. *If ϵ -GPEP is nonempty valued then it is ϵ -conservative internally stable.*

Claim 5.1.2. *Let σ be the largest ϵ -conservative internally stable nonempty valued standard of behavior. Then*

$y \in \sigma(x, t) \implies y$ is a generalized perfect ϵ -equilibrium path in the subgame (x, t) .

Proof of claim 5.1.1. Let $y \in \epsilon\text{-GPEP}(x, t)$. Let f be a generalized perfect ϵ -equilibrium that supports y . That is, $f_{(x, t)}$ is a generalized perfect ϵ -equilibrium in the subgame (x, t) and $p(f) = y$.

We now show that $y \notin \Delta^\epsilon(\epsilon\text{-GPEP}, (x, t))$. Indeed, suppose that z is obtained from y by a unilateral single-stage deviation of i at stage $\tau \geq t$. Denote by η the path generated by the restriction of f to $(z, \tau + 1)$. Then $\eta \in \epsilon\text{-GPEP}(z, \tau + 1)$, and $u^i(\eta) \leq u^i(y) + \epsilon$. Hence, $y \notin \Delta^\epsilon(\epsilon\text{-GPEP}, (x, t))$.

Proof of Claim 5.1.2. Without loss of generality assume $t = 1$. Let $x_0 \in \Pi$ and let $y \in \sigma(x_0, 1)$. Recall by Claim 3.3 that $\sigma(x, 1)$ does not depend on x . We

first construct, recursively, a 1-valued standard of behavior $\bar{\sigma}$ with $\bar{\sigma} \subseteq \sigma$, and $\bar{\sigma}(x_0, 1) = \{y\}$. As $\bar{\sigma}$ is 1-valued we identify it with the single path it contains. In particular, we omit the parentheses $\{ \}$ from its description. For $t = 1$, $\bar{\sigma}(x, 1) = y$ for all $x \in P$. Assume that $\bar{\sigma}(x, \tau)$ was defined for all $x \in P$ and for all $1 \leq \tau \leq t$. To define $\bar{\sigma}(x, t + 1)$ distinguish among the following three cases:

1. $\bar{\sigma}(x, t) \in P(x, t + 1)$: Then define $\bar{\sigma}(x, t + 1) = \bar{\sigma}(x, t)$.²³
2. $\bar{\sigma}(x, t)$ is obtained from x by a unilateral deviation of player i at stage t (i.e., $\{j \in N : x_t^j \neq \bar{\sigma}(x, t)_t^j\} = \{i\}$): Then, by ϵ -conservative internal stability there exists $z \in \sigma(x, t + 1)$ such that $u^i(z) \leq u^i(\bar{\sigma}(x, t)) + \epsilon$. Choose $\bar{\sigma}(x, t + 1) = z$.
3. If neither (1) nor (2) above applies, choose an arbitrary path $\bar{\sigma}(x, t + 1) \in \sigma(x, t + 1)$.

To complete the proof of this claim (and hence of Theorem 5.1) define the n -tuple of strategies f to be such that the n -tuple of actions chosen at subgame (x, t) are those specified by $\bar{\sigma}(x, t)_t$. It can be easily verified that $p(f) = y$ and that f is a generalized perfect ϵ -equilibrium. \square

Proof of Theorem 5.3. With the obvious notational changes the proof mimics the proof of Theorem 5.1. \square

Proof of Theorem 5.6. For every $i \in N$, $x \in \Pi$, and $t \geq 1$, choose $y(i, (x, t)) \in PEP(x, t)$ such that

$$u^i(y(i, (x, t))) = \min_{z \in PEP(x, t)} u^i(z).$$

Define, for $x \in \Pi, t \geq 1$

$$\sigma(x, t) = \{y(i, (x, t)) : i \in N\}.$$

²³Observe that in this case, by Claim 3.3, $\bar{\sigma}(x, t) \in \sigma(x, t + 1)$.

Then, by Remark 5.2 and by the choice of the paths $y(i, (x, t))$, σ is a nonempty valued conservative internally stable standard of behavior. Let $y \in P$. Let σ^y be the standard of behavior obtained from σ by adding y to $\sigma(x, t)$. Then if it is a nonempty valued conservative internally stable standard of behavior, then by Remark 5.2, $\sigma^y \subseteq \text{PEP}$ and therefore y is a perfect equilibrium path. As for the converse, if y is a perfect equilibrium path then σ^y is conservative internally stable by the choice of the paths $y(i, (x, t))$. \square

Verification of Example 5.7. By Remark 5.2, it suffices to prove that P is a conservative internally stable standard of behavior. Indeed, let $y \in P(x, t)$. Assume that player i deviates from y once, at stage τ , using η_τ^i instead of y_τ^i . Since the path $z = (y_1, \dots, y_{\tau-1}, \eta_\tau, y_{\tau+1}, \dots)$ belongs to $P(z, \tau)$, and $u^i(z) = u^i(y)$, we conclude that $\Delta^0(P, (x, t)) = \emptyset$. Therefore P is conservative internally stable. \square

Proof of Theorem 6.1.: We need the following claims. For $\epsilon > 0$, denote by σ^ϵ the unique ϵ -CSSB of G .

Claim 6.1.1.. *Let $\epsilon > 0$. Then σ^ϵ is nonempty valued and*

$$\sigma^\epsilon = \bigcap_{k=1}^{\infty} \sigma_k^\epsilon.$$

proof of Claim 6.1.1.:

It is easily verified that $\Delta_{k_1}^{\epsilon_1}(\sigma, (x, t)) \subseteq \Delta_{k_2}^{\epsilon_2}(\sigma, (x, t))$ for all $k_1 \leq k_2$ and for all $\epsilon_1 \geq \epsilon_2$. Hence, if σ is conservative internally $k_2 - \epsilon_2$ -stable, then σ is conservative internally $k_1 - \epsilon_1$ -stable. In particular $(\sigma_k^\epsilon)_{k \geq 1}$ are nested in k . Moreover, it is obvious that there exists $k(\epsilon)$ such that $\sigma_k^\epsilon = \sigma^\epsilon$ for every $k \geq k(\epsilon)$. Hence the result.

The proof of the following claim goes along the lines of the proof of Theorem 4.6 and hence it is omitted.

Claim 6.1.2.. *Let $k \geq 1$. Then there exists a nonempty valued k -CSSB, and*

$$\sigma_k^L = \bigcap_{\epsilon > 0} \sigma_k^\epsilon.$$

We now return to the proof of Theorem 6.1. By Claim 6.1.1, σ^ϵ is nonempty valued for every $\epsilon > 0$. Therefore by Theorem 4.6, G admits a nonempty valued CSSB and

$$\sigma^L = \bigcap_{\epsilon > 0} \sigma^\epsilon.$$

By Claim 6.1.1,

$$\bigcap_{\epsilon > 0} \sigma^\epsilon = \bigcap_{\epsilon > 0} \bigcap_{k=1}^{\infty} \sigma_k^\epsilon = \bigcap_{k=1}^{\infty} \bigcap_{\epsilon > 0} \sigma_k^\epsilon,$$

by changing the order of taking the intersections. Finally, by Claim 6.1.2,

$$\bigcap_{k=1}^{\infty} \bigcap_{\epsilon > 0} \sigma_k^\epsilon = \bigcap_{k=1}^{\infty} \sigma_k^L. \quad \square$$

Proof of Claim 6.2. Assume, in negation, that $x(k) \notin \sigma_k^L(x, 1)$. Then one of the players, say Player 1, can benefit from a deviation after at most k steps, say at stage τ . Hence, it must be that all paths recommended by σ_k^L after this deviation yield Player 1 more than $u^1(x(k))$. But this is impossible, because $(z^t, e) \in PEP(z, t) = \sigma^L(z, t) \subseteq \sigma_k^L(z, t)$ for every $z \in \Pi$ and $t \leq T$, and (a, a) is an equilibrium for the one-shot Prisoner's Dilemma game. \square

References.

Alesina, A., and H. Rosenthal (1992): "Government Formation and Endogenous Parties," *Harvard D.P.* 1161.

Abreu, D (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 50, 383–396.

Abreu, D., D. Pearce, and E. Stacchetti (1986): "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, 251–269.

Asheim G. B. (1988): “Renegotiation-Proofness in Finite and Infinite Horizon Games through the Theory of Social Situations,” DP A-173, University of Bonn.

Asheim G. B. (1991): “Extending Renegotiation-Proofness to Infinite Horizon Games,” *Games and Economic Behavior*, 3, 278-294.

Asheim G. B. (1992): “Self-Enforcing Agreements in Multistage Games,” *mimeo*.

Bergin, J., and B. Macleod (1993): “Efficiency and Renegotiation in Repeated Games,” *Journal of Economic Theory*, 61, 42-73.

Börgers, T. (1989): “Perfect Equilibrium Histories of Finite and Infinite Horizon Games,” *Journal of Economic Theory*, 47, 218-227.

Börgers, T. (1991): “Upper Hemicontinuity of the Correspondence of Subgame Perfect Equilibrium Outcomes,” *Journal of Mathematical Economics*, 20, 89-106.

Chakrabarti, S. K. (1991): “Finite Complexity and the Folk Theorem in Repeated Games,” *Economics Letters*, 35, 355-358.

Debreu, G. (1952): “A Social Equilibrium Existence Theorem,” *Proceedings of the National Academy of Sciences of the U.S.A.*, 38, 886-893.

Fudenberg, D., and D. Levine (1983): “Subgame Perfect Equilibrium of Finite and Infinite Horizon Games,” *Journal of Economic Theory*, 31, 251-268.

Fudenberg, D., and D. Levine (1986): “Limit Games and Limit Equilibria,” *Journal of Economic Theory*, 38, 261-279.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. The MIT Press.

Gillette, D. (1957): “Stochastic Games with Zero Stop Probabilities,” in *Contributions to the Theory of Games III*, M. Dresher, A.W. Tucker and P. Wolfe (eds.). Princeton University Press.

Greenberg, J. (1989): “An application of The Theory of Social Situations to Repeated Games,” *Journal of Economic Theory*, 45, 278-293.

Greenberg, J (1990): *The Theory of Social Situations: An Alternative Game-Theoretic Approach*. Cambridge University Press.

Grossman, S.J., and O. D. Hart (1986): “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy*, Vol. 94, 41, 691–719.

Guth, W., Leininger, W., and Stephan, G. (1991): “On Supergames and Folk Theorems: A Conceptual Discussion,” in Selten, R. (ed.): *Game Equilibrium Models II*. Springer-Verlag, Berlin, Germany.

Harris, C. (1985 a): “Existence and Characterization of Perfect Equilibrium in Games of Perfect Information,” *Econometrica*, 53, 613–628.

Harris, C. (1985 b): “A Characterization of the Perfect Equilibria of Infinite Horizon Games,” *Journal of Economic Theory*, 37, 99–125.

Harris, C., P. Reny, and A. Robson (1995): “The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: a Case for public randomization,” *Econometrica*, 63, 3, 507-544.

Harsanyi, J. C., and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*. The MIT press, London, England.

Hellwig, M., and W. Leininger (1987): “On the Existence of Subgame Perfect Equilibrium in Infinite-Action Games of Perfect Information,” *Journal of Economic Theory*, 43, 55–75.

Hellwig, M., W. Leininger., P. Reny, and A. Robson (1990): “Subgame Perfect Equilibrium in Continuous Games of Perfect Information: An Elementary Approach to Existence and Approximation by Discrete Games,” *Journal of Economic Theory*, 52, 406–422.

Leininger, W. (1988): “Escalation and Cooperation in Conflict Situations: The Dollar Auction Revisited,” *Journal of Conflict Resolution*, 33, 321–254.

Mariotti, M. (1993): “A Theory of Agreements in Normal Form Games,” *mimeo*, University of London.

Muto, S., and M. Nakayama (1992): “The Resale-Proof Trade of Information as a Stable Standard of Behavior—An Application of the Theory of Social Situations,” *mimeo*, Tohoku University.

Muto, S., and D. Okada (1992a): “von Neumann-Morgenstern Stable Sets in a Price-Setting Duopoly,” *mimeo*, Tohoku University.

Muto, S., and D. Okada (1992b): “von Neumann-Morgenstern Stable Sets in Cournot Competition,” *mimeo*, Tohoku University.

Neyman, A. (1985): “Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner’s Dilemma,” *Economics letters*, 19, 227–229.

Rubinstein, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50, 97-110.

Schmeidler, D. (1973): “Equilibrium Points of Nonatomic Games,” *Journal of Statistical Physics*, Vol. 7, 4, 295–300.

Shitovitz, B. (1994): “Optimistic Stability in Games of Perfect Information,” *Mathematical Social Sciences*, 28, 199–214.

Sorin, S. (1986): “On Repeated Games with Complete Information,” *Mathematics of Operations Research*, 11, 147- 160.

Tadelis, S. (1994): “Optimistic Stability in Repeated Extensive Form Games,” *Journal of Economic Theory*, forthcoming.

van Damme, E. (1987): Stability and Perfection of Nash Equilibrium. *Springer Verlag*.

von Neumann, J. and O. Morgenstern. (1947). Theory of Games and Economic Behavior. Princeton University Press.

Xue, L. (1995): “Coalition Stability under perfect Foresight,” *mimeo*, McGill

University.