Seminar Paper No. 470

ON SELF-ENFORCEMENT IN
EXTENSIVE-FORM GAMES

by

Jörgen W. Weibull

INSTITUTE FOR INTERNATIONAL ECONOMIC STUDIES

University of Stockholm

Seminar Paper No. 470

ON SELF-ENFORCEMENT IN EXTENSIVE-FORM GAMES

by

Jörgen W. Weibull

Seminar Papers are preliminary material circulated to
stimulate discussion and critical comment.

August, 1990

Institute for International Economic Studies
S-106 91 Stockholm
Sweden

ON SELF-ENFORCEMENT IN EXTENSIVE-FORM GAMES*

Jörgen W. Weibull**

Institute for International Economic Studies, Stockholm University;
and Department of Economics, Princeton University.

May 1990

Mailing address: Department of Economics, Princeton University, Princeton, N.J.
08544.

* This study is an elaboration of Weibull (1988) and Weibull (1989b).

** The author is grateful for helpful comments to earlier versions of this paper from Robert Aumann, Kaushik Basu, Hans Carlsson, Eric van Damme, Eddie Dekel, David Kreps, Wolfgang Leininger, Christian Schultz, Reinhard Selten and Jeroen Swinkels. Financial support from the Bank of Sweden Tercentenary Foundation and from the Swedish Council for Research in the Humanities and Social Sciences (HSFR) is gratefully acknowledged.

ABSTRACT

This paper discusses the issue of self-enforcement in finite extensive form games. A new criterion is developed which essentially requires that, for any given degree (ϵ) of accuracy, every player's strategy should be locally optimal at each of his information sets under some probabilistic hypothesis about all players' behavior which at least approximately (within ϵ) agrees with the tested profile at all points in the game where his information set provides no evidence of a deviation. Like sequential equilibrium, this criterion is more stringent than subgame perfection but less stringent than ("trembling hand" extensive-form) perfect equilibrium. In contrast to sequential equilibrium, the approach builds on an individualistic and set-theoretic consistency condition which can be viewed as a version of Occam's razor.

1. Introduction

In game theoretic equilibrium analysis, a strategy serves, at least implicitly and partially, two distinct functions: as a contingent plan for a player's choices and as an expression for other players' contingent expectations (predictions) about his behavior.¹ One extreme position is taken by Nash equilibrium: strategies, and local strategies, off the path of a strategy profile are at most tested in their role as expressions for other player's expectations. The point is illustrated in Figure 1 below. The local strategy T for player 1 in the indicated Nash equilibrium $s^*=(Tt,L)$ is tested (for optimality) in its role as a plan for the player. In contrast, the strategy L for player 2 serves only as an expression for 1's expectation — it is not tested in its role as a plan for 2 under the relevant contingency that her information set has been reached. The local strategy t for player 1 at his second decision node, finally, is tested neither as a contingent plan for player 1 at that node, nor as an expression for 2's expectation about 1's action there.

(Figure 1)

Viewing a (behavior) strategy profile as a kind of potential "complete contract", specifying actions under every contingency that the specification of the game allows for, it seems desirable to require that a strategy profile be "self-enforcing" under every contingency in the game. For example, although the strategy profile $s^*=(Tt,L)$ does not specify that player 1 should play so as to make his second decision node reachable, for s^* to be self-enforcing it ought to be in 1's interest to play according to s^* at

¹ Rubinstein (1988) disusses some of the conceptual difficulties associated with strategy notions.

that node, since otherwise the specified action (t) does not serve in its role as a (reasonable) expectation for 2. Since the profile s^* specifies an inoptimal action at 1's second node, it appears to be not "self-enforcing".

The self-enforcement criterion developed in this study essentially requires of a strategy profile that, for any given degree (ϵ) of accuracy, each player's strategy be locally optimal at each of his information sets under some probabilistic hypothesis about all players' behavior which at least approximately (within ϵ) agrees with the tested profile at all points of the game where his information set provides no evidence of a deviation. In other words, *if* the strategy profile would (at least approximately) encode all players' (contingent) expectations, would the strategies then serve as optimal (contingent) plans for the players? Note however, that this does not mean that "rational" players necessarily should believe that others play according to the tested strategy profile; it only says that the profile should not be self-defeating in the sense that some player, at some information set, is willing to play according to the strategy profile *only* if he holds beliefs about others' strategies that diverge more from the profile than he has evidence for. The criterion developed here can hence be viewed as a candidate for a necessary (but in general insufficient) condition for a strategy profile to be "self-enforcing."

The present approach differs methodologically from, but is much inspired by, the path-breaking works of Selten (1965,1975) and Kreps and Wilson (1982). In fact, it is shown below that every extensive-form ("trembling hand") perfect equilibrium meets the self-enforcement criterion suggested here, although the present criterion is weaker than these established approaches in two respects: (a) players with the same (partial) information may hold differing "beliefs", (b) players are not constrained to believe in infinitesimal deviations once they have evidence that a deviation has occurred.

While Selten's (1975) approach was the first to formally allow for the possibility of "irrationality" — viewing "rationality" as a limiting case of a vanishing degree of a certain form of "irrationality" — the innovative feature of Kreps' and Wilson's (1982) approach is the emphasis on "beliefs" held by players at different information sets. More precisely, Kreps and Wilson (1982) define sequential equilibrium in terms of two conditions: sequential rationality and consistency. Though mathematically attractive, their consistency condition is not fully transparent from the viewpoint of decision theory:

"We shall proceed here to develop the properties of sequential equilibrium as defined above; however, we do so with some doubts of our own concerning what 'ought' to be the definition of a consistent assessment that, with sequential rationality, will give the 'proper' definition of a sequential equilibrium." [op.cit. p.876]

In order to motivate their consistency condition, they discuss a number of intuitively appealing notions of consistency embodied in their formal condition (op.cit. Sect.5). However, their prime argument, that consistency implies that beliefs always respect the information structure of the game, was shown by Kreps and Ramey (1987) to be invalid. This possibility of *structural inconsistency* is disturbing:

"... one may well be very concerned with the loss of structural consistency in the limit, especially insofar as structural consistency is used to advocate particular out-of-equilibrium beliefs. Most applications of sequential equilibrium, in deriving beliefs, make explicit use of structural consistency with respect to strategy conjectures which do lie on the boundary of the strategy simplex. The examples show that this approach is not, in general, compatible with perfect or sequential equilibrium." [Kreps and Ramey (1987, p.1342)]

In contrast, the consistency condition in the present approach does imply structural

consistency.²

Kreps and Ramey (1987) also observe that structural consistency may in certain games conflict with sequential rationality. This conflict can arise when players' information sets "cross", reflecting situations in which players do not know the exact chronological order of the moves during the play of the game. The problem is not that sequential rationality requires optimality, but that it defines optimality under the presumption that other players' subsequent moves are governed by equilibrium play, a presumption that may be falsified by the fact that the information set itself has been reached. As will be seen below, this conflict is resolved in the present approach.

Like (extensive-form) perfect and sequential equilibrium, the present self-enforcement criterion is not invariant under all transformations of an extensive-form game that preserve its (reduced) normal form. Three remarks will be made concerning this extensive-form dependence. First, strategies in different extensive forms, corresponding to the same reduced normal form, are in general not equally detailed, and hence strategy profiles corresponding to the same outcome, but being defined in different extensive forms, may differ in detail. Consequently, the objects tested for self-enforcement may differ, and so the solution sets may differ. Secondly, certain distinguished researchers claim, on the basis of backward-induction rationality axioms, that the normal form is insufficient for strategic analysis, see Chapter 10 in Harsanyi and Selten (1988). Thirdly, in view of the acknowledged difficulties of maintaining a theory of strategic interaction which does not allow for the possibility of irrationality (see von Neumann and Morgenstern (1947,p.32), Rosenthal (1981), Binmore (1987), Harsanyi and Selten (1988), Bicchieri (1989), Basu (1990) etc.), I take the desirability

² Fudenberg and Tirole (1989) develop a notion of perfect Bayesian equilibrium, which uses a weaker consistency condition than that in Kreps and Wilson (1982). Also their condition, although different in nature from the present one, obeys the information structure of the game.

of full invariance as undecided.³

The material is organized as follows. After a short section on basic notation and definitions (Section 2), the notion of a strategy prior is developed (Section 3). The new concept of consistency, central to the present approach, is developed in Section 4, and two self-enforcement criteria, based on optimality under approximately consistent priors, are elaborated in Section 5. These criteria are compared with the concepts of sequential and perfect equilibrium in Section 6. Section 7 discusses relations to weak dominance, and Section 8 suggests directions for further research.

2. Notation and basic definitions

The present analysis concerns finite extensive-form games with perfect recall (as defined in Kuhn (1953), see also Selten (1975) or van Damme (1987)). In particular, all strategies are taken to be behavior strategies. The set of (personal) players will be written $I=\{1,2,\dots,n\}$, and nature is referred to as player 0. The set of nodes in the tree is X , with $x < x'$ denoting that the node $x \in X$ precedes the node $x' \in X$, and $x \leq x'$ signifying that either $x < x'$ or $x = x'$. The initial node of the tree is denoted o , and $Z \subset X$ is the subset of terminal nodes. H_i is the set of information sets h of player i , and H designates the union of all (personal) information sets in the game. As usual, a game is said to have *perfect information* if all information sets are singletons.

A (behavior) *strategy* of a player $i \in I$ is written s_i , and his strategy set is S_i . The Cartesian product of all these sets is written S , so a *strategy profile* is a point s in

³ Of course, if certain transformations of extensive forms are deemed strategically "inessential", even when "irrationality" is allowed for, one could impose corresponding, additional conditions to the present "minimalistic" self-enforcement condition.

S. The Cartesian product of all strategy sets except S_i will be written S_{-i} . The local strategy assigned by $s_i \in S_i$ to an information set $h \in H_i$ is denoted $s_{ih} \in S_{ih}$, i.e. s_{ih} is a probability distribution over the (finite) set C_h of *choices* available at h . The probability assigned by s_{ih} to a choice $c \in C_h$ is written $s_{ih}(c)$, and sometimes the same probability is written $s(c)$ for short. Two strategy profiles are said to *agree* at an information set h if they assign the same local strategy to the set. For any strategy profile $s \in S$ and node $x \in X$, $\rho(x, s)$ denotes the *realization probability* of x , i.e. the probability that x will be reached when s is played. The realization probability of an information set h is accordingly written $\rho(h, s) (= \sum_{x \in h} \rho(x, s))$. A profile s will be said to *reach* an information set h if $\rho(h, s) > 0$, and h will then also be said to be *on the path* of s . (All moves by nature are governed by fixed and positive probabilities.)

The von Neumann–Morgenstern utility for player i at a terminal node z is written $u_i(z)$, and his expected utility, or *payoff*, when $s \in S$ is played is

$$(2.1) \quad U_i(s) = \sum_{z \in Z} \rho(z, s) u_i(z) \quad [\forall i \in I, s \in S].$$

Let $\rho(x' | s, x)$ denote the conditional probability that the node $x' \in X$ is reached, under the condition that $s \in S$ is played and play starts at $x \in X$. The corresponding conditional expected utility for player i is thus

$$(2.2) \quad U_{ix}(s) = \sum_{z > x} \rho(z | s, x) u_i(z) \quad [\forall i \in I, x \in X, s \in S].$$

If $\rho(h, s) > 0$, then s induces a conditional probability distribution μ^s on h , defined by $\mu^s(x) = \rho(x, s) / \rho(h, s) \quad \forall x \in h$, so in this case one may define the conditional expected utility for player i also under the condition that s is played and h is reached

(by play of s):

$$(2.3) \quad U_{ih}(s) = \sum_{x \in h} \mu^s(x) U_{ix}(s) \quad [\forall i \in I, h \in H, \text{ and } s \in S \text{ with } \rho(h, s) > 0].$$

Accordingly, a strategy profile s is a *Nash equilibrium (NE)* if, for every player $i \in I$, $U_i(s) \geq U_i(s \setminus s_i^?) \quad \forall s_i^? \in S_i$, where $s \setminus s_i^?$ is the strategy profile obtained from s when $s_i^?$ is substituted for s_i . Equivalently, $s \in S$ is a Nash equilibrium if, for every player i and information set $h \in H_i$ on the path of s , s_i is a best reply to s at h , i.e. $U_{ih}(s) \geq U_{ih}(s \setminus s_i^?)$ for every $s_i^? \in S_i$ such that $\rho(h, s \setminus s_i^?) > 0$. The *restriction* of a strategy s_i to (information sets in) a subgame Γ_x (with initial node x) will be written $s_i|_x$, and a strategy profile s is a *subgame perfect equilibrium (SPE)* if, for every subgame Γ_x , the subgame profile $s|_x$ constitutes a Nash equilibrium of Γ_x (Selten (1965, 1975)).

3. Probabilistic behavior hypotheses

Any ("rational" or "irrational") behavior of a player in an extensive form game can be represented by a mixed strategy. For any deterministic mode of behavior can be represented by a vector of choices, one for each of the player's information sets, i.e. by a pure strategy, and so any randomized mode of behavior is a mixed strategy. Hence, any probabilistic hypothesis that an observer may postulate about a player's behavior can be modelled as a probability measure over the player's set of mixed strategies.

However, for every (Borel) probability measure ν_j over the set M_j of mixed strategies available to player j there exists a (unique) mixed strategy $\bar{m}_j$ for j which

produces the same probability distribution over j 's pure strategies as does ν_j , viz. the strategy $\bar{m}_j \in M_j$ defined by

$$(3.1) \quad \bar{m}_j(\pi_j) = \int_{M_j} m_j(\pi_j) d\nu_j(m_j) \quad \forall \pi_j \in \Pi_j,$$

where $m_j(\pi_j)$ denotes the probability assigned by a mixed strategy m_j to a pure strategy π_j .⁴ Hence, without loss of generality, we may take any probabilistic hypothesis concerning a player's behavior to be a point in his mixed-strategy simplex.

Moreover, every mixed strategy for player j defines a probability distribution over the choices available at each of j 's (relevant) information sets — hence a behavior strategy for j . More precisely, for any $h \in H_j$, let $\Pi_j(h)$ be the (nonempty) subset of j 's pure strategies π_j that are *possible* at h , i.e. $\Pi_j(h) = \{\pi_j \in \Pi_j: \rho(h, \pi_j) > 0 \text{ for some } \pi' \in \Pi \text{ with } \pi'_j = \pi_j\}$. For any $h \in H_j$ and choice $c \in C_h$, let $\Pi_j(c) = \{\pi_j \in \Pi_j: \pi_{jh} = c\}$. Following Kuhn (1953), the *behavior* induced by a mixed strategy $m_j \in M_j$ is the behavior strategy $p_j \in S_j$ defined at each $h \in H_j$ and $c \in C_h$ by

$$(3.2) \quad p_{jh}(c) = \begin{cases} m_j[\Pi_j(h) \cap \Pi_j(c)] / m_j[\Pi_j(h)] & \text{if } m_j[\Pi_j(h)] > 0 \\ m_j[\Pi_j(c)] & \text{otherwise} \end{cases} \quad .^5$$

Hence, again without loss of generality, we may take any probabilistic

⁴ This equivalence between probability measures over mixed-strategies and mixed strategies was observed in Pearce (1984).

⁵ In games with perfect recall the behavior strategy p_j is realization equivalent with m_j , by Theorem 4 in Kuhn (1953), and the specification of p_{jh} at information sets $h \in H_j$ with $m_j[\Pi_j(h)] = 0$ is irrelevant for realization probabilities.

hypothesis concerning a player's strategic behavior to be a behavior strategy p_j for that player. Moreover, since there for every behavior strategy p_j exists a mixed strategy m_j with behavior p_j (let $m_j(\pi_j) = \prod_{h \in H_j} p_{jh}(\pi_{jh})$, cf. Lemma 3 in Kuhn (1953)), *any* behavior strategy p_j for j corresponds to some probabilistic hypothesis about his strategic behavior. Accordingly, any probabilistic hypothesis concerning all players' behavior (which treats different players' randomizations as statistically independent) can be represented simply as a behavior strategy profile $p \in S$.⁶ Such a probabilistic hypothesis p will be said to be *complete* if it assigns positive probabilities to all choices in the game, i.e. if $p(c) > 0 \ \forall c \in C$, or simply $p > 0$ for short.

Now consider a player i at one of his information sets h . Knowing that the information set h has been reached by actual play, his probabilistic hypothesis p should at least assign positive probability to this event. Let

$$(3.3) \quad P(h) = \{ p \in S ; \rho(h,p) > 0 \}$$

be the (non-empty) set of such probabilistic hypotheses, for convenience to be called *strategy priors* at h .

By construction, every strategy prior p at h defines a conditional probability distribution μ^p over the nodes in h : $\mu^p(x) = \rho(x,p)/\rho(h,p) \ \forall x \in h$. Hence, given any strategy prior $p \in P(h)$, we may define the *expected utility* for player i of playing any of his local strategies $s_{ih} \in S_{ih}$ at h simply by

$$(3.4) \quad U_{ih}(s_{ih}|p) = \sum_{x \in h} \mu^p(x) U_{ix}(p \setminus s_{ih})$$

⁶ For a lucid discussion of the subtle issue of independence vs. correlation, see Aumann (1987).

(cf. (2.3)). Evidently, the expected utility $U_{ih}(s_{ih}|p)$ is functionally independent of the player's assessment p_{ih} of his behavior at h , so it is meaningful to call a local strategy $s_{ih} \in S_{ih}$ *optimal at $h \in H_i$ under $p \in P(h)$* if

$$(3.5) \quad U_{ih}(s_{ih}|p) \geq U_{ih}(s'_{ih}|p) \quad \forall s'_{ih} \in S_{ih}.$$

Note, finally, that while any probabilistic behavior hypothesis is mathematically represented as a single behavior strategy profile, the player at the information set in question may well justify his decision by a more complex (but statistically independent) probabilistic hypothesis about the players' choice vectors. Hence, when we ask whether a certain strategy prior p is "reasonable" or not, we should ask whether or not the corresponding probability distribution over choice vectors, rather than the (reduced form) strategies in p themselves, are reasonable — in contrast to the usual practice in applications of the sequential equilibrium concept.

4. Consistency by Occam's razor

As usual, a *play* of the game is a path through the tree, from its initial node o to one of its terminal nodes $z \in Z$. Accordingly, a play τ "contains" nodes and branches from the tree. Let T denote the set of plays in the game, and let C be the union of all choice sets in the game, and, for any play $\tau \in T$, let $C(\tau)$ be the subset of choices *made* in τ (i.e. choices containing a branch in τ). For any information set $h \in H$, let $T(h) \subset T$ be the subset of plays that *reach* h (i.e. plays that contain some node in h). If, in actual play of the game, an information set $h \in H$ is reached, then the player at h knows

that actual play, τ^a , belongs to the subset $T(h)$. Hence, every play $\tau \in T(h)$ is a potential *explanation* of how the information set was reached and of how play will proceed, and the set $T(h)$ includes all such explanations at the information set h .

For any play $\tau \in T$ and strategy profile $s^* \in S$, the corresponding set of *deviations* from s^* is

$$(4.1) \quad D(\tau, s^*) = \{c \in C(\tau); s^*(c) \neq 0\}.$$

We will say that the player at h has *no evidence* for a deviation from s^* at an information set $h' \in H$ if, for every play τ that reaches h and contains a deviation at h' , there exists another play τ' that reaches h without any deviation at h' , and without any deviation that is not made also in τ . In other words: there is no evidence for a deviation at h' if each explanation which assumes a deviation at h' can be substituted by an explanation which assumes no deviation at h' and yet does not assume any additional deviation. Formally: for any information set $h \in H$ and strategy profile $s^* \in S$, let the subset $H(h, s^*) \subset H$ of *evidence free* information sets be defined by

$$(4.2) \quad H(h, s^*) = \{ h' \in H : \forall \tau \in T(h) \text{ with } C_{h'} \cap D(\tau, s^*) \neq \emptyset \exists \tau' \in T(h) \text{ with } C_{h'} \cap D(\tau', s^*) = \emptyset \text{ and } D(\tau', s^*) \subset D(\tau, s^*) \}.$$

Figure 2 below provides some examples. In diagram (a) there exists a play τ' through the information set h which contains no deviation from the indicated strategy profile s^* (marked by doubled branches). Formally, $D(\tau', s^*) = \emptyset$ and hence $H(h, s^*) = H$, i.e. the player at h has no evidence of any deviation at all. (Likewise $H(\{o\}, s^*) = H(\{x\}, s^*) = H$.) In contrast, the player at h would have evidence of a deviation at o if s^* were as indicated in diagram (b). For there every play through h (there are four

such plays) contains a deviation at the initial information set, so the player at h is sure that a deviation was made at $\{o\}$. But there is no other evidence of a deviation, so $H(h, s^*) = \{\{x\}, h\}$ in this case. (Likewise $H(\{x\}, s^*) = \{\{x\}, h\}$ but $H(\{o\}, s^*) = H$.) In the slightly more complicated situation in diagram (c), the information set h is reached if and only if player 1 and/or player 2 deviates from the indicated strategy profile. Hence, a play $\tau \in T(h)$ containing only a deviation by player 1 (or 2) is "minimal": no deviation-free play τ' reaches h . Thus neither 1's nor 2's information set is evidence-free, so we have $H(h, s^*) = \{h\}$ in this example. Note that although the information set h provides no more information about 1's action than does 2's information set, 1's information set is evidence free as viewed from 2's information set but not as viewed from h . The reason is that while 1's deviation is not part of any "irreducible" explanation of how 2's information set was reached, it is part of such an explanation of how h was reached; unlike player 2, player 3 *knows* that at least one deviation has been made.

(Figure 2)

Having considered a few examples, we now turn to some general observations about the set $H(h, s^*)$ defined in (4.1). First, the player at the information set h can have no evidence about deviations at h itself: $h \in H(h, s) \quad \forall h \in H, s \in S$. Secondly, the player at h has no evidence of any deviation if and only if h lies on the path of s^* :

$$(4.2) \quad \rho(h, s^*) > 0 \iff H(h, s^*) = H$$

(in particular, $H(\{o\}, s) = H \quad \forall s \in S$). To see this, first note that an information set h is reached by a profile s^* iff there exists some play $\tau' \in T(h)$ with $D(\tau', s^*) = \emptyset$. Hence, if

$\rho(h, s^*) > 0$, then the criterion for $h' \in H(h, s^*)$ is met by every information set $h' \in H$. Conversely, if $\rho(h, s^*) = 0$, then no $\tau' \in T(h)$ has $D(\tau', s^*) = \emptyset$ and so there exists (by finiteness of $T(h)$) some play $\tau \in T(h)$ such that $D(\tau, s^*) \neq \emptyset$ and such that $D(\tau'', s^*) \subset D(\tau, s^*)$ for no $\tau'' \neq \tau$ in $T(h)$. But then $D(\tau, s^*) \cap C_h \neq \emptyset$ for some $h' \in H$, so $h' \notin H(h, s^*)$, and thus $H(h, s^*) \neq H$.

Let $P(h, s^*) \subset P(h)$ be the subset of strategy-priors at h that agree with s^* wherever there is no evidence of a deviation from s^* , i.e.

$$(4.3) \quad P(h, s^*) = \{ p \in P(h) : p \text{ agrees with } s^* \text{ on } H(h, s^*) \}.$$

Any prior p in the subset $P(h, s^*)$ will be said to be *consistent* with s^* .

Proposition 4.1: $P(h, s^*) \neq \emptyset \quad \forall h \in H, \forall s^* \in S$.

Proof: There exists at least one τ in the finite set $T(h)$ such that for no $\tau' \in T(h)$ is $D(\tau', s^*)$ a proper subset of $D(\tau, s^*)$. Let τ be such, and let x be the node in h which is contained in τ . Let $p \in S$ agree with s^* at all information sets $h' \in H(h, s^*)$. At each $h' \notin H(h, s^*)$ (if such exists), let p assign positive probabilities to all choices in C_h . Then $\rho(x, p) > 0$. For suppose $h' \in H(h, s^*)$ and $c \in C_h \cap C(\tau)$. If $c \in D(\tau, s^*)$, then, by definition of $H(h, s^*)$, $\exists \tau' \in T(h)$ which has $D(\tau', s^*)$ as a proper subset of $D(\tau, s^*)$, contradicting the construction of τ . Hence $c \notin D(\tau, s^*)$, so $p(c) = s^*(c) > 0$. Thus p assigns positive probabilities to all choices in $C(\tau)$, and hence $\rho(x, p) > 0$. Since p agrees with s^* on $H(h, s^*)$ by construction, we have established $p \in P(h, s^*)$. $\square$

It follows from (4.2) that s^* is the only consistent prior at an information set h if and only if h lies on the path of s^* :

$$(4.4) \quad \rho(h, s^*) > 0 \iff P(h, s^*) = \{s^*\}.$$

Kreps and Ramey (1987) suggest the further consistency requirement that if an information set h of a player lies off the path of a profile s^* , but within reach for him by a unilateral deviation, then he should not hypothesize at h that also other players have deviated. And this extension of (4.4) is indeed implied by the present consistency criterion. For, by perfect recall, the player will then *know* at his information set h that he has deviated, and hence he has no evidence that also others have deviated. Formally:

Proposition 4.2: If $\rho(h, s^*) = 0$ and $\rho(h, s^* \setminus s_i) > 0$ for some $i \in I$, $h \in H_i$ and $s_i \in S_i$, then $\bigcup_{j \neq i} H_j \subset H(h, s^*)$.

Proof: Suppose $h' \in H_j$ for some $j \neq i$, and assume $\tau \in T(h)$ and $c \in C_h \cap D(\tau, s^*)$. By hypothesis, $\exists \tau' \in T(h)$ with $D(\tau', s^*) \subset C_i$, where C_i is the union of i 's information sets. Hence, $C_h \cap D(\tau', s^*) = \emptyset$. By perfect recall at $h \in H_i$, $D(\tau', s^*) \subset D(\tau, s^*)$. Hence, $h' \in H(h, s^*)$. $\square$

Note finally that, just as in Kreps and Wilson (1982), consistency is here defined without any reference to payoffs. Hence, consistency of a player's strategy prior involves no consideration of other players' "rationality" or "intentions".

5. Optimality under approximately consistent priors

We are now in a position to formalize the earlier mentioned criterion for a strategy profile s^* to qualify as "self-enforcing": for any degree $\epsilon > 0$ of accuracy, and at each information set h , the prescribed local strategy s_{ih}^* should be optimal under some probabilistic hypothesis $p \in P(h)$ which lies within distance ϵ from the set $P(h, s^*)$ of probabilistic hypotheses q which agree with s^* at all evidence free information sets $h' \in H(h, s^*)$. Formally, we require $s^* \in E$, where

$$(5.1) \quad E = \{s^* \in S: \forall i \in I, \forall \epsilon > 0, \forall h \in H_i, \exists p \in P(h) \text{ meeting [A] and [B]}\}.$$

$$[A] \quad \|p - q\| < \epsilon \text{ for some } q \in P(h, s^*),$$

$$[B] \quad s_{ih}^* \text{ is optimal at } h \text{ under } p.^7$$

Note, in particular, that a player's prior p may differ slightly from s^* even at his own information sets. Indeed, any theory for strategic interaction not only has to allow for the possibility of others' "irrationality" or "mistakes"; at an information set where a player knows that he has earlier made a mistake (say, a strictly dominated move), we cannot reasonably force him to believe that the probability of a future mistake is zero.

In addition to the requirements in the above definition of the set E , one may require of a "cautious" and "consistent" player that all his local strategies should be optimal under one and the same complete probabilistic hypothesis. Indeed, this is one way to accommodate a theory of "self-enforcement" to the necessity of allowing for

⁷ The norm is here taken to be the max norm:

$$\|p - q\| = \max\{|p_{ih}(c) - q_{ih}(c)| : i \in I, h \in H_i \text{ and } c \in C_h\}.$$

"irrationality": while certain behaviors may contradict a particular strategy profile s^* , no behavior contradicts a probabilistic hypothesis which assigns positive probabilities to all choices in the game (and hence to all nodes and information sets). Formally, this extended requirement amounts to the condition $s^* \in E^*$, where $E^* \subseteq E$ is defined by

$$(5.2) \quad E^* = \{s^* \in S: \forall i \in I, \forall \epsilon > 0 \exists p > 0 \text{ meeting } [A] \text{ and } [B] \forall h \in H_i\}.$$

While intermediate "cautiousness" and "internal consistency" conditions may be natural in many games — requiring positivity only for some components of priors and consistency only on some subcollections of a player's information sets — the present study will focus mainly on the properties of strategy profiles in the subset E^* . The central general observation to be made is that every (extensive-form, "trembling hand") perfect equilibrium belongs to E^* and that every profile in the superset E is subgame perfect.⁸

For this purpose, let PE denote the (nonempty) set of extensive-form perfect equilibria of the game under consideration. Using Selten's (1975) "substitute sequence" characterization of such equilibria, one readily obtains

Proposition 5.1: $PE \subseteq E^*$.

Proof: Suppose $s^* \in PE$. By Theorem 7 in Selten (1975), there exists a sequence (s^t) of behavior strategy profiles $s^t > 0$ such that $s^t \rightarrow s^*$ (in the max-norm used above) and such that each player's local strategy is a best local reply, at the information set in

⁸ Hence, also any intermediate self-enforcement criterion produces a refinement of subgame perfection and a coarsening of extensive-form "trembling hand" perfection.

question, to each profile s^t . Hence, $\forall \epsilon > 0 \exists p > 0$ meeting [B] $\forall i \in I, h \in H_i$, and such that $\|p - s^*\| < \epsilon/2$. However, $\forall \epsilon > 0 \forall h \in H \exists q \in P(h, s^*)$ such that $\|q - s^*\| < \epsilon/2$. To see this, let $q' \in P(h, s^*)$ ($P(h, s^*) \neq \emptyset$ by Prop. 4.1). Then q' and s^* agree on $H(h, s^*)$. For each $h' \in H(h, s^*)$, let $q_{ih'} = q'_{ih'}$, ($= s^*_{ih'}$), while for each $h' \notin H(h, s^*)$, let $q_{ih'} = \lambda \cdot q'_{ih'} + (1-\lambda)s^*_{ih'}$, for $\lambda = \epsilon/2$. Then $q \in P(h', s^*)$ and $\|q - s^*\| < \epsilon/2$. The triangle inequality ensures $\|q - p\| < \epsilon$, i.e. p also meets [A]. $\square$

The claim that each profile s^* in the superset E constitutes a Nash equilibrium can be established in two steps as follows. First, if a profile s^* is not a Nash equilibrium, then there is some player i , information set $h \in H_i$ and strategy s_i such that $s^* \setminus s_i$ reaches h with positive probability, s_i agrees with s^*_i at all information sets $h' \neq h$ which do not come before h , and play of s_{ih} at h yields a higher (conditional) payoff to i than play of s^*_{ih} :

Lemma 5.1: If $s^* \notin NE$, then there exist $i \in I$, $h \in H_i$ and $s_i \in S_i$ such that $s^* \setminus s_i \in P(h, s^*)$ and $U_{ih}(s_{ih} | s^* \setminus s_i) > U_{ih}(s^*_{ih} | s^* \setminus s_i)$.

Proof: If $s^* \notin NE$, then $\exists i \in I, h \in H_i$ and $s_i \in S_i$ such that $\rho(h, s^*) > 0$, $\rho(h, s^* \setminus s_i) > 0$ and

$$(1) \quad U_{ih}(s^* \setminus s_i) > U_{ih}(s^*).$$

For any $h', h'' \in H$, write $h' \leq h''$ iff $x' \leq x''$ for some $x' \in h'$ and $x'' \in h''$. For the information set h in (1), let $A_i(h) = \{h' \in H_i : h \leq h'\}$, and let $F_i(h, s_i) = \{h' \in A_i(h) : \rho(h', s^* \setminus s_i) > 0 \text{ and } s_{ih'} \neq s^*_{ih'}\}$. Without loss of generality, we may assume s_i to agree with s^*_i outside $A_i(h)$, and by (1) $F_i(h, s_i)$ is nonempty. Moreover, $F_i(h, s_i)$ is a finite set, partially ordered by $\leq$, so $\exists h' \in F_i(h, s_i)$ such that $h' < h''$ for no $h'' \in F_i(h, s_i)$. By definition, $p = s^* \setminus s_i \in P(h')$. If

$$(2) \quad U_{ih'}(s_{ih'} | p) > U_{ih'}(s^*_{ih'} | p),$$

then we may, without loss of generality, take p to agree with s_i^* on $A_i(h')$. All $p' = s^* \setminus s_i' \in P(h')$ induce the same conditional probabilities $\mu(x)$ on h' as does p , by perfect recall, and at least one such p' belongs to $P(h', s^*)$, again by perfect recall. Hence, without loss of generality, we may assume $p \in P(h, s^*)$, and so the claim of the Lemma is established for s_i at h' if (2) holds.

If instead (2) is not met, then let $s_i^1 = s_i \setminus s_{ih}^*$. Clearly $s^* \setminus s_i^1 \in P(h)$ and

$$(3) \quad U_{ih}(s^* \setminus s_i^1) \geq U_{ih}(s^* \setminus s_i) > U_{ih}(s^*).$$

Hence, $F_i(h, s_i^1)$ is a nonempty proper subset of $F_i(h, s_i)$. Like before, we may choose a $h' \in F(h, s_i^1)$ such that $h' < h''$ for no $h'' \in F(h, s_i^1)$. Again $p = s^* \setminus s_i^1 \in P(h')$. If

$$(4) \quad U_{ih}(s_i^1, |p) > U_{ih}(s_{ih}^*, |p),$$

then the claim of the Lemma is established for an appropriate modification of p , as shown in the context of (2) above. Otherwise, let $s_i^2 = s_i^1 \setminus s_{ih}^*$, etc. The set $A_i(h)$ being finite, this process stops at some modified strategy s_i^t , obtained from s_i by t substitutions of local strategies s_{ih} for s_{ih}^* , hence establishing the claim of the lemma for s_i^t at the corresponding information set h' . $\square$

As a consequence of this lemma, a strategy profile s^* which is not a Nash equilibrium violates the local optimality condition [B] at some information set h under some prior $p = s^* \setminus s_i \in P(h, s^*)$. But by perfect recall, all priors in $P(h, s^*)$ result in the same conditional payoff to player i , so, by (uniform) continuity of the player's payoff function, [B] is violated for all priors near the set $P(h, s^*)$, and hence $s^* \notin E$, establishing the second step in the proof of our initial claim:

Proposition 5.2: $E \subset NE$.

Proof: Suppose $s^* \notin NE$. By Lemma 5.1, $\exists i \in I, h \in H_i$ and $s_i \in S_i$ such that $p = s^* \setminus s_i \in P(h, s^*)$ and $f(p) > 0$, where $f: S \rightarrow \mathbb{R}$ is defined by $f(p) = U_{ih}(s_{ih} | p) - U_{ih}(s_{ih}^* | p)$. For any $q \in P(h)$ we have $\mu^q(x) = \mu^p(x) \forall x \in h$ by perfect recall. Moreover, any $q \in P(h, s^*)$ agrees with s^* on $A_i(h) = \{h' \in H_i: h \leq h'\}$, and on each set H_j , for $j \neq i$, by Prop. 4.2. Hence, $f(q) = f(p) \forall q \in P(h, s^*)$. Since S is compact and f continuous, f is in fact uniformly continuous, so $\exists \epsilon > 0$ such that $f(q) > 0$ for all $q \in S$ such that $\|q - p\| < \epsilon$ for some $p' \in P(h, s^*)$. Thus, for this ϵ there exists no $q \in P(h)$ meeting [A] and [B]. $\square$

So far, the following chain for arbitrary finite games (with perfect recall) has been established:

$$(5.3) \quad PE \subset E^* \subset E \subset NE.$$

Another property of the present self-enforcement criterion is that it is *backward inductive* in the sense that every solution of a game induces a solution on every subgame. Let E^* be the solution set for some game Γ which has a proper subgame Γ_x , and let E_x^* be the solution set for the subgame:

Proposition 5.3: If $s^* \in E^*$, then $s^*|_x \in E_x^*$.

Proof: To see this, suppose $s^* \in E^*$, and consider an information set $h \in H_i$ in the subgame Γ_x . By hypothesis, $\exists p > 0$ under which every local strategy s_{ih}^* in Γ is optimal, and such that p is near to some $q \in P(h, s^*)$. But then s_{ih}^* is also optimal under the subgame restriction $p_x > 0$ of p . It hence only remains to show that q_x is

a consistent prior in Γ_x , i.e. that $q_x \in P_x(h, s_x^*)$. But the player at h has evidence in Γ of a deviation at an information set h' in Γ_x iff he has evidence in Γ_x of a deviation at h' , since any plays $\tau, \tau' \in T(h)$ in Γ coincide outside Γ_x . $\square$

Propositions 5.2 and 5.3 together imply that each strategy profile s^* in the superset E is a subgame perfect equilibrium. In fact, the requirement $s^* \in E$ is a *refinement* of subgame perfection, as is seen in the famous example by Selten (1975), given in Figure 3 below.

(Figure 3)

This game has no proper subgame, and each of the three players has only one information set. In his informal discussion, Selten disqualifies the SPE s^* indicated in the diagram by the argument that 2's choice should be guided by his conditional payoff—expectation at his information set (op.cit. p.34). Selten then shows how his model of "trembles" in fact rejects s^* . Also the weaker of the present self-enforcement criteria rejects s^* , but on the grounds of Selten's informal rejection. For optimality under an approximately consistent prior demands that 2's strategy be optimal (at least) if 2 were to believe that 3 plays approximately according to s^* . But then player 2 should deviate.

In sum: writing SPE for the set of subgame perfect equilibria of a given game, we have refined the chain in (5.3), applying to every (finite) extensive-form game (with perfect recall), to

$$(5.4) \quad PE \subset E^* \subset E \subset SPE.$$

6. A comparison with sequential and perfect equilibrium

Like the sets E^* and E , the set SE of sequential equilibria of a game Γ is a superset to PE (the set of extensive-form perfect equilibria of Γ) and a subset of SPE (the set of subgame perfect equilibria of Γ). In this section, examples will be given in which the sets E^* and E differ from the set SE , and it will be shown how the present self-enforcement criteria overcome some of the problematic aspects of the sequential equilibrium concept.

Kreps and Wilson (1982) define a *system of beliefs* as a function $\mu: X \rightarrow [0,1]$ such that $\forall h \in H \sum_{x \in h} \mu(x) = 1$, and they define an *assessment* as a pair (μ, s) where μ is a system of beliefs and s a strategy profile. An assessment (μ, s) is *sequentially rational* if, for every $i \in I$ and $h \in H_i$,

$$(6.1) \quad \sum_{x \in h} \mu(x) U_{ix}(s) \geq \sum_{x \in h} \mu(x) U_{ix}(s|s_i^?) \quad \forall s_i^? \in S_i.$$

They call an assessment (μ, s) *consistent* if $(\mu^t, s^t) \rightarrow (\mu, s)$ for some sequence $\langle (\mu^t, s^t) \rangle$ of assessments in which every s^t is completely mixed and induces μ^t , i.e. $\forall h \in H : \mu^t(x) = \rho(x, s^t) / \rho(h, s^t) \quad \forall x \in h$. A *sequential equilibrium* is an assessment which is both sequentially rational and consistent. As is usual, a *profile* s^* will be called a sequential equilibrium, $s^* \in SE$, if there exists some system of beliefs μ such that the assessment (μ, s^*) is sequentially rational and consistent.

Kreps and Wilson (1982) call a system of beliefs μ *structurally consistent* if there for every information set h exists some strategy profile s^h which induces μ on h , i.e. such that $\rho(h, s^h) > 0$ and $\mu(x) = \rho(x, s^h) / \rho(h, s^h) \quad \forall x \in h$. As was shown in

Section 3 above, the requirement of structural consistency is indeed natural, since any (statistically independent) probabilistic hypotheses about strategies can be represented as a single (behavior) strategy profile. However, as shown by Kreps and Ramey (1987), the consistency requirement in the definition of sequential equilibrium, "SE-consistency", does not imply structural consistency. In fact, they provide an example in which *all* sequential equilibria (μ, s^*) are structurally inconsistent, see Figure 4.

(Figure 4)

First, they note that every Nash equilibrium in this game has players 1 and 2 move left, as indicated in the diagram. Moreover, they note that SE-consistency implies $\mu(x)=0$, i.e. player 3 has to assign zero probability to the event that both 1 and 2 deviate. In order to render 1's and 2's strategies sequentially rational, 3 has to mix his strategy (sufficiently), which is sequentially rational for 3 to do only if he assigns positive probabilities to both his other nodes (an assignment compatible with SE-consistency). However, such a conditional probability distribution is clearly structurally inconsistent.⁹ Hence, this game has no structurally consistent sequential equilibrium. Note, however, that all Nash equilibria s^* of this game belong to set E^* . To see this, let 3's prior assign (approximately or exactly) equal probabilities to 1's and 2's deviations — such a strategy prior p is positive, belongs to the set $P(h, s^*)$ of consistent priors at 3's information set h , and renders all 3's strategies optimal. Moreover, for 3's strategy (sufficiently) mixed, it is optimal for each of players 1 and 2 to play L .

⁹ As Kreps and Ramey (1987) point out, such a belief presumes a certain (negative) correlation between 1's and 2's deviations; 1 deviates if and only if 2 does not deviate.

We noted above that SE-consistency requires $\mu(x)=0$ in the game in Figure 4. The variation of Krep's and Ramey's (1987) example given in Figure 5 (a) below illustrates how this consistency requirement may lead to rejection of certain strategy profiles in the set E^* .

(Figure 5)

The indicated strategy profile, $s^*=(L,L,L)$, is a Nash equilibrium. Moreover, it meets the present self-enforcement criterion, for any payoff value $v>0$. For example, if 3 assigns the same conditional probability δ to 1's and 2's deviations, even that his information set is reached, then L is his unique optimal choice if and only if $\delta > 2/(v+2)$. Hence, for every $v>0$ there exist $\delta \in (0,1)$ for which $s_3=L$ is optimal, in which case $s^* \in E^*$. However, for no value of v is s^* a sequential equilibrium. In fact, the outcome $"(2,2,0)"$ is incompatible with sequential equilibrium. For if a profile has players 1 and 2 play L , then player 3 has to assign probability zero to his right-most node, by SE-consistency, and hence he should play R , by sequential rationality. But then it is not sequentially rational for players 1 and 2 to play L . A fortiori, the strategy profile $s^*=(L,L,L)$ and the outcome $(2,2,0)$ is incompatible with ("trembling hand") perfect equilibrium (extensive-form and normal-form perfection being equivalent in this game). Note that the reason why the outcome $"(2,2,0)"$ is imperfect is not that it is non-robust with respect to "trembles" per se, but that it is non-robust with respect to infinitesimal trembles. To see this, just note that, for any estimate $\delta \in (0,1)$ that player 3 may make of (independent and equal) tremble probabilities by players 1 and 2, there exist payoff values v for which $s_3^*=L$ is optimal for player 3, viz. for all $v \geq 2(1-\delta)/\delta$, see Figure 5 (b). For every $v>0$, let

$\delta(v) \in [0,1]$ be 3's probabilistic hypothesis. Unless we can rule out a priori the possibility that $\delta(v) > 2/(v+2)$ for some $v > 0$, we are not justified to discard the Nash equilibrium s^* .

The example in Figure 5 above showed that the set E^* is not always a subset of the set SE . The following example, due to van Damme (1987), establishes that SE is not always a subset even of the superset E , see Figure 6.

(Figure 6)

In this variation of Kreps' and Ramey's example, the strategy profile $s^* = (L, L, M)$ is a sequential equilibrium: let μ assign conditional probabilities $(\frac{1}{2}, \frac{1}{2}, 0)$ to 3's nodes. However, the strategies (L, L) for players 1 and 2 are optimal only if M is chosen with a probability of at least $\frac{1}{2}$, and M is optimal for player 3 *only* under the structurally inconsistent belief $\mu = (\frac{1}{2}, \frac{1}{2}, 0)$ at his information set. Hence, under no strategy prior is it optimal for player 3 to play M . And if 3 plays M with a probability below $\frac{1}{2}$, then players 1 and 2 optimally play R . Thus, while the outcome $(1, 1, 0)$ is compatible with sequential equilibrium, it is incompatible even with $s^* \in E$. (In contrast to the earlier examples, this example is non-generic.)

Another example in Kreps and Ramey (1987) shows that structural consistency may conflict with sequential rationality, see Figure 7 below. In this game, nature has the first move, and gives the first personal move to either player 1 or 3, with equal probabilities. The selected player can either end the game (Q), or give the move to player 2 (P). In the latter case, player 2 will not know who gave him the move, and he can then either end the game, in two alternative ways (T or B), or (P) give the move to the player (1 or 3) who did not yet move. Hence, if a strategy profile has players 1

and 3 choose Q, then the information set of player 2 is not reached. As Kreps and Ramey show, this game has only one NE outcome, and in this outcome players 1 and 3 do select Q with probability 1, so 2's information set is indeed not reached in any Nash equilibrium. To support this profile as a sequential equilibrium, player 2 must select T and B with probabilities not exceeding $1/3$. By sequential rationality, he must thus assign positive priors to both nodes in his information set. Hence, he must assume that both players 1 and 3 play P with positive probability, i.e. that they *both* deviate from equilibrium play. However, sequential rationality also requires that 2's move be a best reply to equilibrium moves by players 1 and 3, i.e. player 2 should hypothesize that both 1 and 2 play Q, a contradiction.

(Figure 7)

In contrast to sequential rationality, the present self-enforcement criterion does not require that players always presume equilibrium play at all "later" information sets. In Figure 7, neither the first nor the third player's information set is evidence-free, as seen from the information set of player 2. Hence, the Nash equilibrium $s^*=(Q,P,Q)$ belongs to E^* . (It is optimal for 2 to play P if he, for example, assigns equal probabilities to 1's and 3's deviations.)

Let us finally consider the example in Figure 8. In this game, player 1 first decides whether or not to take an outside option (T). If he does not take this outside option, then he has two options (A and B), both leading to 2's information set. Precisely one of 2's choices (D) leads to 3's information set, so player(s) 2 and 3 possess exactly the same (imperfect) information about 1's move. Such situations of shared imperfect information are common in economic models. The indicated strategy

profile $s^* = (T, C, E)$ is a Nash equilibrium, giving the "potential entrant" (player 1) a sure payoff of 1, instead of a payoff in the set $\{0, 2, 3\}$ if he would "enter" and interact with the "incumbants" 2 and 3.

(Figure 8)

Now consider any strategy profile which has player 1 play T, and any system of beliefs μ . SE-consistency of any such assessment (μ, s^*) imposes no restriction on the conditional probability $\mu(x)$, but, irrespective of which value $\mu(x)$ takes, $\mu(x')$ must take precisely the same value. To see this, just note that, for any completely mixed strategy profile s^t , the realization probabilities in 3's information set are proportional to those in 2's, since one and the same choice by 2 is made in all plays through 3's information set. Hence, SE-consistency implies $\mu(x') = \mu(x)$, as claimed. It follows that the Nash equilibrium $s^* = (T, C, E)$ is not a sequential equilibrium, since $s_2^* = C$ is sequentially rational for 2 if and only if $\mu(x) \geq 2/3$, while $s_3^* = E$ is sequentially rational for 3 if and only if $\mu(x') \leq 1/2$.¹⁰ In contrast, we evidently have $s^* \in E^*$ since this self-enforcement criterion does not require players 2 and 3 to hold identical beliefs.

Also this example shows that E^* is not necessarily a subset of SE, but this time because the present self-enforcement criterion requires less interpersonal consistency in beliefs than does sequential equilibrium. (Also Kohlberg's and Mertens' (1986) stability criterion requires full interpersonal consistency in beliefs, and it rejects the Nash equilibrium outcome (1,5,5) in the game in Figure 8 since all strategy profiles in a stable set are normal-form perfect, a condition which is equivalent to

¹⁰ In fact, it is easily verified that the outcome "(1,5,5)" is incompatible with sequential equilibrium.

extensive-form perfection in this example).¹¹

7. Cautious play and weak dominance

It is well-known that, unlike normal form perfection, extensive-form perfection does not exclude play of weakly dominated strategies. The set E^* being a superset of this set PE , this "weakness" of extensive-form perfection is shared by the present self-enforcement criterion. However, by means of some simple concluding examples, we will question whether it is indeed a sound practice to reject all weakly dominated strategies.

In the game in Figure 10 (a), the "unreasonable" strategy profile $s^* = (L, L)$ is not a member even of the superset E_a , while the same profile indeed is a member of the set E_b associated with the similar game in Figure 10 (b).

(Figure 9)

The two games in question clearly have the same normal form, and some game theorists would argue that they are strategically equivalent. The reason why the profile s^* is rejected in game (a) is evidently that if player 2 is called upon to make a move, then 2 *knows* that 1 has deviated, and 2's unique optimal choice is then to deviate too. In game (b), on the other hand, player 2 *can* believe that 1 did not deviate, a belief that does not appear to be "irrational" per se and which justifies 2's

¹¹ For a more detailed discussion of interpersonal consistency in beliefs, see Weibull (1989a). Fudenberg, Kreps and Levine (1988) develop a normal-form refinement of Nash equilibrium, c-perfection, which does allow for interpersonal differences in beliefs, as well as for correlation. The above profile s^* is c-perfect.

choice to play L . However, 2's play of L could be said to be insufficiently cautious, and the weakly dominated strategy L is indeed rejected by the stricter criterion $s^* \in E^*$.

Consider now the game in Figure 10 (a). The only weakly undominated strategy available to player 1 is Ll , resulting in the outcome $(2,0)$. However, refutation of the subgame perfect outcome $(2,2)$ is justified only if player 1 at his first decision node assigns four times higher mistake-probability to 2's move than to his own future move. (By changing the payoffs, this critical ratio between subjectively assigned mistake probabilities can be changed arbitrarily.) While the strategy profile $s^* = (Rl, l)$ is refuted by normal-form perfection, it is accepted by extensive-form perfection, and accordingly $s^* \in E^*$.

(Figure 10)

Adding an initial move to player 2, as in Figure 10 (b), makes the rejection of weakly dominated strategies even more effectful and dubious: only the outcome $(0,1)$ in that game is compatible with undominated Nash equilibrium, and hence with normal-form perfection, while the extensive-form perfect outcome $(2,2)$ appears "self-enforcing".

8. Directions for further research

The self-enforcement criterion $s^* \in E^*$ discussed in this study requires that each player should be willing to play his strategy *if* the tested strategy profile would (at least approximately) encode his beliefs about how the game is being played. A

further question thus is whether a "rational" player indeed should hold such beliefs. Hence, a natural extension of the present approach is to elaborate players' expectation formation with respect to intended versus unintended deviations.¹²

Another avenue for further research is to extend the present approach to infinite (action) games in extensive form.

¹² See van Damme (1989) for a discussion of stability and forward induction, and see Weibull (1988) for a sketch of how forward induction reasoning can be incorporated in an approach such as the present.

REFERENCES

- Aumann R. (1987), "Correlated equilibrium as an expression of Bayesian rationality", *Econometrica* 55, 1–18.
- Basu K. (1990), "On the non-existence of a rationality definition for extensive games", *International Journal of Game Theory* (forthcoming).
- Bernheim D. (1984), "Rationalizable strategic behavior", *Econometrica* 52, 1007–1029.
- Bicchieri C. (1989), "Self-refuting theories of strategic interaction: a paradox of common knowledge", *Erkenntnis* 30, pp.69–85.
- Binmore K. (1987), "Modeling rational players: part I", *Economics and Philosophy* 3, pp.179–214.
- van Damme E. (1987), *Stability and Perfection of Nash Equilibria*, Springer Verlag, Berlin.
- van Damme E. (1989), "Stable equilibria and forward induction", *Journal of Economic Theory* 48, pp.476–496.
- Fudenberg D. and J. Tirole (1989), "Perfect Bayesian and sequential equilibrium", mimeo.
- Fudenberg D., Kreps D. and D. Levine (1988), "On the robustness of equilibrium refinements", *Journal of Economic Theory* 44, 351–380.
- Harsanyi J. and R. Selten, *A General Theory of Equilibrium Selection in Games*, M.I.T. Press, 1988.
- Kohlberg E. and J.-F. Mertens (1986), "On the strategic stability of equilibria", *Econometrica* 54, 1003–1037.
- Kreps D. and G. Ramey (1987), "Structural consistency, consistency and sequential rationality", *Econometrica* 55, 1331–1348.
- Kreps D. and R. Wilson (1982), "Sequential equilibria", *Econometrica* 50, 863–894.
- Kuhn H. (1953), "Extensive games and the problem of information", in H. Kuhn and A. Tucker (eds.), *Contributions to the Theory of Games*, Princeton Univ. Press, Vol.2.
- von Neumann J. and O. Morgenstern (1947), *Theory of Games and Economic Behavior*, 2nd ed., Princeton University Press.
- Pearce D. (1984), "Rationalizable strategic behavior and the problem of perfection", *Econometrica* 52, 1029–1050.

- Rosenthal R.W. (1981), "Games of perfect information, predatory pricing and the chain store paradox", *Journal of Economic Theory* 25, 92–100.
- Rubinstein A. (1988), "Comments on the interpretation of game theory", Discussion Paper No. TE/88/181, London School of Economics.
- Selten R. (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, Teil I", *Zeitschrift für die gesamte Staatswissenschaft* 12, pp. 301–324 and 667–689.
- Selten R. (1975), "Reexamination of the perfectness concept for equilibrium points in extensive games", *Int. Journal of Game Theory* 4, 25–55.
- Weibull J.W. (1988), "Refinements of subgame perfection – without 'trembles'", Bonn University, Discussion Paper No. A–146.
- Weibull J.W. (1989a), "Interpersonal consistency of beliefs in extensive-form games", Institute for International Economic Studies, Stockholm University, Seminar Paper No. 433.
- Weibull J.W. (1989b), "A new self-enforcement criterion for extensive-form games", Institute for International Economic Studies, Stockholm University, Seminar Paper No. 445.

Figure 1

(a)

(b)

(c)

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Figure 7

Figure 8

(a)

(b)

Figure 9

(a)

(b)

Figure 10