

Subjective Equilibrium in Repeated Games

Author(s): Ehud Kalai and Ehud Lehrer

Source: *Econometrica*, Vol. 61, No. 5 (Sep., 1993), pp. 1231-1240

Published by: [The Econometric Society](#)

Stable URL: <http://www.jstor.org/stable/2951500>

Accessed: 21/08/2011 06:07

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The Econometric Society is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*.

SUBJECTIVE EQUILIBRIUM IN REPEATED GAMES¹

BY EHUD KALAI AND EHUD LEHRER

1. INTRODUCTION

A NASH EQUILIBRIUM, of an n -person infinitely repeated game with discounting, is a vector of behavior strategies, $f = (f_1, f_2, \dots, f_n)$, with each player's strategy, f_i , being a best response to his opponents' strategies, f_{-i} . One may think of it as an objective notion since any two players attribute to a third player, say, k , the same, and the correct strategy, f_k .

A *subjectively rational* player, on the other hand, chooses his strategy f_i to be a best response to his individual beliefs about his opponents' strategies, $f^i = (f_1^i, f_2^i, \dots, f_n^i)$. Except for knowing his own strategy, $f_i^i = f_i$, his assessments of opponents are not necessarily correct, nor do they coincide with other players' assessments. For example, all three strategies, f_3^1 , f_3^2 , and f_3 , may be different when players one and two disagree about the strategy of player three and both are wrong.

Nevertheless, a vector of subjectively rational strategies f can be in equilibrium if the play it induces coincides with the plays induced by the beliefs of each player i , as described by his belief vector, f^i . For instance, if both f and f^i consist of pure strategies and generate the same single play path, then player i 's belief, even if wrong off the play path, can only be reinforced as the game progresses. Similarly, randomizing strategies f and f^i can be *realization equivalent*, in a probabilistic sense, if the distributions they induce on future play paths coincide. Again, since the distributions induced by f and f^i may differ only after histories that have zero probability, any statistical inference conducted by player i can only reinforce his belief that the vector f^i is being played.

With the above interpretation in mind, we think of a *subjective equilibrium* as being stable with respect to learning and optimization. Players placed at such an equilibrium will not alter their beliefs and will have no incentive to alter their strategies.

Notions of subjective equilibrium are not new in economics and game theory (see Battigalli et al. (1992) for a survey). Von Hayek (1937) already discussed the differences between subjective and objective knowledge. His test for equilibrium was "whether the individual subjective sets of data correspond to the objective data, and whether in consequence the expectations in which plans were based are born out by the facts." Hahn (1973) assumed that agents maximize their utility relative to their subjective theories about the future evolution of the economy. He defined a *conjectural equilibrium* as a situation where the signals generated by the economy do not alter the agents' individual theories, nor do they induce them to change their policies. Battigalli (1987) and Battigalli and Guaitoli (1988) formalized and studied the game theoretic version of Hahn's conjectural equilibrium. Rubinstein and Wolinsky (1990) defined rationalizable versions of conjectural equilibrium.

Fudenberg and Levine (1993b) introduced a notion of *self-confirming equilibrium* defined for finite extensive form games. A player in such a game chooses a strategy to maximize his expected payoff given his subjective beliefs about opponents' strategies. These beliefs allow the possibility that the opponents' strategies are correlated, and being defined for general extensive games this notion allows for imperfect information

¹This paper was previously titled "Private-Beliefs Equilibrium." The authors wish to thank the referees of this journal for suggesting substantial improvements, particularly in the proof of the main result. This research was partly supported by Grants Nos. SES-9011790 and SES-9022305 from the National Science Foundation, Economics Program.

the players obtain and use to confirm their subjective beliefs. Fudenberg and Levine (1993a) motivated this notion in a model of overlapping generations where players of different ages are randomly matched to play a fixed extensive form game. Considering steady state learning equilibrium of such a system, they showed that as the life length of the players approaches infinity, the group's mixed strategies must approach a self-confirming equilibrium of the underlying extensive game.

The notion of subjective equilibrium studied in this paper emerges in a learning model of n players engaged in a fixed infinitely repeated perfect-monitoring game. A player in this model chooses a strategy to maximize expected payoff, given his subjective beliefs about the strategies of each of his opponents. Kalai and Lehrer (1993a) showed, under an assumption of compatibility of the beliefs with the truth, that after sufficiently long time the players must play a subjective ε -equilibrium for arbitrarily small ε . At such an equilibrium the beliefs each player holds about the future play of the game essentially coincide with the actual play.

The result just stated suggests that there exist dynamic processes of learning that lead individual utility maximizers to subjective equilibrium. Also, while subjective equilibrium is reached in the limit of such processes, only ε -subjective equilibrium is attained in finite times. For these reasons, it is important to study the group behavior induced by these types of equilibria, and to compare it to the better-known objective counterparts, Nash and ε -Nash equilibrium. The main contribution of the paper is in identifying conditions under which near coincidence is obtained.

It is well known that even in one-player games, i.e., decision problems, there is a serious discrepancy between subjective optimization and true (objective, Nash) optimization. A multi-arm bandit player (see Whittle (1982)) can be engaged repeatedly in one activity whose payoff is lower than the expected payoff of a competing alternative. The player's subjective beliefs assign the activity used a correct payoff distribution but assign the competing unused activity a false low payoff distribution. In such a situation, his assessments are reinforced and he never finds out that he is wrong off the play path and that his play is suboptimal. In other words, he does not follow a Nash equilibrium of the complete information one person game. The current paper, however, will rule out such situations by assuming that players know their correct payoff distributions and uncertainty is restricted to be strategic, i.e., concerning opponents' strategies.

But even under strategic uncertainty alone there are serious discrepancies between behavior induced by Nash and behavior induced by subjective equilibrium. Revealing examples of extensive games exhibiting this subtle phenomenon were described by Fudenberg and Kreps (1988), and Fudenberg and Levine (1993b). In order to rule out such examples, the current paper assumes that the infinitely repeated game is played with perfect monitoring (of players' actions). This assumption, together with the earlier ones, suffices to close the gap between the behavior of Nash and of subjective equilibrium.

When we restrict ourselves to Nash equilibrium and subjective equilibrium, ignoring their approximated ε versions, it is easy to see, under the conditions stated above, that the two notions induce identical behavior patterns. Starting with a subjective equilibrium, one modifies the strategies of all players as follows: (1) on the support of the original play paths no modification is done; (2) in subgames following a unilateral deviation (from the support of his strategy) by player i , all players will play according to player i 's subjective beliefs; (3) in subgames following multi-player deviation, any vector of strategies can be assigned. Following such a modification, we do not change the distribution on future play paths, since we create the incentives for each player to continue playing according to his original strategy. Thus, we end up with a Nash equilibrium which keeps the same exact distribution on future play paths as the original subjective equilibrium. In an earlier version of our learning paper (Kalai and Lehrer (1990)) we used these arguments to prove the convergence of behavior to ε -Nash equilibrium behavior in repeated two person games.

When dealing, however, with an n -person game and with subjective ε -equilibrium, where the subjective beliefs are only accurate up to ε , this simple technique does not work. One reason is that there is no more unanimity in the players' subjective minds as to whether player i deviated or not. Indeed, we can no longer prove that identical distribution is induced by some Nash equilibrium or even ε -Nash equilibrium. We have to restrict ourselves to a statement that a close distribution is induced by an ε -Nash equilibrium. Important differences in notions of closeness of measures and of ε -optimization, and their implications for issues of learning and stability, are discussed in portions of this paper.

The notion of subjective equilibrium used in this paper is closely related to the notion of self-confirming equilibrium of Fudenberg and Levine (1993b) discussed earlier. Self-confirming equilibrium is defined on a general class of finite extensive form games. It requires, as does subjective equilibrium, that the distribution on the finite play paths of the game coincide with the players' beliefs. However, it is more general, since it allows a player to believe that his opponents' actions off the positive probability play paths are correlated. Fudenberg and Levine showed, under assumptions closely related to ours, that coincidence of self confirming equilibrium behavior with Nash behavior is obtained.

Our objective in this paper is to describe general sufficient conditions under which subjective equilibrium behavior coincides with Nash behavior, and subjective ε -equilibrium behavior is ε -close to ε -Nash equilibrium behavior in the space of infinite play paths. Since correlations off the play paths will have to be assumed away in the statements of our main results, we prefer the simplicity gained by assuming them away in the definition of subjective equilibrium. For this reason we restrict the beliefs in a subjective equilibrium to consist entirely of (independent) behavior strategies.

It is important to note that if one player's beliefs regarding an opponent's strategy were described by a mixed strategy, i.e., believing that his opponent chose randomly one strategy from a set of behavior strategies, then by using the standard Kuhn (1953) method we could replace his beliefs with an equivalent single behavior strategy to fit the model of this paper. Disallowing correlations, as discussed above, will restrict us to the use of individually mixed strategies and thus rule out the mixing of strategies in a correlated way across players.

2. THE REPEATED GAME

First, we briefly review the standard model of an n -player discounted repeated game with perfect monitoring. An n -person stage game is described by a set of *action combination* $\Sigma = \times_{i=1}^n \Sigma_i$ with Σ_i denoting a finite *action set* of player i . Functions $u_i: \Sigma \rightarrow \mathbb{R}$ describe the *stage game payoffs* of the players.

The set of *histories of length t* , H_t , is defined to be the set of all $t \geq 1$ tuples of elements of Σ , i.e., Σ^t , and H_0 is a singleton set containing the "empty history."

2A. THE FINITELY REPEATED GAME

For any integer l we define the finitely repeated game which consists of l iterations of the stage game just described. Let H^l be the set of all histories with length less than l . That is, $H^l = \cup_{t=0}^{l-1} H_t$. A (*behavior*) *strategy* of player i is a function $f_i^l: H^l \rightarrow \Delta(\Sigma_i)$ with $\Delta(\Sigma_i)$ denoting the set of probability distributions over Σ_i . Notice that, in the definition of a strategy, we implicitly assume that the game is played with perfect monitoring.

A *strategy vector* $f^l = (f_1^l, f_2^l, \dots, f_n^l)$ is a vector consisting of individual strategies. Such a strategy vector induces a probability distribution μ_{f^l} over H^l , which is also the set of all the play paths of the finite game, as defined below.

With every history h in H^l we associate a cylinder set $c(h)$ consisting of all the play paths going through h , i.e., having h as a prefix. We will use h to denote the history and also to denote the event $c(h)$ when we think of the space of play paths.

We define μ_{f^l} of the empty history to be one and proceed inductively. For a history h and an action combination a , we define $\mu_{f^l}(ha) = \mu_{f^l}(h) \prod_i f_i^l(h)(a_i)$. Thus, the probability of the history consisting of h followed by the action combination a is the probability of h times the product of the conditional probabilities of each player taking his action a_i given the history h .

We assume that each player has a discount parameter $\lambda_i, 0 < \lambda_i < 1$, by which he evaluates the payoff received along play paths. Thus, if $z^l = (z_1, z_2, \dots, z_l)$ is a play path, we define

$$u_i(z^l) = \sum_{t=1}^l \lambda_i^{t-1} u_i(z_t^l).$$

Now we complete the definition of the repeated game by defining individual payoffs for each strategy vector f^l ,

$$U_i(f^l) = E u_i(z^l) = \int u_i(z^l) d\mu_{f^l}(z^l).$$

Equivalently, one can define the expected stage payoffs and take the discounted sum of these.

2B. THE INFINITELY REPEATED GAME

The set of all finite length histories is denoted by H . I.e., $H = \bigcup_{t=0}^{\infty} H_t$. A (behavior) strategy of player i in the infinitely repeated game is a function f_i from H to $\Delta(\Sigma_i)$. Note that any f_i induces a strategy, f_i^l , in the corresponding l -fold repeated game. The f_i^l is simply the restriction of f to the smaller domain, H^l , and it is called the l -truncation of f .

For every strategy vector $f = (f_1, \dots, f_n)$ we define a probability distribution, μ_f , over the set of all infinite play paths Σ^∞ . The σ -algebra for this set is defined to be the smallest one that contains all the cylinder sets, $c(h)$. Following a standard probability formulation it suffices to assign probabilities to all the cylinder sets in order to obtain a probability distribution over the set of all play paths.

For every finite history $h \in H_l$ we define $\mu_{f^l}(h) = \mu_f(h)$. As in the finite case we define

$$u_i(z) = \sum_{t=1}^{\infty} \lambda_i^{t-1} u_i(z_t)$$

for every infinite play path $z = (z_1, z_2, \dots)$ and

$$U_i(f) = E u_i(z) = \int u_i(z) d\mu_f(z).$$

Unless specified, statements in the sequel refer to both the finitely and infinitely repeated game.

We say that a sequence of player i 's strategies, $f_{i,n}$, converges to f_i if $f_{i,n}(h)$ converges to $f_i(h)$ for every finite history h .

3. DESCRIPTION OF THE MAIN RESULT

In Kalai and Lehrer (1993a) it was shown that subjectively rational players, with truth compatible beliefs, must converge in time to play subjective equilibria and converge in finite time to play subjective ε -equilibria, as defined below.

The rationality of such a player i is expressed by the fact that he chooses his strategy g_i to be a best (maximizing expected utility) response to his beliefs that strategies $g^i = (g_1^i, \dots, g_n^i)$ will be played. Truth compatibility of beliefs is expressed by the fact that the true distribution on the induced play of the game, μ_g , is absolutely continuous with respect to the belief distribution, μ_{g^i} . In other words, no positive probability event in the game is assigned probability zero by the beliefs.

Given that a player maximizes expected utility relative to a prior probability distribution, it follows that he must be maximizing expected utility relative to his Bayes updated posterior beliefs after positive probability histories. As was shown earlier by Blackwell and Dubins (1962), and by the recent Kalai and Lehrer paper, absolute continuity guarantees that Bayes posteriors converge to the true distribution. So in the limit the players will predict the future correctly and will play a subjective equilibrium; and in finite time they will predict the future correctly only up to ε and will play subjective ε -equilibrium. We begin by considering the case of correct predictions in the future, and hence of subjective equilibrium.

As usual, we say that a strategy f_i is a best response to g_{-i} if $U_i(g_1, \dots, k_i, \dots, g_n) - U_i(g_1, \dots, f_i, \dots, g_n) \leq 0$ for every strategy k_i . If the right side 0 is replaced by ε we say that f_i is an ε -best response to g_{-i} .

DEFINITION 1: A *subjective equilibrium* is a strategy vector g with a beliefs matrix $(g_j^i)_{1 \leq i, j \leq n}$ satisfying for each player i :

- (0) $g_i^i = g_i$;
- (1) g_i is best response to g_{-i}^i ; and
- (2) $\mu_g = \mu_{g^i}$.

In this case, we say that the matrix (g_j^i) sustains g .

The idea is that the i th row, g^i , represents the subjective assessment of player i about the strategy vector that is played. Condition (0) requires that every player knows his own strategy. Condition (1) expresses the usual utility maximization assumption where g_{-i}^i is the $(n - 1)$ vector consisting of all the entries of g^i excluding the i th. Condition (2) states that g plays like g^i . It expresses the idea that any statistical study will serve only to strengthen player i 's belief that g^i is being played.

Obviously, in the above definition, if we let $g^i = g$ for all i we have a *Nash equilibrium* as a special case. In subjective equilibrium, however, we allow for inaccuracy of beliefs regarding opponents' actions after histories that are never observed.

As mentioned in the Introduction, describing player i 's beliefs about j 's strategy by a single behavior strategy, g_j^i , is not a serious restriction. By a well-known theorem of Kuhn (1953) (see also Aumann (1964) and Selten (1975)), a mixed strategy of player j , i.e., a probability distribution over his behavior strategies, can be replaced by a single equivalent behavior one. (The equivalence is strong since the Kuhn's constructed strategy plays exactly as the mixed strategy with every choice of strategies by j 's opponents.) Thus, a belief of player i given by a probability distribution over j 's strategies can be replaced by an equivalent belief consisting of a single behavior strategy.

As was already outlined in the Introduction, one can show that the behavior induced by a subjective equilibrium coincides with the behavior induced by some Nash equilibrium.

PROPOSITION 1: Let $g = (g_1, \dots, g_n)$ be a subjective equilibrium. There is Nash equilibrium $f = (f_1, \dots, f_n)$ with $\mu_f = \mu_g$.

We omit the proof of this proposition and refer the readers to Fudenberg and Levine (1993b), and Battigalli et al. (1992) for earlier references.

Returning to the subjectively rational model, where players start with private beliefs, in any finite time the Bayesian updating process will typically become only approximately correct. This means that after finite time, even if arbitrarily long, we can only assume that the players play approximate subjective equilibria which we proceed to define.

DEFINITION 2: Let $\varepsilon > 0$ and let μ and $\tilde{\mu}$ be two probability measures defined on the same space. We say that μ is ε -close to $\tilde{\mu}$ if there is a measurable set Q satisfying:

- (i) $\mu(Q)$ and $\tilde{\mu}(Q)$ are greater than $1 - \varepsilon$; and
- (ii) for every measurable set $A \subseteq Q$ $|\mu(A) - \tilde{\mu}(A)| \leq \varepsilon \tilde{\mu}(A)$.

REMARK 1: In Blackwell and Dubins' (1962) paper on merging of measures, closeness of measures was expressed by $|\mu(A) - \tilde{\mu}(A)| \leq \varepsilon$ for every event A (not just in Q). Their easy-to-state condition seems weaker, since it implies little for small probability events. For example, $\mu(A)$ could equal $2\tilde{\mu}(A)$ and still satisfy the Blackwell-Dubins closeness provided that $\mu(A)$ is sufficiently small. It turns out, however (see Kalai and Lehrer (1993b)), that the two notions are asymptotically equivalent, and the results that follow can be stated using either condition. The notion stated in Definition 2, despite its length, has the advantage of being explicit on small probability events. It states that they are approximated well, provided that we stay away from the small excluded set $\Omega - Q$. Thus, when μ and $\tilde{\mu}$ are close in this sense we can think of $\tilde{\mu}$ as being "probably approximately correct" as an estimator of μ . We can think of the notion used by Blackwell and Dubins as being approximately correct for all large events. For repeated games, being correct on small probability events is important since even significant events may have small probability if they occur late in the game. It implies having approximately correct conditional probability, even when conditioning on histories that occur arbitrarily late in the game. Thus, predictions on the play of arbitrarily late subgames will probably be correct.

DEFINITION 3: Let f and g be two joint strategies, and let $\varepsilon > 0$. We say that g plays ε -like f if μ_g is ε -close to μ_f .

The above notion of playing ε -like is strong, since it uses a strong notion of closeness of measures. Except for a set of infinite play paths with measure of at most ε , it guarantees that even for very long histories, ones that have small probability of being reached, f and g will assign close probabilities, with ratio close to one. In other words, such f, g are very likely to play almost the same throughout an infinite game and in its positive probability subgames.

DEFINITION 4: An n -vector of strategies, g , is a subjective ε -equilibrium if there is a matrix of strategies $(g^i)_{1 \leq i, j \leq n}$ such that for every player i :

- (a) $g^i = g_i$;
- (b) g_i is a best response to g^i_{-i} ; and
- (c) g plays ε -like g^i .

An n -vector of strategies, f , is an ε -Nash equilibrium, if each f_i is an ε -best response to f_{-i} .

Our main result deals with the relation between subjective ε -equilibrium and ε -Nash equilibrium.

THEOREM 1: *In infinitely repeated games, for every $\varepsilon > 0$ there is $\bar{\eta} > 0$ such that for all $\eta \leq \bar{\eta}$, if g is a subjective η -equilibrium, then there exists f , such that*

- (i) g plays ε -like f , and
- (ii) f is an ε -Nash equilibrium.

Theorem 1 states that a behavior induced by a subjective ε -equilibrium must be close to a behavior induced by an ε -Nash equilibrium. An ε -Nash equilibrium requires each player to choose a strategy that is ε -best response against the precise strategies used by his opponents, i.e., his payoff should be within ε of the optimally possible against theirs. On the other hand, the subjective ε -equilibrium requires precise optimization but against beliefs that are almost accurate.

The easy proof of Proposition 1 outlined in the Introduction made use of the precise coincidence of the play and conjectured play of all the players. However, in Theorem 1, with only ε -precision, this is no longer the case. Instead, our construction takes advantage of the fact that, in a repeated game with discounting, the payoff function of a player is continuous. That is, small changes in the strategies played, in particular after long histories, will not much affect the overall payoff. As an intermediate step we prove a stronger version of Theorem 1 for finitely repeated games. The result is interesting in its own right. Later we use the result of the finite case for proving the infinite case (see Fudenberg and Levine (1986) for finite truncations of infinite games and their equilibrium).

PROPOSITION 2: *In finitely repeated games, for every $\varepsilon > 0$ there is $\bar{\eta} > 0$ such that for all $\eta < \bar{\eta}$, if g is a subjective η -equilibrium, then there exists f such that*

- (i) g plays ε -like f , and
- (ii) f is a Nash equilibrium.

Note that, in Proposition 2, a subjective η -equilibrium ε -plays like Nash (not ε -Nash) equilibrium.

4. PROOFS OF THE MAIN RESULT

PROOF OF PROPOSITION 2: Suppose to the contrary that there is $\varepsilon > 0$ and a sequence of strategy vectors $g(m)$ such that

- (i) $g(m)$ is subjective η_m -equilibrium, where $\eta_m \rightarrow 0$ as $m \rightarrow \infty$, and
- (ii) $g(m)$ does not play ε -like any Nash equilibrium.

Since $g(m)$ is a subjective η_m -equilibrium, there is a matrix $(g(m)_i^j)$ which sustains it. In finitely repeated games each player has a finite number of pure strategies. Therefore, the set of behavior strategies is sequentially compact. Thus, without loss of generality, the sequences $\{g(m)\}_m$ and $\{(g(m)_i^j)\}_m$ are converging to, say, g and to (g_i^j) . As the payoff functions are continuous, g is subjective equilibrium sustained by (g_i^j) . Moreover, if η_m is close enough to zero $g(m)$ ε -plays like g .

Using Proposition 1 we can find a Nash equilibrium f which plays 0-like g . Thus, if η_m is sufficiently small, $g(m)$ ε -plays like f , which is a Nash equilibrium. This is a contradiction. *Q.E.D.*

REMARK 2: In Definition 4(b) we required that g_i be a best response to g_{-i}^i . One can define δ -subjective ε -equilibrium by replacing “best response” with “ δ best response.” A similar proof to the one of Proposition 2 shows that, in a finitely repeated game, for

every $\varepsilon > 0$ there is $\bar{\eta}$ such that if $\eta < \bar{\eta}$ then for every δ -subjective η -equilibrium, g , there is f such that

- (i) g plays ε -like f , and
- (ii) f is a δ -Nash equilibrium.

Notice that the relation between ε and η is independent of δ .

We are now ready to prove Theorem 1. Starting with a subjective η -equilibrium g we consider its truncation to the finitely repeated game of length l . If l is large then the truncated g is an approximate subjective η -equilibrium of the finite game, and by the above remark, it must approximately play like some Nash equilibrium f of the finite game. We extend f to the infinite game by making it coincide with g after all histories longer than l . This extension makes g play close to f in the infinite game, and exploiting again the fact that l is large, we conclude that f must be an approximate Nash equilibrium of the infinite game.

PROOF OF THEOREM: Let $\varepsilon > 0$. Observe first that there is an integer $l = l(\varepsilon)$ such that: (i) if a strategy k_i^l is δ -best response to k_{-i}^l in the l -fold repeated game, then any strategy k_i of the infinitely repeated game whose l -truncation (see 2a above) coincides with k_i^l is an $(\delta + \varepsilon/2)$ -best response to any k_{-i} , whose l -truncation coincides with k_{-i}^l ; and (ii) if k_i is a best response to k_{-i} in the infinitely repeated game, then k_i^l is an $\varepsilon/2$ -best response to k_{-i}^l .

In view of Remark 2 there exists an $\bar{\eta}$ which corresponds to $\delta = \varepsilon/2$ and to the l -fold repeated game. Let g be a subjective η -equilibrium for some $\eta < \bar{\eta}$. g^l is therefore an $\varepsilon/2$ -subjective η -equilibrium in the l -fold repeated game. Therefore, by Remark 2 applied to $\varepsilon/2$, it plays $\varepsilon/2$ -like some $\varepsilon/2$ -Nash equilibrium, say, f^l .

In order to conclude the proof we should define a strategy f of the infinitely repeated game whose l -truncation coincides with f^l and, moreover, have g ε -plays like it. Thus, we have only to define f on histories longer than l . Let h be such a history ($h \in H_{l'}$, for $l' > l$); define $f_i(h) = g_i(h)$.

Denote the partition of Σ^∞ induced by H_l by \mathcal{P}_l . The fact that g^l $\varepsilon/2$ -plays (therefore ε -plays) like f^l means that there exists a set Q which is a union of atoms of \mathcal{P}_l satisfying

- (i) $\mu_{f^l}(Q) > 1 - \varepsilon$ and $\mu_{g^l}(Q) > 1 - \varepsilon$, and
- (ii) $(1 - \varepsilon)\mu_{f^l}(C) \leq \mu_{g^l}(C) \leq (1 + \varepsilon)\mu_{f^l}(C)$ for every atom $C \subseteq Q$.

To show that g plays ε -like f consider an event A of infinite paths in Q . Notice that by the definition of f $\mu_f(A|C) = \mu_g(A|C)$ for any atom $C \in \mathcal{P}_l$. Therefore,

$$\begin{aligned} (1 - \varepsilon)\mu_f(A) &= (1 - \varepsilon) \sum_{C \in \mathcal{P}_l, C \subseteq Q} \mu_f(A \cap C) \\ &= (1 - \varepsilon) \sum_{C \in \mathcal{P}_l, C \subseteq Q} \mu_f(A|C)\mu_f(C) \\ &= (1 - \varepsilon) \sum_{C \in \mathcal{P}_l, C \subseteq Q} \mu_g(A|C)\mu_f(C) \\ &\leq \sum_{C \in \mathcal{P}_l, C \subseteq Q} \mu_g(A|C)\mu_g(C) = \mu_g(A). \end{aligned}$$

For a similar argument, $\mu_g(A) \leq (1 + \varepsilon)\mu_f(A)$ which concludes the proof that g plays ε -like f .

Recall that f^l is $\varepsilon/2$ -Nash equilibrium in the l -fold repeated game. Therefore, f is $(\varepsilon/2 + \varepsilon/2)$ -Nash equilibrium in the infinite repeated game, which completes the proof of the theorem. Q.E.D.

REMARK 3: It is easy to find examples where the behavior of an ε -Nash equilibrium is not ε -close to any Nash equilibrium under the strong notion of closeness we use.

EXAMPLE: Suppose that in a two person game each player has two actions, say, a and b . Suppose furthermore that the pair (a, a) is the unique Nash equilibrium of the stage game (as in the prisoners' dilemma). Consider the following time dependent and not history dependent strategy: play always a and only at time t play b . Denote this strategy by g^t . Note that since future payoffs are discounted, for very $\varepsilon > 0$ there is t large enough so that $g = (g^t, g^t)$ is an ε -Nash equilibrium.

As g^t is a pure strategy, μ_g is concentrated in one play path. If the discount factor is close to 0, there is no Nash equilibrium with which it plays ε -like, even for very large t 's.

Notice that in the above example g^t is an ε -best response because of the discounting and the fact that the suboptimal action is taken late in the game. The pair (g^t, g^t) will fail, however, to be a *time-consistent ε -Nash equilibrium*. At such an equilibrium, after every positive probability history, the induced vector of strategies must be an ε -equilibrium of the induced game, computed from the local time on.

To obtain in Theorem 1 a behavior approximating time-consistent ε -Nash equilibrium, we can use a notion of *strong subjective equilibrium*. Such a vector of strategy vectors (g, g^1, \dots, g^n) will satisfy as before $g_i = g_i^i$ (knowing your own strategy) and g_i is a best-response to g_{-i}^i (subjective optimization). But it will be required to satisfy a stronger consistency of the beliefs with the actual observed play, i.e.,

$$|\mu_g(h) / \mu_{g^i}(h) - 1| < \varepsilon$$

for all μ_g -positive probability histories h (not just for histories h in a large set Q).

5. A WEAKER NOTION OF CLOSENESS

As mentioned above, our notion of closeness of measures was inspired by the results regarding convergence of beliefs to the truth in a repeated game. Under this notion, a subjective ε -equilibrium plays ε -like some ε -Nash equilibrium. Does it play ε -like Nash, rather than ε -Nash equilibrium? With a weaker notion of closeness it does.

DEFINITION 5: Let g and f be two strategy vectors, $\varepsilon > 0$, and l be an integer. We say that g plays (ε, l) like f if $|\mu_f(A) - \mu_g(A)| < \varepsilon$ for every event A which consists of histories of length l .

THEOREM 2: For any (ε, l) there is $\bar{\eta}$ such that if $\eta < \bar{\eta}$, then every subjective η -equilibrium g plays (ε, l) like some Nash equilibrium.

PROOF: In view of Proposition 1, it suffices to show that g plays (ε, l) like subjective equilibrium. We proceed by assuming that the theorem is incorrect. Thus, there is (ε, l) and a sequence g_n of subjective η_n -equilibrium such that $\eta_n \rightarrow 0$ and g_n does not (ε, l) play like any subjective equilibrium.

The limit of g_n , say, g , is clearly a subjective equilibrium. Moreover, on finite histories g_n and g are very close when n is sufficiently large. Therefore, g_n plays (ε, l) like g when n is large enough. This is a contradiction. Q.E.D.

REMARK 4: Two different notions of "approximate playing like" were used in the previous theorem. The weaker new one was used for the approximation as is explicitly stated, but the old one was still implicitly used in the definition of subjective η -equilibrium. One can strengthen Theorem 2 by using also the weaker notion of closeness in

the definition of a subjective η -equilibrium. The added advantage of consistency, using the same definition of closeness throughout, is attractive. However, it would require the introduction of yet another version of subjective equilibrium, which we chose to avoid.

Dept. of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, Il. 60208, U.S.A.

Manuscript received March, 1991; final revision received January, 1993.

REFERENCES

- AUMANN, R. J. (1964): "Mixed and Behaviour Strategies in Infinite Extensive Games," in *Advances in Game Theory*, Annals of Mathematics Studies, 52, ed. by M. Dresher, L. S. Shapley and A. W. Tucker. Princeton: Princeton University Press, pp. 627–650.
- BATTIGALLI, P. (1987): "Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali," Unpublished Dissertation, Bocconi University, Milano.
- BATTIGALLI, P., M. GILLI, AND M. C. MOLINARI (1992): "Learning Convergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey," *Ricerche Economiche*, forthcoming.
- BATTIGALLI, P., AND D. GUAITOLI (1988): "Conjectured Equilibria and Rationalizability in a Macroeconomic Game with Incomplete Information," Bocconi University.
- BLACKWELL, D., AND L. DUBINS (1962): "Merging of Opinions with Increasing Information," *Annals of Mathematical Statistics*, 38, 882–886.
- FUDENBERG, D., AND D. KREPS (1988): "A Theory of Learning, Experimentation, and Equilibrium in Games," Stanford University.
- FUDENBERG, D., AND D. LEVINE (1986): "Limit Games and Limit Equilibria," *Journal of Economic Theory*, 38, 261–279.
- (1993a): "Steady State Learning and Nash Equilibrium," *Econometrica*, 61, 547–573.
- (1993b): "Self Confirming Equilibrium," *Econometrica*, 61, 523–545.
- HAHN, F. (1973): *On the Notion of Equilibrium in Economics: An Inaugural Lecture*. Cambridge: Cambridge University Press.
- HAYEK, F. A., VON (1937): "Economics of Knowledge," *Economica*, 4, 33–54.
- KALAI, E., AND E. LEHRER (1990): "Learning to Play Equilibrium," Northwestern University.
- (1993a): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.
- (1993b): "Weak and Strong Merging of Opinions," *Journal of Mathematical Economics*, forthcoming.
- KUHN, H. W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games, Vol. II*, Annals of Mathematics Studies, 28, ed. by H. W. Kuhn and A. W. Tucker. Princeton: Princeton University Press, pp. 193–216.
- RUBINSTEIN, A., AND A. WOLINSKY (1990): "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability," *Games and Economic Behavior*, forthcoming.
- SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.
- WITTLE, P. (1982): *Optimization Over Time*, Vol. 1. New York: Wiley.