



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 03-021

“Contemporaneous Perfect Epsilon-Equilibria”

by

George J. Mailath, Andrew Postlewaite and Larry Samuelson

<http://ssrn.com/abstract=436821>

Contemporaneous Perfect Epsilon-Equilibria*

George J. Mailath,[†] Andrew Postlewaite[‡] Larry Samuelson[§]

August 15, 2003

Abstract

We examine *contemporaneous perfect ε -equilibria*, in which a player's actions after every history, evaluated at the point of deviation from the equilibrium, must be within ε of a best response. This concept implies, but is stronger than, Radner's ex ante perfect ε -equilibrium. A strategy profile is a contemporaneous perfect ε -equilibrium of a game if it is a subgame perfect equilibrium in a perturbed game with nearly the same payoffs, with the converse holding for pure equilibria. *Keywords:* Epsilon equilibrium, ex ante payoff, multistage game, subgame perfect equilibrium. *JEL classification numbers* C70, C72, C73.

1. Introduction

Analyzing a game begins with the construction of a model specifying the strategies of the players and the resulting payoffs. For many games, one cannot be positive that the specified payoffs are precisely correct. For the model to be useful, one must hope that its equilibria are close to those of the real game whenever the payoff misspecification is small.

To ensure that an equilibrium of the model is close to a Nash equilibrium of *every* possible game with nearly the same payoffs, the appropriate solution concept in the model is some version of strategic stability (Kohlberg and Mertens (1986)). In this note, we take the alternative perspective of an analyst seeking to ensure that no Nash equilibria of the real game are neglected. The appropriate solution concept in the model is then ε -Nash equilibrium: It is a straightforward exercise (Proposition 3 below) that a

*We thank Drew Fudenberg, David Levine, two referees, and an Associate Editor for helpful comments and the National Science Foundation for financial support (under grants SES-9911219, SES-0095768 and SES-0241506).

[†]University of Pennsylvania, gmailto:mailath@econ.upenn.edu

[‡]University of Pennsylvania, apostlew@econ.upenn.edu

[§]University of Wisconsin, larrysam@ssc.wisc.edu

strategy profile is an ε -Nash equilibrium of a game if it is a Nash equilibrium of a game with nearly the same payoffs.

When dealing with extensive-form games, one is typically interested in sequential rationality. Radner (1980) defined a *perfect ε -equilibrium* as a strategy profile in which each player, following every history of the game and taking opponents' strategies as given, is within ε of the largest possible payoff. In Radner's perfect ε -equilibria, however, the gain from deviating from the proposed equilibrium strategy is evaluated after having been discounted back to the beginning of the game. We accordingly refer to this as an *ex ante perfect ε -equilibrium*. If the game is played over time and the players' payoffs are the discounted present values of future payments, then there may be ex ante perfect ε -equilibria in which a player has a deviation strategy that yields a large increase in his payoff in a distant period, but a quite small gain when discounted to the beginning of the game. At the beginning of the game, the player's behavior will then be ex ante ε optimal. But conditional on reaching the point where the deviation is to occur, the gain will be large. As a result, such an ε -equilibrium will *not* be a subgame perfect equilibrium of any game with nearly the same payoffs.

We propose an alternative definition of approximate equilibrium that requires every player to be within ε of his optimal payoff after every history, where the evaluation of strategies is made contemporaneously, that is, evaluations are made at the point that an alternative strategy deviates from the proposed strategy. We call a vector of strategies that satisfies this criterion a *contemporaneous perfect ε -equilibrium*. Following the preliminaries presented in Sections 2–4, Section 5 shows that any subgame perfect equilibrium in a nearby game is a contemporaneous perfect ε -equilibrium in the game in question, with the converse holding for pure strategies.

2. Multistage Games with Observed Actions

We consider multistage games with observed actions (Fudenberg and Tirole (1991, ch. 4)). There is a potentially infinite number of periods $\{0, 1, \dots\}$. In each period, some subset of the players simultaneously choose from nontrivial feasible sets of actions, knowing the history of the game through the preceding period. In each period, the feasible sets may depend upon the period and the history of play, and some players may have no choices to make. Let G denote such a game. This class of games includes both repeated games and dynamic games like Rubinstein's (1982) alternating-offers bargaining game.

Every information set for a player in period t corresponds to a particular history of actions taken before period t . The converse need not be true, however, since a player may be constrained to do nothing in some periods or after some histories. In addition, some finite histories may correspond to terminal nodes ending the game, as is the case after an agreement in alternating-offers bargaining. The set of histories corresponding

to an information set for player i is denoted H_i^* , and the set of all histories is denoted H . Histories in $H \setminus \cup_i H_i^*$ are called *terminal histories*.

The set of actions available to player i at the information set $h \in H_i^*$ is denoted by $A_i(h)$. We assume that each $A_i(h)$ is finite. Let σ_i be a behavior strategy for player i , so that $\sigma_i(h) \in \Delta(A_i(h))$ associates a mixture over $A_i(h)$ to h . Endow each $\Delta(A_i(h))$ (a subset of a Euclidean space) with the standard topology and endow the set of strategy profiles with the product topology.

Let σ_i^h be the strategy σ_i , modified at (only) i 's information sets preceding h so as to take those pure actions consistent with play generating the history h . In multistage games with observed actions, the actions specified by σ_i^h and σ_{-i}^h are unique at each of the preceding information sets. The length of the history h is denoted $t(h)$. Since the initial period is period 0, actions taken at the information set h are taken in period $t(h)$.

In a dynamic environment, players may receive payoffs at different times. We are interested in the difference between a decision with immediate monetary or physical consequences and a decision with the same monetary or physical consequences, but realized at some point in the future. To capture this distinction, we formulate payoffs in terms of a discounting scheme and a reward function.

The reward function is denoted by $r_i : H \rightarrow \mathbb{R}$, where $r_i(h)$ is the reward player i receives after the history h . We emphasize that the reward $r_i(h)$ is received in period $t(h) - 1$ (recall that the initial period is period 0), and that it can depend on the *entire* sequence of actions taken in the preceding $t(h)$ periods.

Player i discounts period t rewards to period $t - 1$ using the factor $\delta_{it} \in (0, 1]$. Define $\delta_i^{(t', t'')} \equiv \prod_{\tau=t'+1}^{t''} \delta_{i\tau}$, so that a reward r_i received in period t'' has value $\delta_i^{(t', t'')} r_i$ in period t' . We sometimes write $\delta_i^{(t)}$ for $\delta_i^{(0, t)}$. We set $\delta_{i0} = \delta_i^{(t, t)} = 1$. Finally, for notational simplicity, if the game has a finite horizon T , we set $\delta_{it} = \beta$ for some fixed $\beta \in (0, 1)$ for all $t \geq T$.

We assume that players discount, in that there exists $D < \infty$ such that, for all i ,

$$\sup_T \sum_{t=T}^{\infty} \prod_{\tau=T}^t \delta_{i\tau} \leq D. \quad (1)$$

This discounting formulation is sufficiently general as to impose very little restriction on the payoffs of the game. For example, the possibility of different discount factors in different periods allows us to capture games like Rubinstein's alternating-offers bargaining game, where (using our numbering convention for periods) offers are made in even periods, acceptance/rejections in odd periods, and $\delta_{it} = 1$ for all even t . In addition, we have imposed no bounds on the reward functions r_i . Hence, by allowing rewards to grow sufficiently fast, we can model games in which future payoffs have larger present values than current ones, even with discounting. However, the discounting scheme is

essential in capturing the player's relative evaluation of rewards received in different periods, and hence to our study of ex ante and contemporaneous perfect ε -equilibria.

The set of pure strategy profiles is denoted by Σ , the outcome path induced by the pure strategy profile $s \in \Sigma$ is denoted by $a^\infty(s)$, and the initial $t + 1$ period history is denoted by $a^t(s)$. For notational simplicity, if $a^\infty(s)$ is a terminal history of length T , we define $r_i(a^t(s)) = 0$ for all $t \geq T$. Player i 's payoff function, $\pi_i : \Sigma \rightarrow \mathbb{R}$, is given by

$$\pi_i(s) = \sum_{t=0}^{\infty} \left(\prod_{\tau=0}^t \delta_{i\tau} \right) r_i(a^t(s)) = \sum_{t=0}^{\infty} \delta_i^{(t)} r_i(a^t(s)). \quad (2)$$

We assume the reward function is such that this expression is well-defined for all $s \in \Sigma$. We extend π_i to the set of behavior strategy profiles, Σ^* , in the obvious way.

This representation of a game is quite general. In Rubinstein's alternating-offers bargaining game, $r_i(h)$ equals i 's share if an agreement is reached in period $t(h)$ under h , and zero otherwise. In the T -period centipede game, we let $\delta_{it} = 1$ for $t < T$ (since there are only finitely many periods, this satisfies our discounting assumption) and let $r_i(h)$ equal i 's payoff when the game is stopped in period $t(h)$ under h .

Define $\pi_i(\sigma|h)$ as the continuation payoff to player i under the strategy profile σ , conditional on the history h . For pure strategies $s \in \Sigma$, we have (recall that $\delta_i^{(t,t)} = 1$):

$$\begin{aligned} \pi_i(s|h) &= r_i(a^{t(h)}(s^h)) + \sum_{t=t(h)+1}^{\infty} \left(\prod_{\tau=t(h)+1}^t \delta_{i\tau} \right) r_i(a^t(s^h)) \\ &= \sum_{t=t(h)}^{\infty} \delta_i^{(t(h),t)} r_i(a^t(s^h)). \end{aligned} \quad (3)$$

Note that $a^{t(h)}(s^h) \equiv (a^{t(h)-1}(s^h), a_{t(h)}(s^h))$ is the concatenation of the history of actions that reaches h and the action profile taken in period $t(h)$.

3. Epsilon Equilibria

The strategy profile σ^h specifies a unique history of length $t(h)$ that causes information set h to be reached, allowing us to write:

$$\pi_i(\sigma^h) = \sum_{t=0}^{t(h)-1} \delta_i^{(t)} r_i(a^t(\sigma^h)) + \delta_i^{(t(h))} \pi_i(\sigma^h|h).$$

In other words, for a fixed history h and strategy profile σ_{-i} , $\pi_i(\sigma_{-i}^h, \cdot)$ is a player i payoff function on the space of player i 's strategies of the form σ^h that is a positive affine transformation of the payoff function $\pi_i(\sigma_{-i}, \cdot|h)$.

Definition 1 For $\varepsilon > 0$, a strategy profile $\hat{\sigma}$ is an ε -Nash equilibrium if, for each player i and strategy σ_i ,

$$\pi_i(\hat{\sigma}) \geq \pi_i(\hat{\sigma}_{-i}, \sigma_i) - \varepsilon.$$

A strategy profile $\hat{\sigma}$ is an ex ante perfect ε -equilibrium if, for each player i , history h , and strategy σ_i ,

$$\pi_i(\hat{\sigma}^h) \geq \pi_i(\hat{\sigma}_{-i}^h, \sigma_i^h) - \varepsilon.$$

A strategy profile $\hat{\sigma}$ is a contemporaneous perfect ε -equilibrium if, for each player i , history h , and strategy σ_i ,

$$\pi_i(\hat{\sigma}|h) \geq \pi_i(\hat{\sigma}_{-i}, \sigma_i|h) - \varepsilon.$$

Ex ante ε -perfection appears in Radner (1980) and Fudenberg and Levine (1983).¹

Any contemporaneous perfect ε -equilibrium is an ex ante perfect ε -equilibrium, and the two concepts coincide in the absence of discounting or when $\varepsilon = 0$ (in which case they also coincide with subgame perfection). Radner studies ε -equilibria in a repeated oligopoly that are ex ante but not contemporaneous ε -equilibria. We use the finitely repeated prisoners' dilemma to capture the spirit of his analysis, showing that ex ante and contemporaneous perfect ε -equilibria for the same value of $\varepsilon > 0$ can be quite different:

Example 1: The finitely repeated prisoners' dilemma. The stage game is given by

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	-1, 3
<i>D</i>	3, -1	0, 0

This game is played $N + 1$ times, with payoffs discounted according to the common discount factor $\delta < 1$. The unique Nash (and hence subgame perfect) equilibrium features perpetual defection. Consider “trigger” strategies that specify cooperation after every history featuring no defection, and defection otherwise. If δ is sufficiently close to 1, the only potentially profitable deviation will be to defect in period N . As long as N is sufficiently large that

$$\delta^N < \varepsilon, \tag{4}$$

the benefit from this defection is below the ε threshold, and the trigger strategies are an ex ante perfect ε -equilibrium. However, for any $\varepsilon < 1$, the unique contemporaneous perfect ε -equilibrium is to always defect.² ||

¹Watson (1994) considers an intermediate concept that requires ε -optimality conditional only on those histories that are reached along the equilibrium path.

²Radner (1980, p. 153) also defines an alternative notion of perfect ε -equilibrium in which the utility of a continuation strategy is calculated relative to the period at which the decision is being made. This

In Example 1, for sufficiently small ε (in particular, so that (4) is violated), both players must defect in every period in any ex ante perfect ε -equilibrium of a finitely repeated prisoners' dilemma. More generally, in finite horizon games, for sufficiently small ε , ex ante perfect ε -equilibria and contemporaneous perfect ε -equilibria coincide:

Proposition 1 *Suppose G is a finite game (so that it has finite horizon and finite action sets). For sufficiently small ε , the sets of ex ante perfect pure-strategy ε -equilibria and of contemporaneous perfect pure-strategy ε -equilibria coincide, and they coincide with the set of pure-strategy subgame perfect equilibria.*

Proof. Observe that any subgame perfect equilibrium is necessarily both an ex ante and a contemporaneous perfect ε -equilibrium. Suppose then that \hat{s} is not a subgame perfect equilibrium. We will show that for ε sufficiently small, \hat{s} is neither an ex ante nor a contemporaneous perfect ε -equilibrium.

Since \hat{s} is not subgame perfect, there is some player i , history h and strategy s_i such that

$$\pi_i(\hat{s}_{-i}, s_i|h) - \pi_i(\hat{s}|h) > 0.$$

Since the game is finite, there exists an ε' sufficiently small such that, for all such h , i , and s_i ,

$$\pi_i(\hat{s}_{-i}, s_i|h) - \pi_i(\hat{s}|h) > \varepsilon'.$$

But,

$$\begin{aligned} \pi_i(\hat{s}_{-i}^h, s_i^h) - \pi_i(\hat{s}^h) &= \delta_i^{t(h)} [\pi_i(\hat{s}_{-i}, s_i|h) - \pi_i(\hat{s}|h)] \\ &> \delta_i^{t(h)} \varepsilon' \end{aligned}$$

and consequently, the profile \hat{s} is not an ex ante perfect $\delta_i^{t(h)} \varepsilon'$ -equilibrium. Choosing $\varepsilon = \min_i \{\delta_i^{(T)} \varepsilon'\}$, where T is the length of the game, shows that the profile is also not a contemporaneous perfect ε -equilibrium. ||

This result appears to conflict with Radner's demonstration that there exist ex ante perfect ε -equilibria featuring cooperation in the finitely repeated prisoners' dilemma *for arbitrarily small* ε . However, Radner's result is achieved by allowing the number of

is certainly in the spirit of contemporaneous perfect ε -equilibrium. However, Radner works with payoffs that are undiscounted averages of the payoffs received in the current and future periods. Consider strategies that cooperate until some period t , and defect thereafter. When considering a deviation in period t , a player evaluates the average of all subsequent payoffs. As long as t is not too close to the end of the game, it is then still ε -optimal by this criterion to cooperate in period t , yielding equilibria featuring cooperation in some (possibly very long) sequence of initial periods followed by subsequent defection.

periods T to grow sufficiently rapidly, as ε falls, that the ex ante value of foregoing defection in period T remains always below ε .

Proposition 1 implies that the sets of ex ante perfect ε -equilibria and contemporaneous perfect ε -equilibria can differ for arbitrarily small ε only in infinite horizon games. The next example illustrates this possible difference:

Example 2: An infinite game. Consider a potential surplus whose contemporaneous value in time t is given by $2\delta^{-t}$ for some $\delta \in (0, 1)$. In each period, two agents simultaneously announce either *take* or *pass*. The game ends with the first announcement of *take*. If this is a simultaneous announcement, each agent receives a contemporaneous payoff of $\frac{1}{2}(2\delta^{-t} - 1)$. We can think of this as the agents splitting the surplus, after paying a cost of 1. If only one agent announces *take*, then that agent receives $\frac{1}{2}2\delta^{-t}$, while the other agent receives nothing. Hence, a single *take* avoids the cost, but provides a payoff only to the agent doing the taking. The agents' (common and constant) discount factor is given by δ .

This game has a unique pure-strategy contemporaneous perfect ε -equilibrium, in which both players *take* in the first period, for any $\varepsilon < \frac{\delta}{2}$. To verify this, suppose that both agents' strategies stipulate that they *take* in period $t > 0$. Then the period $t - 1$ contemporaneous payoff gain to playing *take* in period $t - 1$ is given by

$$\left(\frac{1}{2}2\delta^{-(t-1)}\right) - \delta \left(\frac{1}{2}(2\delta^{-t} - 1)\right) = \frac{\delta}{2} > \varepsilon. \quad (5)$$

Hence, a simultaneous *take* can appear only in the first period. If the first play of *take* occurs in any period $t > 0$ and is a *take* on the part of only one player, then it is a superior (contemporaneous) response for the other player to *take* in the previous period, since

$$\frac{1}{2}2\delta^{-(t-1)} - 0 > \varepsilon.$$

The only possible pure-strategy contemporaneous equilibrium thus calls for both agents to *take* in every period. It remains only to verify that such strategies are a best reply, which follows from the observation that

$$\frac{1}{2}(2\delta^0 - 1) > 0.$$

A straightforward variation on this argument shows that the only (pure or mixed) Nash (and hence subgame perfect) equilibrium outcome of the game also calls for both agents to *take* in the first period.

In contrast, let τ satisfy $\delta^\tau < 2\varepsilon$. Given any such τ , there exists a pure-strategy ex ante perfect ε -equilibrium in which both players *pass* in every period $t < \tau$ and *take* in

every period $t \geq \tau$. In particular, the most profitable deviation for either player is to choose *take* in period $\tau - 1$, for an ex ante payoff increment of

$$\delta^{\tau-1} \frac{1}{2} 2\delta^{\tau-1} - \delta^{\tau} \frac{1}{2} (2\delta^{\tau} - 1) = \frac{1}{2} \delta^{\tau},$$

which by construction is smaller than ε . ||

In contrast to Proposition 1, Example 2 shows that in infinite games, ex ante and contemporaneous perfect ε -equilibria can be quite different for arbitrarily small ε . Fudenberg and Levine (1983, p. 261) introduce a condition under which sufficiently distant future periods are relatively unimportant, making infinite games “approximately” finite:

Definition 2 *The game is continuous at infinity if for all i ,*

$$\lim_{t \rightarrow \infty} \sup_{\substack{s, s', h \text{ s.t.} \\ t=t(h)}} \left| \delta_i^{(t(h))} [\pi_i(s|h) - \pi_i(s'|h)] \right| = 0.$$

Equivalently, a game is continuous at infinity if two strategy profiles give nearly the same payoffs when they agree on a sufficiently long finite sequence of periods. A sufficient condition for continuity at infinity is that the reward function $r_i(a^t(s))$ be bounded and the players discount.

Fudenberg and Levine’s (1983) Lemma 3.2 can be easily adapted to give:³

Proposition 2 *In a game that is continuous at infinity, every converging (in the product topology on the set of strategy profiles) sequence of ex ante perfect $\varepsilon(n)$ -equilibria (and hence every converging sequence of contemporaneous perfect $\varepsilon(n)$ -equilibria) with $\varepsilon(n) \rightarrow 0$ converges to a subgame perfect equilibrium.*

Proof. We argue to a contradiction. Suppose $\{\sigma(n)\}$ is a sequence of ex ante perfect $\varepsilon(n)$ -equilibria, where $\varepsilon(n) \rightarrow 0$, converging to a strategy $\hat{\sigma}$ that is not a subgame perfect equilibrium. Because $\hat{\sigma}$ is not a subgame perfect equilibrium, there exists an information set h for player i , strategy σ_i and $\gamma > 0$ such that

$$\pi_i(\hat{\sigma}_{-i}^h, \sigma_i^h) = \pi_i(\hat{\sigma}^h) + \gamma \tag{6}$$

³See Fudenberg and Levine (1986) for extensions of Fudenberg and Levine’s (1983) results and Borgeers (1989) and Harris (1985) for related work. Fudenberg and Levine (1983, Theorem 3.3) go further to show that an equilibrium σ is subgame perfect in an infinite-horizon game if and only if there is a sequence $\sigma(n)$ of ex ante perfect $\varepsilon(n)$ -equilibria in finite-horizon truncations of length $T(n)$ of the original game, with $\sigma(n) \rightarrow \sigma$, $\varepsilon(n) \rightarrow 0$, and $T(n) \rightarrow \infty$. The fact that, for sufficiently small ε , the finitely-repeated prisoners’ dilemma (of *any* length) has a unique contemporaneous perfect ε -equilibrium, regardless of length, shows that the same is not true for contemporaneous perfect ε -equilibrium.

while $\sigma(n)$ must be an ex ante perfect $\gamma/4$ -equilibrium for all sufficiently large n , requiring

$$\pi_i(\sigma_{-i}^h(n), \sigma_i^h) \leq \pi_i(\sigma^h(n)) + \frac{\gamma}{4}. \quad (7)$$

Because the game is continuous at infinity, we can find n sufficiently large that⁴

$$\left| \pi_i(\hat{\sigma}_{-i}^h, \sigma_i^h) - \pi_i(\sigma_{-i}^h(n), \sigma_i^h) \right| < \frac{\gamma}{4}$$

and

$$\left| \pi_i(\hat{\sigma}^h) - \pi_i(\sigma^h(n)) \right| < \frac{\gamma}{4}.$$

Combining with (7), this gives

$$\pi_i(\hat{\sigma}_{-i}^h, \sigma_i^h) \leq \pi_i(\hat{\sigma}^h) + \frac{3\gamma}{4},$$

contradicting (6). ||

Example 2 shows that in games that are not continuous at infinity, Proposition 2 does not hold for ex ante perfect ε -equilibria. The following example shows that, without continuity at infinity, it also need not hold for contemporaneous perfect ε -equilibria:

Example 3. A single player, after every nonterminal history, chooses between L and R . The player discounts future payoffs at constant rate $\delta_{it} = \delta \in (0, 1)$. A choice of R in period t ends the game with a period- t reward of $\delta^{-t} - \delta^t$. A choice of L leads to the next period. For any ε , it is a contemporaneous perfect ε -equilibria to choose L in every period t for which $t \leq \ln \varepsilon / \ln \delta$ (i.e., every period in which $\delta^t \leq \varepsilon$) and R in every period for which $t > \ln \varepsilon / \ln \delta$.⁵ However, as ε goes to zero, the sequence of such equilibria converges to always choosing L , which is not a subgame-perfect equilibrium. Instead, this game has no subgame perfect equilibrium. ||

In a finite game, ex ante and contemporaneous perfect ε -equilibria coincide for sufficiently small ε (Proposition 1). The observation that any ex ante perfect ε -equilibria is also a contemporaneous perfect ε -equilibria, together with Proposition 2's convergence result for both concepts, raises the possibility that the following counterpart of this finite-horizon equivalence might hold for infinite games that are continuous at infinity: for every ε there is an $\hat{\varepsilon}(\varepsilon) \geq \varepsilon$ such that every ex ante perfect ε equilibrium is a contemporaneous perfect $\hat{\varepsilon}(\varepsilon)$ equilibrium, with $\lim_{\varepsilon \rightarrow 0} \hat{\varepsilon}(\varepsilon) = 0$. However, this is not the case, as the following example illustrates.

⁴Intuitively, by choosing n sufficiently large, we can make all behavior differences arbitrarily small except those that are discounted so heavily as to have an arbitrarily small effect.

⁵For any period $\tau \leq t$, choosing R in period τ gives a payoff (evaluated in period τ) of $\delta^{-\tau} - \delta^\tau$, which is no larger than the payoff $\delta^{t-\tau}(\delta^{-t} - \delta^t)$ of adhering to the equilibrium strategy. In any period $\tau > t$, choosing R also gives a payoff $\delta^{-\tau} - \delta^\tau$, while waiting until some later period t' to choose R gives a payoff of $\delta^{t'-\tau}(\delta^{-t'} - \delta^{t'})$, which exceeds the former by less than ε when $\delta^\tau \leq \varepsilon$.

Example 4. As in Example 3, a single player, after every nonterminal history, chooses between L and R . The player discounts future payoffs at constant rate $\delta_{it} = \delta \in (0, 1)$. A choice of R in period t ends the game with a period- t reward of -1 . A choice of L leads to the next period. Since $\delta < 1$ and payoffs are bounded, the game is continuous at infinity. For any ε and $\tau \geq \ln \varepsilon / \ln \delta$, it is an ex ante perfect ε -equilibrium to choose L for all $t < \tau$ and R for all $t \geq \tau$. However, for $\varepsilon < 1$, the only contemporaneous perfect ε -equilibrium is to choose L at every opportunity (which is also the subgame perfect equilibrium). We thus have ex ante perfect ε -equilibria for arbitrarily small ε that are contemporaneous perfect $\hat{\varepsilon}$ -equilibria only for large (> 1) values of $\hat{\varepsilon}$. \parallel

4. Nearby Games

For the remainder of the paper, we fix the game form and the discounting scheme and identify games with their associated sequence of reward functions. In this view, two games are close if the reward functions are close.⁶ Formally, we define two metrics on games:

$$d_N(G, \hat{G}) = \sup_{i,h} \left| \frac{r_i(h) - \hat{r}_i(h)}{t(h)} \right|$$

and

$$d_P(G, \hat{G}) = \sup_{i,h} |r_i(h) - \hat{r}_i(h)|. \quad (8)$$

Let r_i^k and r_i be player i 's reward functions in G^k and G respectively. The following lemma is an immediate consequence of the definitions:

Lemma 1

(1.1) Suppose that, for a sequence of games $\{G^k\}$ and game G , $\lim_{k \rightarrow \infty} d_N(G^k, G) = 0$ and there is $M \in \mathbb{R}$ such that the associated reward functions $\{r^k\}$ and r take values in $[-M, M]$. Then

$$\sup_{i,\sigma} \left| \pi_i^{G^k}(\sigma) - \pi_i^G(\sigma) \right| \rightarrow 0.$$

(1.2) Suppose that, for a sequence of games $\{G^k\}$ and game G , $\lim_{k \rightarrow \infty} d_P(G^k, G) = 0$. Then

$$\sup_{i,\sigma,h} \left| \pi_i^{G^k}(\sigma|h) - \pi_i^G(\sigma|h) \right| \rightarrow 0.$$

⁶More generally, we might define two games to be close if their game forms, discounting schemes, and reward functions are close. Börgers (1991, p. 95) introduces a such a measure, defined in terms of the game form and the payoffs $\pi(\sigma)$. Given our interest in the implications of different timing of rewards for ε -optimization, it is most revealing to fix the game form and discounting scheme while examining perturbations of the reward function.

Convergence under d_P is equivalent to uniform convergence of the reward functions. Given the assumed bound on payoffs in the Lemma, convergence under d_N is equivalent to pointwise convergence of the reward functions. Without this bound, d_N implies, but is stronger than, pointwise convergence.

5. Approximating Equilibria in Nearby Games

It is straightforward that, for static games, ε -Nash equilibria of a given game G approximate Nash equilibria of nearby games. A similar result holds for multistage games (recall that D , the bound from (1), does not depend on the reward function). The only complication in extending the observation from static to multistage games is that our notions of closeness for games examine the reward functions, while optimality is based on the discounted sums of rewards.

Since players discount and the concept of a Nash equilibrium depends only on ex ante payoffs, under a slight strengthening of (1), it is not necessary for the result that the rewards be uniformly (in t) close (as required by d_P):

Proposition 3 *Fix a game G .*

(3.1) *If the strategy profile $\hat{\sigma}$ is a Nash equilibrium of game G' with $d_P(G', G) < \varepsilon/2$, then $\hat{\sigma}$ is an εD -Nash equilibrium of game G . Moreover, if for each i*

$$\limsup_{t \rightarrow \infty} \left(\prod_{\tau=0}^t \delta_{i\tau} \right)^{\frac{1}{t}} < 1, \quad (9)$$

then there exists D' (independent of the reward function of G) such that if the strategy profile $\hat{\sigma}$ is a Nash equilibrium of game G' with $d_N(G', G) < \varepsilon/2$, then $\hat{\sigma}$ is an $\varepsilon D'$ -Nash equilibrium of game G .

(3.2) *If $\hat{\sigma}$ is a pure-strategy ε -Nash equilibrium of game G , then there exists a game G' with $d_P(G', G) < \frac{\varepsilon}{2}$ (and hence $d_N(G', G) < \varepsilon/2$) for which $\hat{\sigma}$ is a Nash equilibrium.*

The proof of this proposition follows that of the next proposition.

The restriction to pure strategy equilibria cannot be dropped in Proposition 3.2. For example, in the game,

	L	R
T	$\begin{array}{ c } \hline 0, 0 \\ \hline \end{array}$	$\begin{array}{ c } \hline 1, 1 \\ \hline \end{array}$
B	$\begin{array}{ c } \hline 1, 0 \\ \hline \end{array}$	$\begin{array}{ c } \hline 2, 1 \\ \hline \end{array}$

the strategy profile $((\varepsilon \circ T + (1 - \varepsilon) \circ B), R)$ is an ε -equilibrium. However, in any Nash equilibrium of any game $\varepsilon/2$ -close to this game, player 1 must choose B with probability 1. The “problem” mixed strategies are those that, as in the example, put small

probability on an action that is far from optimal.⁷

Returning to Example 2, it is straightforward that in all games that are sufficiently close, as measured by d_P , the strict inequality in (5) continues to hold. Hence, there is a unique subgame perfect equilibrium in such games, in which both players immediately play *take*. Contrasting this observation with the variety of ex ante perfect ε equilibria that appear in Example 2, we are led to the conclusion that if one seeks an approximate equilibrium concept capturing subgame perfect equilibria of nearby games, contemporaneous perfection is the appropriate concept:

Proposition 4 *Fix a game G .*

(4.1) *If the strategy profile $\hat{\sigma}$ is a subgame perfect equilibrium of game G' with $d_P(G', G) < \varepsilon/2$, then $\hat{\sigma}$ is a contemporaneous perfect εD -equilibrium of game G .*

(4.2) *If $\hat{\sigma}$ is a pure-strategy contemporaneous perfect ε -equilibrium of game G , then there exists a game G' with $d_P(G', G) < \varepsilon/2$ for which $\hat{\sigma}$ is a subgame perfect equilibrium.*

The first statement of Proposition 4 guarantees that the contemporaneous perfect ε -equilibria of a game include all subgame-perfect equilibria of nearby games. The second guarantees that *every* pure strategy contemporary perfect ε -equilibrium is a subgame perfect equilibrium of a nearby game. We emphasize again, however, that neither Proposition 3 nor Proposition 4 contains an if and only if result, since the ε 's for the two parts are not the same. We could replace εD in the first statements with ε , making the two statements symmetric, if we had also replaced d_P with the metric

$$\tilde{d}_P(G, \hat{G}) = \sup_{s, h, i} \left| \sum_{t=t(h)}^{\infty} \prod \delta_i^{(t(h), t)} r_i(a^t(s^h)) - \sum_{t=t(h)}^{\infty} \prod \delta_i^{(t(h), t)} \hat{r}_i(a^t(s^h)) \right|.$$

Since d_P yields a more transparent notion of closeness for games in terms of the reward function, we have chosen to work with d_P rather than \tilde{d}_P .

Proof of Proposition 4. (4.1) Let $\hat{\sigma}$ be a subgame perfect equilibrium of game G' with $d_P(G', G) < \varepsilon/2$. It follows from (1) and (2) that, for any strategy profile σ , player i and history h ,

$$|\pi_i^G(\sigma|h) - \pi_i^{G'}(\sigma|h)| < \frac{\varepsilon}{2} \sum_{t=t(h)}^{\infty} \prod_{\tau=t(h)}^t \delta_{i\tau} \leq \frac{\varepsilon}{2} D. \quad (10)$$

⁷A referee noted that, if we replace player 1's mixture in this game with a node at which Nature first draws a real number from $[0, 1]$ and then player 1 chooses T for draws less than ε and B for higher draws, then these “problem” strategy profiles are no longer ε -Nash equilibria. This suggests that Proposition 3 could be formulated as an equivalence if we insisted on purifying mixed strategies in this way. Doing so raises the inconvenience of dealing with infinite numbers of finite-length histories.

We then have, for any player i , history h and strategy σ_i ,

$$\begin{aligned}
& \pi_i^G(\hat{\sigma}|h) - \pi_i^G(\hat{\sigma}_{-i}, \sigma_i|h) \\
&= \left(\pi_i^G(\hat{\sigma}|h) - \pi_i^{G'}(\hat{\sigma}|h) \right) + \left(\pi_i^{G'}(\hat{\sigma}|h) - \pi_i^{G'}(\hat{\sigma}_{-i}, \sigma_i|h) \right) \\
&\quad + \left(\pi_i^{G'}(\hat{\sigma}_{-i}, \sigma_i|h) - \pi_i^G(\hat{\sigma}_{-i}, \sigma_i|h) \right) \\
&\geq -\varepsilon D,
\end{aligned}$$

giving the result.

(4.2) Let \hat{s} be a pure strategy contemporaneous perfect ε -equilibrium of G . For notational purposes, assume that $A_i(h)$ and $A_i(h')$ are disjoint for all h and $h' \in H_i^*$, so that the action a_i uniquely identifies a history. For all information sets $h \in H_i^*$ for player i , $\hat{s}_i^{a_i}$ denotes the strategy that agrees with \hat{s}_i at every information set other than h , and specifies the action a_i at h . In other words, $\hat{s}_i^{a_i}$ is the one-shot deviation

$$\hat{s}_i(h') = \begin{cases} \hat{s}_i(h'), & \text{if } h' \neq h, \\ a_i, & \text{if } h' = h. \end{cases}$$

Since \hat{s} is a contemporaneous perfect ε -equilibrium,⁸

$$\gamma_i(h) \equiv \max_{a_i \in A_i(h)} \pi_i(\hat{s}_{-i}, \hat{s}_i^{a_i}|h) - \pi_i(\hat{s}_{-i}, \hat{s}_i|h) < \varepsilon.$$

The idea in constructing the perturbed game is to increase the reward to player i from taking the specified action $\hat{s}_i(h)$ at his information set $h \in H_i^*$ by $\gamma_i(h)$, and then lowering all rewards by $\varepsilon/2$. (Discounting guarantees that subtracting a constant from every reward still yields well-defined payoffs.) However, care must be taken that the one-shot benefit takes into account the other adjustments. So, we construct a sequence of games as follows.

For fixed T , we define the adjustments to the rewards at histories of length less than or equal to T , $\gamma_i^T(h)$. The definition is recursive, beginning at the longest histories, and proceeding to the beginning of the game. For information sets $h \in H_i^*$ satisfying $t(h) = T$, set

$$\gamma_i^T(h) = \gamma_i(h).$$

Now, suppose $\gamma_i^T(h'')$ has been determined for all $h'' \in H_i^*$ satisfying $t(h'') = \ell \leq T$. For h' satisfying $t(h') = \ell + 1$, define

$$r_i^T(h') \equiv \begin{cases} r_i(h'', \hat{s}(h'')) + \gamma_i^T(h''), & \text{if } h' = (h'', \hat{s}(h'')) \text{ for some } h'' \in H_i^* \\ & \text{such that } t(h'') = \ell, \\ r_i(h'), & \text{otherwise.} \end{cases}$$

⁸Since $s_i(h) \in A_i(h)$, $\gamma_i(h) \geq 0$.

This then allows us to define for h satisfying $t(h) = \ell - 1$,⁹

$$\pi_i^T(s|h) \equiv r_i(a^{t(h)}(s^h)) + \sum_{t=t(h)+1}^{\infty} \left(\prod_{\tau=t(h)+1}^t \delta_{i\tau} \right) r_i^T(a^t(s^h)),$$

and

$$\gamma_i^T(h) \equiv \max_{a_i \in A_i(h)} \pi_i^T(\hat{s}_{-i}, \hat{s}_i^{a_i}|h) - \pi_i^T(\hat{s}_{-i}, \hat{s}_i|h).$$

Proceeding in this way determines $r_i^T(h)$ for all h .

We claim that for any T , $r_i^T(h) - r_i(h) < \varepsilon$. To see this, recall that \hat{s} is a contemporaneous perfect ε -equilibrium and note that the adjustment at any h can never yield a continuation value (under \hat{s}) larger than the maximum continuation value at h .

Moreover, the sequence $\{r_i^T\}_T$ of reward functions has a convergent subsequence (there is a countable number of histories, and for all $h \in H_i^*$, $r_i^T(h) \in [r_i(h), r_i(h) + \varepsilon]$). Denote the limit by r_i^* . Note that there are no profitable one-shot deviations from \hat{s} under r_i^* , by construction. As a result, because of discounting, \hat{s} is subgame perfect.

Finally, we subtract $\varepsilon/2$ from every reward. Equilibrium is unaffected, and the resulting game is within $\varepsilon/2$ under d_P . ||

Proof of Proposition 3. The proofs of the statements about d_P are the same arguments as in the proof of Proposition 4, but applied only to the initial history.

Suppose now that the discounting scheme satisfies (9). From the root test ((1976, Theorem 3.33)), there exists D' such that

$$\sum_{t=0}^{\infty} t \prod_{\tau=0}^t \delta_{i\tau} < D'.$$

The proof of the d_N result is now again a special case of that of Proposition 4.1, with the exception that the first inequality in (10) in the statement of that proposition is now replaced by

$$|\pi_i^G(\sigma|h) - \pi_i^{G'}(\sigma|h)| < \varepsilon \sum_{t=0}^{\infty} t \prod_{\tau=0}^t \delta_{i\tau},$$

which is less than $\varepsilon D'$. ||

6. Discussion

The set of contemporaneous perfect ε -equilibria of a game G includes the set of subgame perfect equilibria of nearby games. Examining contemporaneous perfect ε -equilibria

⁹Note that the history $a^{t(h)+1}(s^h) = a^\ell(s^h)$ is of length $\ell + 1$.

thus ensures that one has not missed any subgame perfect equilibria of the real games that might correspond to the potentially misspecified model.

Examples 2 and 3 show that, for games that are not continuous at infinity, examining either ex ante perfect ε -equilibria or subgame perfect equilibria (respectively) can give a misleading picture of the set of contemporaneous perfect ε -equilibria of a game, and hence subgame perfect equilibria of nearby games, including too many equilibria in the first case and too few in the second. Suppose, however, that we restrict attention to games that are continuous at infinity. As ε gets small, the set of ex ante perfect ε -equilibria and the set of contemporaneous perfect ε -equilibria of the model converge to the set of subgame perfect equilibria of the model.

In light of this, why not simply dispense with ε altogether and examine subgame perfect equilibria of the model? Since our goal is to ensure that no subgame perfect equilibrium from the real game is neglected, we would need that *every* subgame perfect equilibrium of the real game is close to some subgame perfect equilibrium of close-by models. That is, suppose a modeler, after fixing $\varepsilon > 0$, postulates his best-guess model of the real game and calculates its subgame perfect equilibria. Can we be assured that, if the real game is ε -close to the model, all of its subgame perfect equilibria will be captured? No, as the following simple example illustrates: Suppose the model has two choices with player 1 receiving 0 from L and $\varepsilon/2$ from R . The only subgame perfect equilibrium is to play R . However, if in the true game, player 1 receives $\varepsilon/2$ after L and 0 after R , the only subgame perfect equilibrium is L . In contrast, from Proposition 4, we in fact know that *every* subgame perfect equilibrium of the real game is in fact a contemporaneous perfect ε -equilibrium of the model.

References

- BÖRGER, T. (1989): “Perfect Equilibrium Histories of Finite and Infinite Horizon Games,” *Journal of Economic Theory*, 47, 218–227.
- (1991): “Upper Hemicontinuity of the Correspondence of Subgame-Perfect Equilibrium Outcomes,” *Journal of Mathematical Economics*, 20(1), 89–106.
- FUDENBERG, D., AND D. K. LEVINE (1983): “Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games,” *Journal of Economic Theory*, 31, 251–268.
- (1986): “Limit Games and Limit Equilibria,” *Journal of Economic Theory*, 38, 261–279.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, MA.
- HARRIS, C. (1985): “A Characterization of the Perfect Equilibria of Infinite Horizon Games,” *Journal of Economic Theory*, 37, 99–125.

- KOHLBERG, E., AND J.-F. MERTENS (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54, 1003–1037.
- RADNER, R. (1980): “Collusive behaviour in noncooperative epsilon-equilibria of oligopolies with long but finite lives,” *Journal of Economic Theory*, 22, 136–154.
- RUBINSTEIN, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50, 97–109.
- RUDIN, W. (1976): *Principles of Mathematical Analysis*. McGraw-Hill, New York, 3rd edn.
- WATSON, J. (1994): “Cooperation in the Infinitely Repeated Prisoners’ Dilemma with Perturbations,” *Games and Economic Behavior*, 7(2), 260–285.



ABSTRACT, DOCUMENT,
AUTHOR CONTACT INFORMATION

[SSRN Home](#)

[Search Abstracts](#)

[Subscription Request Form](#)

The paper you requested was not found. Please check your search criteria and try again.

[Return to search page](#)

Copyright © 2003 Social Science Electronic Publishing, Inc. All Rights Reserved