PIER Working Paper 02-004

"Contemporaneous Perfect Epison-Equilibria"

by

George J. Mailath, Andrew Postlewaite and Larry Samuelson

http://ssrn.com/abstract_id=305894

Contemporaneous Perfect Epsilon-Equilibria*

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

February 11, 2002.

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 not implied by 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 game achieved by perturbing payoffs by at most $\varepsilon/2$, 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 problems, 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.

If an analyst would like to ensure that an equilibrium of the model is close to a Nash equilibrium of *every* possible game with nearly the same payoffs, then the appropriate solution concept in the model is some version of strategic stability (Kohlberg and Mertens [6]). 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 2 below) that a strategy profile is an ε-Nash equilibrium of a game if it is a Nash equilibrium of a game with nearly the same payoffs.

[†]University of Pennsylvania, gmailath@econ.upenn.edu.
[‡]University of Pennsylvania, apostlew@econ.upenn.edu.
[§]University of Wisconsin, larrysam@ssc.wisc.edu.

When dealing with extensive-form games, sequential rationality is typically imposed. Radner [7] 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 evaluation of the gain from deviating from the proposed equilibrium strategy is done ex ante, that is, prior to the start 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 can be an ex ante perfect ε-equilibrium 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 back 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 for every history, every player is within ε of his optimal payoff, where the evaluation of strategies is made contemporaneously, that is, evaluations are made at the point that the 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 [4, ch. 4]). There is a potentially infinite number of periods $\{0, 1, \ldots, \}$. In each period, players simultaneously choose from feasible sets of actions knowing the history of the game through the preceding period. The feasible sets may depend upon the period and the history of play, and may constrain some players to make only the trivial choice "do nothing." This class of games includes both repeated games and dynamic games like Rubinstein's [8] 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 histories may correspond to terminal nodes of the game (i.e., the game is over after that history, such as 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

$A_i(h)$. 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. 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, and the consequences of decisions taken at any point in time may be felt at that time or later. Players can reach agreement in an alternating-offer bargaining game more or less quickly, they choose when to end the centipede game, and in each period of the repeated oligopoly game analyzed by Radner, players decide whether to continue cooperating or to defect. We are concerned with the differences between the time decisions are made and the time the consequences of those decisions are realized. We capture this distinction via discounting and the reward function $r_i : H \to \Re$, 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 payoffs to period $t-1$ using the factor $\delta_{it} \in (0,1]$. Define $\delta_i^{(t',t'')} \equiv \prod_{\tau=t'+1}^{t''} \delta_{it}$; 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 $\delta_i^{(0,t)}$. Finally, we set $\delta_{i0} = \delta_i^{(t,t)} = 1$. We allow for different discount factors in different periods 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.

The set of pure strategies is denoted Σ, the outcome path induced by the pure strategy $s \in \Sigma$ is denoted $a^\infty(s)$, and its first $t+1$ periods is denoted $a^t(s)$. Player i's payoff function, $\pi_i : \Sigma \to \Re$, is given by

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

We assume this expression is well-defined for all $s \in \Sigma$ and we assume that players discount, in the sense that

$$\sum_{t=0}^{\infty} \prod_{\tau=0}^{t} \delta_{i\tau} < \infty.$$

We extend π_i to the set of behavior strategies, 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, and in games without discounting such as the standard centipede

game, $\delta_i = 1$ and $r_i(h)$ equals i's payoff when the game is stopped in period $t(h)$ under h, and zero otherwise.

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[1]

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

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 $t(h)$.

Fudenberg and Levine [3, p. 261] introduce the following:

Definition 1 *The game is* continuous at infinity *if for all i,*

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

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.

3. Epsilon Equilibria

The strategy σ^h induces 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}, $\pi_i(\sigma^h_{-i}, \cdot)$ 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 $\pi_i(\sigma_{-i}, \cdot|h)$.

---

[1]Recall that $\delta_i^{(t,t)} = 1$.

Definition 2 *For $\varepsilon > 0$, a strategy profile $\hat{\sigma}$ is an ε-Nash equilibrium if, for each player i and strategy σ_i,*

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

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

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

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 [7] and Fudenberg and Levine [3].[2] It is easy to see that any contemporaneous perfect ε-equilibria is an ex ante perfect ε-equilibrium. The two concepts coincide in the absence of discounting or when $\varepsilon = 0$, while Radner's ε-equilibria in a repeated oligopoly with discounting are ex ante but not contemporaneous ε-equilibria. Note that if $\varepsilon = 0$, an ex ante or contemporaneous perfect ε-equilibrium is subgame perfect.

Fudenberg and Levine's [3] Lemma 3.2 can be easily adapted to give:[3]

Proposition 1 *In a game that is continuous at infinity, every converging sequence of ex ante perfect $\varepsilon(n)$-equilibria 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 $\varepsilon(n) \to 0$, 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

$$\pi_i(\hat{\sigma}^h_{-i}, \sigma^h_i) = \pi_i(\hat{\sigma}^h) + \gamma \tag{2}$$

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

$$\pi_i(\sigma^h_{-i}(n), \sigma^h_i) \leq \pi_i(\sigma^h(n)) + \frac{\gamma}{4}. \tag{3}$$

Because the game is continuous at infinity, we can find n sufficiently large that[4]

$$\left| \pi_i(\hat{\sigma}^h_{-i}, \sigma^h_i) - \pi_i(\sigma^h_{-i}(n), \sigma^h_i) \right| < \frac{\gamma}{4}$$

---

[2]Watson [9] considers an intermediate concept that requires ε-optimality conditional only on those histories that are reached along the equilibrium path.

[3]See also Harris [5], Fudenberg and Levine [2], and Borgers [1] for extensions of Fudenberg and Levine's results.

[4]Intuitively, 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.

and

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

Combining with (3), this gives

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

contradicting (2). ‖

4. Nearby Games

We now formalize two senses in which games with the same game form can be close. Once we fix a game form, two games can only differ in payoffs. Since we take the reward functions and discounting as primitives, it is natural to view two games as being close if the discount factors and reward functions are close. Define

$$d_N(G, \hat{G}) = \sup_{i.t} \left|\frac{\delta_{it} - \hat{\delta}_{it}}{t}\right| + \sup_{i,h} \left|\frac{r_i(h) - \hat{r}_i(h)}{t(h)}\right|$$

and

$$d_P(G, \hat{G}) = \sup_{i.t} \left|\delta_{it} - \hat{\delta}_{it}\right| + \sup_{i,h} \left|r_i(h) - \hat{r}_i(h)\right|.$$

Convergence under d_N is equivalent to pointwise convergence of the discount factors and reward functions, while convergence under d_P is equivalent to uniform convergence of the discount factors and reward functions.

The following lemma is an immediate consequence of continuity at infinity:

Lemma 1 *1. If G is continuous at infinity and $G^k \to G$ under d_N, then*

$$d_N^*(G^k, G) \equiv \max_{i,\sigma} \left|\pi_i^{G^k}(\sigma) - \pi_i^G(\sigma)\right| \to 0.$$

2. If $G^k \to G$ under d_P, then

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

While d_N^* and d_P^* are pseudometrics on the space of games (as defined by discount factors and reward functions), they need not separate points and hence they are not metrics. This is somewhat of a departure from standard practice in noncooperative game theory: Typically, if different strategic environments have the same game form with the same payoffs on that game form, one would say that these different strategic environments are represented by the same game. However, because the timing of decisions and receipt of rewards plays a central role in our analysis, we take a more "structural" view, discriminating between games with the same implied payoffs π.

5. Approximating Equilibria in Nearby Games

It is straightforward that ε-Nash equilibria approximate Nash equilibria of nearby games:

Proposition 2 *The strategy profile σ^* is an ε-Nash equilibrium of game G if there is a game G' such that $d_N(G', G) \leq \varepsilon/2$ for which σ^* is a Nash equilibrium. If σ^* is pure, then the implication is if and only if.*

Proof. Let σ^* be a Nash equilibrium of game G', so that $\pi_i^{G'}(\sigma^*) \geq \pi_i^{G'}(\sigma^*_{-i}, \sigma_i)$. Then as long as payoffs in G are within $\varepsilon/2$ of those in G', we must have $\pi_i(\sigma^*) \geq \pi_i(\sigma^*_{-i}, \sigma_i) - \varepsilon$ in game G, and hence σ^* is an ε-Nash equilibrium of game G. Conversely, let s^* be a pure strategy ε-equilibrium of game G. Construct a new game G' by reducing player i's payoff by $\varepsilon/2$ at any strategy profile (s^*_{-i}, s_i) where $s_i \neq s_i^*$, and increasing i's payoff at (s^*_{-i}, s_i^*) by $\varepsilon/2$.[5] Then s^* will be a Nash equilibrium in G', where $d_N(G, G') = \varepsilon/2$. ‖

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

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

,

the strategy profile $\left((\varepsilon \circ T + (1-\varepsilon) \circ B), R \right)$ is an ε-equilibrium. However, in any Nash equilibrium of in 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.

The relationship between equilibria of nearby games is more complicated when dealing with perfect ε-equilibria. We first use a simplified version of Radner's original analysis to show that ex ante and contemporaneous perfect ε-equilibria for the same value of ε can be quite different.

Example 1: The finitely repeated prisoners' dilemma. Radner [7] used the notion of ex ante perfect ε-equilibrium to examine finitely repeated oligopoly games, which we simplify to the finitely repeated prisoners' dilemma. Let the stage game be

	C	D
C	$2,2$	$-1,3$
D	$3,-1$	$0,0$

.

---

[5]For example, if $\delta_i^{(t,t+1)} \equiv \delta_i < 1$, then increase each $r_i\left(a^t(s^*_{-i}, s_i^*)\right)$ by $(1-\delta_i)\frac{\varepsilon}{2}$ and decrease $r_i\left(a^t(s^*_{-i}, s_i)\right)$ by $(1-\delta_i)\frac{\varepsilon}{2}$. A general algorithm for adjusting rewards is presented in the proof of Proposition 4.

Play this game $N+1$ times, with payoffs discounted according to the common discount factor $\delta < 1$. 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.[6]

The unique Nash (and hence subgame perfect) equilibrium in the finitely repeated prisoners' dilemma features perpetual defection. This example does not contradict Proposition 2, because for sufficiently small ε (in particular, so that (4) is violated), both players must defect in every period in any ε-equilibrium of a finitely repeated prisoners' dilemma. Radner exhibits a sequence of games featuring cooperation, for ever smaller values of ε, by continually lengthening the game as ε falls, so that the ex ante value of foregoing defection in period N declines sufficiently rapidly as to stay below ε. For fixed N and small ε, ex ante perfect ε-equilibria and contemporaneous perfect ε-equilibria coincide.

Indeed, in finite horizon games, for sufficiently small ε, the two concepts must coincide:

Proposition 3 *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 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.

There is some player i, history h and strategy s_i such that

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

---

[6]Radner [7, p. 153] defines an alternative notion of perfect ε-equilibrium in which the utility of a continuation strategy is calculated from the point of view of the period at which the decision is being made. This is certainly in the spirit of contemporaneous perfect ε-equilibrium. However, Radner works with payoffs that are undiscounted averages of the payoffs received in each period. 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 the prisoners' dilemma.

Since the game is finite, there exists an ε' sufficiently small such that, for all h, i, and $s_i \neq \hat{s}_i$,

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

But,

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

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

Hence it is only in infinite horizon games that the two concepts can differ for arbitrarily small ε. It is also worth noting that it is only for infinite horizon games that pointwise convergence of the reward functions is a strictly weaker notion than uniform convergence. We next present an example for which contemporaneous and ex ante perfect ε-equilibria differ for all $\varepsilon > 0$.

Example 2: An infinite game. Consider a potential surplus whose contemporaneous value in time t is given by δ^{-t} for some $\delta \in (0, 1)$. In each period, each agent can 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 $\frac{1}{2}(\delta^{-t} - 1)$. 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 $\frac{1}{2}\delta^{-t}$, 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' discount factor is given by δ.

This game has a unique Nash and contemporaneous perfect ε-equilibrium outcome, in which at least one player *takes* in the first period, for any $\varepsilon < \frac{\delta}{2}$. To verify this, suppose that both agents' equilibrium 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

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

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

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

9

However, for any period t for which $\delta^t < 2\varepsilon$, it is an ex ante perfect ε-equilibrium for both players to choose *take* in period t.

This example shows that there are ex ante perfect ε-equilibria that are not subgame perfect equilibria in any nearby game. If one wants to use ε-equilibria to approximate equilibria in nearby games, and one is concerned with perfection, then contemporaneous perfection is the appropriate concept:

Proposition 4 *The strategy profile $\hat{\sigma}$ is a contemporaneous perfect ε-equilibrium of game G if there is a game G' such that $d_P(G', G) \leq \varepsilon/2$ for which $\hat{\sigma}$ is a subgame perfect equilibrium. Moreover, if $\hat{\sigma}$ is pure, then the implication is if and only if.*

Proof. Mimicking the analogous argument from Proposition 2, it is immediate from the definition of d_P that a subgame perfect equilibrium of any game G' with $d_P(G', G) \leq \varepsilon/2$ is a contemporaneous perfect ε-equilibrium of G.

Conversely, let $\hat{s}$ be a pure strategy contemporaneous perfect ε-equilibrium of G. 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:

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

Since $\hat{s}$ is a contemporaneous perfect ε-equilibrium,[7]

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

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

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

---

[7] Since $s_i(h) \in A_i(h)$, $\gamma_i(h) \geq 0$.

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

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

This then allows us to define for h satisfying $t(h) = \ell - 1$,[8]

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

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 $\left\{ r_i^T \right\}_T$ 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. $\quad\quad ||$

6. Discussion

Epsilon Equilibria. The set of contemporaneous perfect ε-equilibria of a game G includes the set of subgame perfect equilibria of games whose payoffs are within $\varepsilon/2$ of G. Examining contemporaneous perfect ε-equilibria thus ensures that one has not missed any subgame perfect equilibria of the real games that might correspond to the potentially misspecified model. But as ε gets small, the set of contemporaneous perfect ε-equilibria of the model converges to the set of subgame perfect equilibria of the model. Why not just dispense with ε and examine the latter?

The difficulty is a lack of uniformity. No matter how small ε, one can find a model game whose set of subgame perfect equilibria does not contain an equilibrium close to some of the subgame perfect equilibria of a game whose payoffs lie within $\varepsilon/2$ of

---

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

those of the model.[9] An analyst who simply examines subgame perfect equilibria thus cannot be certain that no subgame perfect equilibria of the true game associated with the potentially misspecified model have been missed unless he can take ε to zero, i.e., unless he is certain that the model's payoffs are correct. But an analyst who is certain only that he knows payoffs with an error of $\varepsilon/2$ can examine contemporaneous perfect ε-equilibria of the model and be assured of identifying every subgame perfect equilibria of every game within $\varepsilon/2$ of the model.

Finite and infinite horizons. Fudenberg and Levine [3] consider a sequence of finite truncations of an infinite game that converge to the infinite game. They show that for any perfect equilibrium $\hat{\sigma}$ of the infinite game, there exists a sequence of ex ante $\varepsilon(n)$-perfect equilibria $\sigma(n)$ of the finitely truncated games that converge to $\hat{\sigma}$ as the length n of the finite games approaches infinity and as $\varepsilon(n) \to 0$.

It is easy to see that the analog of this result for contemporaneous ε-perfect equilibria does not hold. We need only consider Radner's example of the repeated prisoners' dilemma game. As Radner pointed out, trigger strategies that prescribe cooperation in each stage if and only if both players have cooperated in all previous periods will be an *ex ante ε-equilibrium* if the number of stages is sufficiently large relative to the discount rate. Hence, the equilibrium in the infinitely played prisoners' dilemma game consisting of such trigger strategies will be the limit of a sequence of ex ante ε-equilibria of finite truncations of the game with $\varepsilon(n) \to 0$.[10] However, it is clear that for small ε, the only contemporaneous ε-equilibria of the finitely played prisoners' dilemma game prescribe defection in all periods *regardless of the number of times the game is played*.

Bounded rationality. Bounded rationality on the part of the players in a game is often offered as a motivation for ε-equilibria. It is useful to explore this motivation in the context of our results.

Suppose first that players are able to compute precisely the consequences of their strategy choices, but for some reason satisfy themselves with strategies that get them within ε of their best response. Then consider the ex ante perfect ε-equilibrium in the finitely played prisoners' dilemma game. Player 1's calculations are predicated on Player 2 playing the given strategy. However, if 1 is forward-looking, he should recognize that 2 will have an incentive to deviate from the given trigger strategy at some point in the future when the gains from doing so are no longer small. Then how can we justify

---

[9]For example, the one-player game in which strategy 1 has a payoff of 0 and strategy 2 a payoff of $-\varepsilon$ has no equilibrium close to that which chooses the second strategy in the game in which both strategies yield a payoff of 0.

[10]In taking this limit, it is important that the length of the game grow sufficiently rapidly compared to the rate at which ε approaches zero. If ε approaches zero too rapidly, then all ex ante equilibria feature only defection.

Player 1's calculations? There are (at least) two possibilities. First, Player 1 might fail to be forward-looking, and hence fail to recognize the future deviation. Second, players' strategy choices may be irrevocable, so that player 2 will be unable to deviate in the future despite the nontrivial gain from doing so at that point in time. Each of these justifications raises questions. There is a tension in assuming that players are sufficiently rational to do the computations that assure that they are within ε of a best response, but irrational enough that looking forward is beyond them. If, as in the second case, players' strategy choices are irrevocable, it is difficult to see why one should be interested in perfection.

Contemporaneous ε-equilibrium avoids the schizophrenia implicit in a player who calculates the payoffs associated with strategy vectors in a sophisticated way but fails to be forward-looking. Also, in the words of Fudenberg and Levine [3], it captures the idea "that agents will not be misled by opponents' threats, but will instead compute their opponents' future actions from their knowledge of the structure of the game." However, contemporaneous perfect ε-equilibrium raises problems of its own. If player 1 faces no limitations on the calculations he is able to perform, why is 1 content with coming within ε of a best response? Doing the calculations required to verify ε-optimality also provides the information required to choose an optimal strategy, at which point one would expect player 1 to optimize.

Two possibilities again suggest themselves. First, it may be that players do not calculate at all; rather, an evolutionary process selects players who achieve relatively high payoffs at the expense of those who achieve lower payoffs. If the evolutionary process is subject to noise or drift, then it may only be capable of ensuring ε-optimality. Alternatively, if selection acts more forcefully on large payoff differences, the length of time required to squeeze the last ε of optimization out of the system may be so long as to make ε optimization what is likely to be observed. However, one would expect an evolutionary process to operate on ex ante rather than contemporaneous payoffs, calling the notion of contemporaneous perfect ε-equilibria into doubt. Secondly, it may again be that players cannot do the calculations required to assess strategies, but infer the potential payoffs of alternative strategies by observing the payoffs of others who play these strategies. Noise in observing the payoffs of others or difficulties in observing fine differences in payoffs would plausibly lead to ε-optimality. Again, however, one would expect this process to operate on ex ante payoffs much more effectively than contemporaneous payoffs.

Perhaps a more compelling motivation would be to assume that players make mistakes in their calculations, computing the payoff of each strategy to be its actual payoff plus the realization of a random variable, but then optimize given these mistaken calculations. Once again, however, it is not clear that this is compatible with the ability to flawlessly perform backward induction calculations.

References

[1] Tilman Borgers. Perfect equilibrium histories of finite and infinite horizon games. *Journal of Economic Theory*, 47:218–227, 1989.

[2] Drew Fudenberg and David Levine. Limit games and limit equilibria. *Journal of Economic Theory*, 38:261–279, 1986.

[3] Drew Fudenberg and David K. Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31:227–250, 1983.

[4] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

[5] Chris Harris. A characterization of the perfect equilibria of infinite horizon games. *Journal of Economic Theory*, 37:99–125, 1985.

[6] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1038, 1986.

[7] Roy Radner. Collusive behaviour in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22:136–154, 1980.

[8] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.

[9] Joel Watson. Cooperation in the infinitely repeated prisoners' dilemma with perturbations. *Games and Economic Behavior*, 7(2):260–285, September 1994.