Reprinted from JOURNAL OF ECONOMIC THEORY

Vol. 44, No. 1, February 1988
Printed in Belgium

Open-Loop and Closed-Loop Equilibria in Dynamic Games with Many Players*

DREW FUDENBERG† AND DAVID K. LEVINE‡

†*Department of Economics, University of California, Berkeley, California 94720 and*
‡*Department of Economics, University of California, Los Angeles, California 90024*

Received February 7, 1985; revised January 9, 1987

If players are small, one might expect that optimal reactions to one-player deviations are negligible, so that the open- and closed-loop equilibria are approximately the same. We investigate the circumstances in which this is true. *Journal of Economic Literature* Classification Numbers: 022, 026.

1. INTRODUCTION

The terms "open-loop" and "closed-loop" refer to two different information structures for multi-stage dynamic games. In the open-loop model, players cannot observe the play of their opponents; in the closed-loop model, all past play is common knowledge at the beginning of each stage. Open-loop and closed-loop equilibria are then the perfect equilibria corresponding to the two information structures. (Caution: This terminology is widespread but not universal. Some authors use "closed-loop equilibrium" to refer to all the *Nash* equilibria of the closed-loop model.) Open-loop equilibria are more tractable than closed-loop equilibria, because players need not consider how their opponents would react to deviations from the equilibrium path. For this reason, economists have sometimes analyzed open-loop equilibria, even when the closed-loop concept is more appropriate.

If players are small, one might expect that optimal reactions to one-player deviations are negligible, so that the two sets of equilibria are approximately the same. We investigate the circumstances in which this is true. The intuition here is very similar to that underlying the literature on

* We would like to thank Robert Anderson, Jean Tirole, Xavier Vives, and two referees for helpful comments. Research support from the National Science Foundation under Grants SES85-09697 and SES85-09484, the Sloan Foundation, and the UCLA Academic Senate is gratefully acknowledged. Current address of D.F.: Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02139.

0022-0531/88 $3.00

the limit points of monopolistic competition (Novshek and Sonnenschein [9], Roberts [10], Mas–Colell [7]), and our work can be viewed as extending that literature to more general games, as we explain below.

We consider three formulations of the proposition that, with many players, open-loop and closed-loop equilibrium sets are approximately the same. In Section 3 we consider nonatomic games in which no player's payoff is affected by the actions of a single rival. In this idealized or limiting version of a game with many players, open-loop equilibria are always closed-loop equilibria, but closed-loop equilibria may fail to be open-loop equilibria. If, however, there is a unique Nash equilibrium in every subgame, then the open- and closed-loop equilibria coincide.

In Sections 4 to 6, we consider the limits of finite games approaching a nonatomic game. We begin with an example in Section 4. Following this, our goal is to show that equilibria in the nonatomic game are approximately the same as those in the approaching finite games. This breaks into two parts. In Section 5 we demonstrate upper hemi-continuity: the limit of equilibria is an equilibrium, and every sequence of equilibria has a limit. Therefore if the nonatomic game has unique open- and closed-loop equilibria, all open and closed equilibria approaching the limit must be near each other. Our second goal, lower hemi-continuity, is less fully realized; Section 6 reports some results for the differentiable case from Fudenberg and Levine [3]. Section 7 considers a second approach to the question of lower-hemi-continuity. Here we show that in a game that is "almost" nonatomic, an open-loop equilibrium is an approximate closed-loop equilibrium.

Through most of the paper, we use two-period models for convenience. In Section 8 we explain how our results extend directly to any finite number of periods. We find that as the strategic possibilities increase with the length of the game and the set of closed-loop equilibria expands, it becomes increasingly difficult to guarantee that closed-loop equilibria are near to open-loop equilibria. We conclude in Section 9 by relating our results to some earlier research in this area.

2. The Model

The set of potential players is $P \equiv [0, 1]$. The set of probability measures on P is M. For now there will be only two periods; later we will discuss the n-period case. The set of actions available to player p in each period $t = 1, 2$ is A, which we take to be a compact convex subset of a Euclidean space. An *outcome* of the game in period t is a measurable map $x_t \colon P \to A$; the space of all such maps is x_t, and $X \equiv X_1 \times X_2$. In place of $x_t(p)$ we shall write x_t^p.

We will need to view $D \equiv X \times M$ as a topological space. We view D as the space of measures over $A^2 \times P$ which are degenerate in the first component: the support of $d = (x, \mu) \in D$ is the closure of the graph of x. For technical reasons, it will frequently be convenient to work not with D but with the space of all measures over $A^2 \times P$, which we denote $\bar{D}$. We endow $\bar{D}$ (and thus D) with the weak* topology, which is defined as follows. Each continuous real valued function $g: A^2 \times P \to \mathbb{R}$ induces a pseudometric ρ_g on $\bar{D}$; if $\Delta, \Delta' \in \bar{D}$, the distance is given by

$$\rho_g(\Delta, \Delta') \equiv \left| \int g(x, p)\, \Delta(dx, dp) - \int g(x, p)\, \Delta'(dx, dp)) \right|.$$

The weak* topology is the weakest (that is, coarsest) topology which makes all of these pseudometrics continuous. Consequently, $d_n \to d$ iff $\rho_g(d_n, d) \to 0$ for all continuous functions g. The space D is not closed in this topology: a sequence of degenerate measures can converge to a non-degenerate one. Indeed, it can be verified that any measure can be approximated by a sequence of degenerate ones,[1] so that the closure of D is $\bar{D}$.

The payoff function $\pi: P \times A^2 \times D \to \mathbb{R}$ gives each player's payoff as a function of his own action and the joint distribution of types and actions. We also use the notation $\pi^p(x^p, x, \mu)$ in place of $\pi(p, x^p, x, \mu)$. We will assume that

(A1) π is uniformly continuous with respect to ρ_g for *some* continuous function $g: A^2 \times P \to \mathbb{R}$.

Observe that (A1) implies that π is continuous. Moreover, uniform continuity implies that π has a unique extension[2] to $P \times A^2 \times \bar{D}$. This will be important when we consider the possibility of equilibrium distributions, that is, $d \in \bar{D}$, $d \notin D$. The relevant payoff function for this type of equilibrium will be the unique extension of π.

We will focus on two polar cases of relative player size. We say that μ is *nonatomic* if $\mu(p) = 0$ for all p; μ is *finite* if it has finite support. Note that the payoff π^p depends on p's own action both directly and indirectly through the effect that it has on the joint distribution of players and

[1] This can be shown by approximating the given measure by one with finite support, which is approximated in turn by a continuous function.

[2] If $\bar{d} \in \bar{D}$, then for every $\varepsilon > 0$, the ε-ball $(d \mid \rho_g(d, \bar{d}) < \varepsilon)$ contains points in D. Extract a sequence d_n with $\rho_g(d^n, \bar{d}) \to 0$, and let $\pi^p(x^p, \bar{d})$ be any limit of $\pi^p(x^p, d^n)$. Any convergent net $d^\alpha \to \bar{d}$ must have the property that for all $\varepsilon > 0$ there is an α' such that $\alpha > \alpha'$ implies $\rho_g(d^\alpha, \bar{d}) < \varepsilon$, so we can use the uniform continuity of π with respect to ρ_g to conclude that the extension is uniquely defined. Continuity now follows directly. For details on uniform continuity see Kelley [6].

actions. For notational simplicity we will sometimes combine these two effects and write

$$\pi^p[x, \mu] = \pi^p(x^p, x, \mu).$$

Continuity implies that if $x = \hat{x}$ almost everywhere w.r.t. μ, then $\pi^p(x^p, x, \mu) = \pi^p(x^p, \hat{x}, \mu)$. With nonatomic players this means that no single player has any influence on another's payoff. Including p's own action as a separate argument allows an individual player to have a substantial impact on his own payoff. Notice that the nonatomic assumption rules out games in which a player has little effect on most rivals, but has a large effect on just a few of them. Thus, this type of game may be a good model of a large impersonal market, but it cannot be a good model of a large number of interrelated small markets. (With finite players, the indirect effect of p's own action on his payoff though D need not be trivial.)

The first equilibrium concept we shall consider is that of an *open-loop* equilibrium. This is the Nash equilibrium of the game in which players precommit themselves to a sequence of actions at the beginning of the game. It will be convenient to have a concise notation for the outcome obtained from $\hat{x}$ by replacing its pth component with x^p. We will denote this outcome $\hat{x}\backslash x^p$. Let supp μ denote the support of μ. Formally

DEFINITION 2.1. An *open-loop* ε-equilibrium relative to μ is an $\hat{x} \in X$ such that for all $p \in \operatorname{supp} \mu$ and $x^p \in X^p$,

$$\pi^p[\hat{x}, \mu] \geqslant \pi^p[\hat{x}\backslash x^p, \mu] - \varepsilon.$$

This says that given the actions of all other players, no player can gain more than ε by changing his action. If $\varepsilon = 0$ we shall speak simply of an open-loop equilibrium, and indeed we shall not consider the case $\varepsilon > 0$ until the end of the paper. We ignore players $p \notin \operatorname{supp} \mu$ as these players are effectively not present in the game described by μ.

Now we wish to generalize the idea of an ε-equilibrium to the closure $\bar{D}$ of D. (Recall that π^p has a unique extension to $\bar{D}$.) Any $\bar{d} \in \bar{D}$ gives rise to a measurable correspondence y: $P \rightrightarrows A^2$ which assigns to each p in the support of μ the set of actions y such that (p, y) is contained in the support of $\bar{d}$. It also gives rise to a marginal measure $\bar{\mu}$ over players. Note that the following definition is only applicable to nonatomic games.

DEFINITION 2.2. An open-loop ε-equilibrium distribution relative to (nonatomic) μ is a $\bar{d} \in \bar{D}$ such that

(i) $\bar{\mu} = \mu$ (where $\bar{\mu}$ is the marginal over P derived from $\bar{d}$), and

(ii) for all $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$, $\pi^p(\hat{x}^p, \bar{d}) \geqslant \pi^p(x^p, \bar{d}) - \varepsilon$.

Again, if $\varepsilon = 0$, we refer simply to an open-loop distribution.

For $\bar{d} \in D$ this definition differs slightly from Definition 2.1. If $x \in X$ fails to be continuous as a function of p, then its graph is not closed. Since the support of a distribution is by definition the smallest closed set of full measure, the graph of the correspondence y is the *closure* of the graph of x. Thus Definition 2.2 appears to be stronger than Definition 2.1 in that it requires that (ii) hold for all $\hat{x}^p$ in the closure of the graph of x. However, since π is continuous, it is apparent that (ii) must continue to hold on the boundary. Consequently, the difference in the definitions is purely technical. Alternatively, it would suffice to assume in (ii) that there exists a dense subset of the support of $\bar{d}$, and define y to be the correspondence induced by this subset. By the continuity of π this is equivalent to both Definitions 2.2 and 2.1.

As noted, we are concerned with equilibrium distributions only for nonatomic games. These distributions should not be thought of as arising from randomizations by the players, as we have not introduced mixed strategies. Rather, these distributions are idealizations of situations in which nearby players play quite different actions. For example, imagine that the distribution on players is uniform, and $y_t^p = \{(0), (1)\}$ for all p and t. This is the limit of a distribution in which players between 0 and $1/n$ play zero, those between $1/n$ and $2/n$ play 1, $2/n$ to $3/n$ play zero, and so on.

The following proposition shows that open-loop distributions of nonatomic games can always be approximately "purified." That is, they correspond to the limits of open-loop ε-equilibria.

PROPOSITION 2.1. *If μ is nonatomic, then for every open-loop equilibrium distribution $\bar{d}$ there is a sequence x^n of open-loop ε^n-equilibria such that $(x^n, \mu) \to \bar{d}$ and $\varepsilon^n \to 0$.*

Proof. For each n, divide the support of μ into n nonintersecting intervals $(P_1, ..., P_n)$ of measure $1/n$ and divide $A^2 = A \times A$ into n^2 nonintersecting squares of equal size A_j. Let $s_{jk} = \bar{d}(A_j, P_k)$. For each j renumber the A_k so that s_{jk} is positive for A_1 through A_{k_j}. Subdivide P_j into k_j subintervals P_{jk}, such that $\bar{d}(\cdot, P_{jk}) = s_{jk}$. For each k, fix an $a_k \in A_k$. Now define x^n on supp μ by $x^n(p) = a_k$ for $p \in P_{jk}$. Thus x^n is a degenerate approximation of y. Clearly $(x^n, \mu) \to \bar{d}$; that x^n is an ε^n-equilibrium with $\varepsilon^n \to 0$ follows from the continuity of π. Q.E.D.

We will relate the open-loop equilibria to the perfect equilibria of the game in which all players observe first-period play before choosing their second-period actions. To define these "closed-loop equilibria" we introduce the concept of a reaction function $r: X_1 \to X_2$. The reaction function specifies how second-period play will respond to the first-period outcome.

In the game μ, we will require that r be insensitive to actions by players outside of supp μ, as such players are not actually present. We *do*, however, allow r to depend on the play of a single player, even if μ is nonatomic. This is because we do not wish to confound the players' *payoffs*, which are insensitive to measure-zero deviations, with their *information*, which includes the complete specification of first-period play. One might imagine that with nonatomic players, optimal reactions should ignore measure-zero deviations. The next section explores the conditions under which this is true. We do not wish to place a priori restrictions on the reaction functions for nonatomic μ because with small but finite players optimal reactions can be large.

In a closed-loop equilibrium, we require that x_1 be rational given r and that r represents a rational response to first-period deviations.

DEFINITION 2.3. A reaction function r is a *second period* ε-*equilibrium* relative to $\hat{x}_1$ and μ if for all p and q in supp μ and all $x_1^q, x_2^q \in A$,

$$\pi^p[\hat{x}_1 \backslash x_1^q, r(\hat{x}_1 \backslash x_1^q), \mu] \geqslant \pi^p[\hat{x}_1 \backslash x_1^q, r(\hat{x}_1 \backslash x_1^q) \backslash x_2^p, \mu] - \varepsilon.$$

In the above definition and subsequently the expression "$r(\hat{x}_1 \backslash x_1^q) \backslash x_2^p$" means the second-period outcome obtained by replacing the pth component of $r(\hat{x}_1 \backslash x_1^q)$ by x_2^p.

DEFINITION 2.4. A closed-loop ε-*equilibrium path* relative to μ is an $\hat{x} \in X$ such that there exists a reaction function r satisfying

(i) r is a second-period ε-equilibrium relative to $\hat{x}_1$ and μ, and

(ii) for all $p \in$ supp μ,

$$\pi^p[\hat{x}_1, r(\hat{x}_1), \mu] \geqslant \pi^p[\hat{x}_1 \backslash x_1^p, r(\hat{x}_i \backslash x_1^p), \mu] - \varepsilon.$$

A *closed-loop equilibrium* is a closed-loop equilibrium path together with a reaction function satisfying (i) and (ii).

Notice that this differs a bit from the standard formulation: usually r must be an optimal response to *all* initial x_1, not merely those which represent a deviation by one player alone. However, these additional responses are irrelevant to the equilibrium determination of $\hat{x}_1$. If an equilibrium reaction function r exists for *all* x_1 then our definition and the usual one yield the same sets of equilibrium *paths*. If an equilibrium reaction function r fails to exist for some x_1 then there can be *no* perfect equilibrium in the usual sense, but there may be in our sense.

3. The Nonatomic Case

We are primarily interested in the relationship between open- and closed-loop equilibria.

THEOREM 3.1. *If μ is nonatomic, then every open-loop equilibrium $\hat{x}$ is a closed-loop equilibrium path.*

Proof. We need only define r for initial profiles of the form $\hat{x}_1 \backslash x_1^p$. We set $r(\hat{x}_1 \backslash x_1^p) = \hat{x}_2 \backslash x_2^p$, where x_2^p is an optimal response by p to his own deviation. By nonatomicity this clearly means that $\hat{x}$ is a closed-loop equilibrium path. Q.E.D.

If we had required that second-period equilibrium hold for *all* x_1, and not merely for one-player deviations from $\hat{x}_1$, the theorem would be: every open-loop equilibrium is a closed-loop equilibrium path *if and only if* an optimal response function r exists.

It is *not* true that every closed-loop equilibrium path is an open-loop equilibrium: even though a deviation by player p has no effect on the second-period decision problem of rivals, it need not be true that the rivals do not react to his choice. Consider, for example, the pair of matrix games

$$\begin{bmatrix} (2,2) & (-4,4) \\ (4,-4) & (0,0) \end{bmatrix} \text{ in period 1}$$

$$\begin{bmatrix} (4,4) & (-10,-10) \\ (-10,-10) & (0,0) \end{bmatrix} \text{ in period 2.}$$

Each player on $P \equiv [0, 1]$ selects a pure strategy and receives the μ weighted average payoff against all opponents summed over both periods. Suppose μ is uniform. Obviously players are nonatomic. Since the games are unconnected and the first-period game is the prisoner's dilemma, in any open-loop equilibrium the first-period equilibrium is always $(0, 0)$. This is also a closed-loop equilibrium path. However, there is also a closed-loop equilibrium path in which the prisoner's dilemma is resolved and $(2, 2)$ is played in the first period. In the second period $(0, 0)$ and $(4, 4)$ are both equilibria if played by everyone. The reaction function is that all players play $(4, 4)$ in the second period if everyone played $(2, 2)$ in the first period, otherwise everyone plays $(0, 0)$ in the second period. Thus the reward to defecting in period 1 is 2 but the penalty in the second period is 4, so no defections will occur. Notice that although the game is nonatomic the reaction function is atomic: if any *one* player defects then *everyone* else responds.

As the example shows, for open- and closed-loop equilibrium paths to coincide a deviation by only one player must not affect the reactions of others.

DEFINITION 3.1. A reaction function r is nonatomic if $r^q(\hat{x}_1 \backslash x_1^p) = r^q(\hat{x}_1)$ for almost all $q \neq p$. Otherwise, we call the reaction function *atomic.*

THEOREM 3.2. *Suppose that for given nonatomic* μ *and each* $\hat{x}_1$ *there is a* unique *second-period equilibrium, that is,* $\hat{x}_2$ *such that*

$$\pi^p[\hat{x}_1, \hat{x}_2, \mu] \geqslant \pi^p[\hat{x}_1, \hat{x}_2 \backslash x_2^p, \mu] \qquad \text{for almost all}$$

$$p \in \operatorname{supp} \mu \text{ and all } x_2^p.$$

Then the unique reaction function is nonatomic.

Proof. Follows directly from the nonatomicity of π.

THEOREM 3.3. *With nonatomic* μ, *every closed-loop equilibrium path with nonatomic reaction function is an open-loop equilibrium.*

Note that an open-loop equilibrium is always a closed-loop equilibrium path supported by a nonatomic reaction function.

4. AN EXAMPLE

This section presents a two-period game in which each player's payoff depends on his own actions and the average actions of his opponents. For a large but finite number of players, the game has a unique open-loop equilibrium and two closed-loop ones. As the number of players goes to infinity, one of the closed-loop equilibria has the same limiting path as the open-loop one, while the other converges to an atomic closed-loop equilibrium. In this latter equilibrium the second-period outcome is responsive to deviations by single players, even though in the limit game no single player can influence the payoff of any other. From Theorem 3.2 we know that the limit game must have multiple second-period equilibria for the equilibrium first-period outcome. Yet for any finite number of players there is a unique second-period equilibrium. This illustrates two related points. First, atomic reactions in games with a continuum of players are not pathological: they can arise as the limits of finite-player equilibria. Second, the uniqueness of second-period equilibrium need not be preserved in passing to the nonatomic limit.

In our example, all players have the same payoff function $\pi^p(x^p, x, \mu)$, which depends on x only through its average $\bar{x}(\mu) \equiv \int x(p)\, d\mu$. We will only consider the distributions μ^n which place weight $1/n$ on n (identical) players $(0, 1/n, ..., (n-1)/n)$, and the limit μ^* which is uniform on $[0, 1]$. Thus $\bar{x}(\mu^n) = \Sigma x^p/n$, and $\bar{x}(\mu^*) = \int_0^1 x(p)\, dp$.

In the example, all actions are considered to lie in the interval [0, 10]. The payoffs are

$$\begin{aligned}\pi^p(x^p, x, \mu) &\equiv \pi^p(x^p, \bar{x}(\mu)) \\ &\equiv (x_1^p - 1)^2/2 - (x_2^p)^2/2 - (x_2^p - 2)^2\,[I(x_2^p \geqslant 2)] \\ &\quad + \lambda(x_2^p \bar{x}_2 - \bar{x}_2^2 + x_2^p \bar{x}_1),\end{aligned}$$

where $I(x_2^p \geqslant 2)$ is an indicator function, taking on the value 1 if $x_2^p \geqslant 2$ and 0 otherwise, and $0 \leqslant \lambda \leqslant 1$.

Note that each player's payoffs are strictly concave in his or her own actions and that cross-player effects work through averages. This implies that all the open-loop and second-period distributional equilibria are symmetric. In this section we consider the case $\lambda = 1$.

We first study the limit game μ^*. In an open-loop equilibrium, $\bar{x}_2$ is independent of first-period play, so $x_1^p = 1$ is a dominant strategy, and the unique open-loop equilibrium is (1, 5/2). This is a closed-loop equilibrium path as well. To look for other closed-loop equilibria we solve for the second-period equilibria as a function of $\bar{x}_1$. Straightforward computation shows that for $\bar{x}_1 > 0$, we must have $x_2^p = \bar{x}_2 = (4 + \bar{x}_1)/2$, while for $\bar{x}_1 = 0$, any $\bar{x}_2$ between 0 and 2 corresponds to a second-period equilibrium. We know that the only closed-loop equilibrium with a nonatomic reaction function is the open-loop equilibrium, so other closed-loop equilibria must have $\bar{x}_1 = 0$. It is easy to see that $\bar{x}_1 = 0$ can occur: take the reaction function $\bar{x}_2 = 1$ if all players set $x_1^p = 0$, and $\bar{x}_2 = 2$ if any player unilaterally deviates. Other second-period outcomes can also arise in equilibrium, as can a range of first-period outcomes whose average is 0. Note in particular that there can be nonsymmetric and even distributional outcomes in the first period.

Clearly atomic closed-loop equilibria cannot be approximated by open-loop ones, but this would not matter if the atomic equilibria were somehow "pathological." They are not, as can be seen by considering the finite-player version of the game.

We begin by solving for the unique second-period equilibrium relative to μ^n as a function of $\bar{x}_1$. These equilibria, which are symmetric, are given by

$$\bar{x}_2^n = r^n(\bar{x}_1) = \begin{cases} n\bar{x}_1, & 0 \leqslant \bar{x}_1 < 2/n \\ 2, & \bar{x}_1 = 2/n \\ n\bar{x}_1/(1+2n) + 4n/(1+2n), & \bar{x}_1 > 2/n. \end{cases}$$

We also note that the unique open-loop equilibrium has $x_1^{OL} = (2n+5)/(1+2n)$, which converges to 1 as $n \to \infty$.

While for any n and $\bar{x}_1$ there is a unique equilibrium, there is a sense in which $r^n(0)$ converges to a multi-valued limit as $n \to \infty$. As shown in Fig. 1, as n grows, $r^n(\bar{x}_1)$ becomes steeper and steeper near $\bar{x}_1 = 0$, and "converges" to $r^*(\bar{x}_1)$, the reaction function for μ^*, which is vertical at 0.

Another way of putting this is that $\rho^n(\bar{x}_1) \equiv dr^n(\bar{x}_1)/dx_1^p = 1$ in the neighborhood of $\bar{x}_1 = 0$, no matter how large n is. Even though in the n-player game, each player has very little effect on his opponents, the equilibrium remains very sensitive to any one player's actions. For $\bar{x}_1$ strictly greater than zero, we eventually have $\bar{x}_1 > 2/n$, so that $\rho(\bar{x}_1) = 4/(1+2n)$ converges to zero. This is the limiting behavior that one expects: second-period play becomes insensitive to any individual player's first-period action. One would expect that closed-loop equilibria in this region would be similar to open-loop ones, and that there is a sequence of closed-loop equilibria whose path $(\bar{x}_1, r_n(\bar{x}_1))$ converges to the same limit $(1, 5/2)$ as the open-loop ones.

It is easy to check that $\bar{x}_1 = 1/(n-1)$, $r_n(\bar{x}_1)$ is another closed-loop equilibrium, whose path converges to $(0, 1)$. It is also true, although harder to show, that these two are the only symmetric closed-loop equilibria of the game.

In the example, the open- and closed-loop correspondences are both well-behaved in passing to the continuum-of-players limit. Nevertheless, the two concepts differ, even in the limit, but for natural reasons. In order that all the closed-loop equilibria can be approximated by open-loop ones, we would need stronger conditions which guaranteed that the continuum game has a unique second-period equilibrium after every outcome *and* that in the finite-players games the relative size of the cross-effects between players is small compared to their own effects. We provided conditions of this sort in an earlier version of this paper [2].

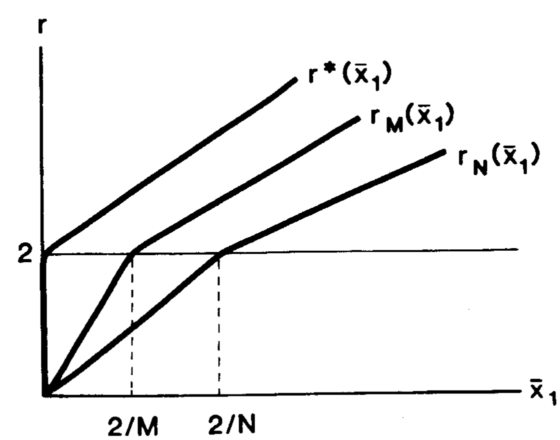

FIGURE 1

5. Upper Hemi-continuity

To relate open- and closed-loop equilibria in games with finite but small players, we will relate each of them to the corresponding equilibria with nonatomic players. If the open- and closed-loop equilibria coincide in the limit game, and both equilibrium correspondences are continuous with respect to the measure over player types, then we can conclude that the two sets of equilibria are close to each other when players are small but not infinitesimal.

The set of open-loop equilibrium distributions lies in $\bar{D}$ and is thus compact. In this section we introduce the notion of closed-loop equilibrium distributions, so that the closed-loop paths are compact as well. Thus the question of upper hemi-continuity reduces to one of whether the graphs of the equilibrium correspondences are closed. For open-loop equilibria this is almost immediate. For closed-loop equilibria, we use the compactness of $\bar{D}$ to argue that any limit of equilibrium paths can indeed be supported as a closed-loop equilibrium by *some* optimal reactions.

The conclusion that we draw is that every limit of equilibria is an equilibrium and every sequence of equilibria has a limit. If, with nonatomic players, there is a unique closed-loop and unique open-loop equilibrium, not only are they the same, but, as the players become "small," all open- and closed-loop equilibria must be close to one another.

THEOREM 5.1. *Suppose that* (A1) *holds, that* $\hat{x}^n$ *is an open-loop equilibrium for* μ^n, *and that* $d^n = (\hat{x}^n, \mu^n) \to \hat{d} \in \bar{D}$. *Then* $\hat{d}$ *is an open-loop equilibrium distribution of* μ.

Proof. As the proof is straightforward we give only a sketch. Since $d^n \to \hat{d}$, for each pair (p, y) in supp $\hat{d}$, there is a sequence $(p^n, \hat{x}^n(p))$ with p^n in supp μ^n such that p^n converges to p and $\hat{x}^n(p)$ converges to y. We then argue that a profitable deviation for p against y in game μ would, by continuity, be a profitable deviation for p^n against $\hat{x}_n$ in μ^n for n sufficiently large. Q.E.D.

COROLLARY 5.1. *If* $\mu^n \to \mu$ *and* $\hat{x}^n$ *is an open-loop equilibrium for* μ_n *then there is a convergent subsequence such that* $(\hat{x}^n, \mu^n) \to \hat{d} \in \bar{D}$ *and* $\hat{d}$ *is an open-loop equilibrium distribution for* μ.

We must now introduce the notion of a closed-loop distribution $\hat{d} \in \bar{D}$. As before, these distributions idealize situations in which nearby players play different actions. It is important that distributions have this interpretation, rather than being interpreted as mixed strategies. If players use mixed strategies in period one, then a closed-loop equilibrium path will be a first-period distribution and map from first-period realizations to second-period distributions, rather than a distribution jointly over both periods and players.

The notion of a second-period equilibrium distribution is similar to that of an open-loop distribution. New conceptual issues arise, however, in formulating the notion of optimal responses to a deviation from the first-period marginal distribution $\bar{d}_1$. As before, a distribution $\bar{d}_1$ gives rise to a correspondence y_1 which specifies the various actions "played by player p." The interpretation is that the distribution is an idealization of a situation in which nearby players play different actions. Thus, instead of specifying that there is a fixed reaction function $r(\hat{x}_1 \backslash x_1^p)$, we specify a family of reactions to player p's first-period play, with each reaction corresponding to a different action that player p is "supposed" to be playing, that is, for each $(p, \hat{x}_1^p) \in \operatorname{supp} \bar{d}_1$, the reaction function is $R(p, \hat{x}_1^p, x_1^p)$. In equilibrium, we will require that each player p is willing to play $\hat{x}_1^p$ when reactions are given by $R(p, \hat{x}_1^p, \cdot)$. As with open-loop distributions, closed-loop distributions can only be interpreted as idealized equilibria in nonatomic games. For notational reasons, we shall suppose that $R(p, \hat{x}_1^p, x_1^p) \in \bar{D}$, that is, a joint distribution over both periods, rather than a distribution only in the second period. We shall require that the marginal of R over $A \times P$, $R_1(p, \hat{x}_1^p, x_1^p)$, shall actually equal $\bar{d}_1$, since the first-period marginal distribution is not affected by a single player deviating.

Definition 5.1. For a nonatomic distribution μ, $R(p, \hat{x}_1^q, x_1^q)$ is a second-period ε-equilibrium distribution relative to a marginal first-period distribution $\bar{d}_1$ (over $A \times P$) if

(i) $\bar{\mu} = \mu$, (where $\bar{\mu}$ is the marginal over P derived from $\bar{d}$),

(ii) $R_1(p, \hat{x}_1^p, x_1^p) = \bar{d}_1$, and

(iii) for $(p, \hat{x}^p) \in \operatorname{supp} R(q, \hat{x}_1^q, x_1^q)$ and any $x_2^p \in A$,

$$\pi^p(\hat{x}^p, R(q, \hat{x}_1^q, x_1^q)) \geqslant \pi^p((\hat{x}_1^p, x_2^p), R(q, \hat{x}_1^q, x_1^q)) - \varepsilon.$$

Since the definition is intended to make sense only in the nonatomic case (hence the use of $\pi^p(\ )$ rather than $\pi^p[\ ]$), the requirement is simply that a player p playing $\hat{x}^p$ cannot gain by deviation in the second period.

Definition 5.2. A distribution $\bar{d} \in \bar{D}$ is a closed-loop ε-equilibrium distribution relative to a nonatomic μ if

(i) $\bar{\mu} = \mu$ (where $\bar{\mu}$ is the marginal over P derived from $\bar{d}$),

and there exists an $R(p, \hat{x}_1^p, x_1^p)$ such that

(ii) if $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$, then $R(p, \hat{x}_1^p, \hat{x}_1^p) = \bar{d}$,

(iii) for all x_1^p, $R(p, \hat{x}_1^p, x_1^p)$ is a second-period ε-equilibrium, and

(iv) for all $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$ and all $x^p \in A$,

$$\pi^p(\hat{x}^p, R(p, \hat{x}_1^p, \hat{x}_1^p)) \geq \pi^p(x^p, R(p, \hat{x}_1^p, x_1^p)) - \varepsilon.$$

Condition (ii) requires that first-period play consistent with $\bar{d}$ does not provoke a reaction in period 2. As in the definition of an open-loop distributional equilibrium, the notion of a closed-loop equilibrium in Definition 5.2 differs in a technical way from the earlier Definition 2.4 when $\bar{d}$ is a unit mass on a point in D. Specifically, Definition 5.2 requires that (iii) hold for any limit of $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$, while Definition 2.4 does not. Although R may not itself be continuous, the fact that π is continuous is enough to establish the equivalence of the two definitions. All that is required is that *some* $R(p, \hat{x}_1^p, \cdot)$ exist satisfying (iii). If $(p^n, \hat{x}(p^n)) \to (p, \hat{x}^p)$ define $R(p, \hat{x}_1^p, x_1^p)$ to be any limit of a convergent subsequence of $R(p^n, \hat{x}_1^{(p^n)}, x_1^p)$. Evidently, by continuity of π, (iii) will be satisfied.

THEOREM 5.2. *Under* (A1) *if* $(\hat{x}_1^n, \hat{x}_2^n)$ *is a sequence of closed-loop equilibrium paths relative to* μ^n, *and* $(\hat{x}_1^n, \hat{x}_2^n, \mu_n) \to \bar{d}$, *then* $\bar{d}$ *is a closed-loop equilibrium distribution relative to* $\bar{\mu}$. (*Notice that in the statement and proof of this theorem superscript n's refer to the position in the sequence of equilibria, while superscript p's refer to players.*)

Proof. We must construct a reaction function R such that (ii) and (iii) are satisfied. For each $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$ let $p^n \to p$ be such that $\hat{x}^n(p^n) \to \hat{x}^p$ (such a sequence must exist since $(\hat{x}^n, \mu) \to \bar{d}$). If $x_1^p \neq \hat{x}_1^p$, define $R(p, \hat{x}_1^p, x_1^p)$ to be any accumulation point of $(\hat{x}_1^n, r^n(\hat{x}_1^n \backslash x_1^{(p^n)}))$ and define $R(p, \hat{x}_1^p, \hat{x}_1^p) = \bar{d}$. Theorem 5.1 shows that the R thus constructed always prescribes equilibrium play. Now we must argue that no player can gain by deviating in the first period. If there was a gain, there would be $(p, \hat{x}^p) \in \operatorname{supp} \bar{d}$ and x^p such that $\pi^p(\hat{x}^p, \bar{d}) < \pi^p(x^p, R(p, \hat{x}_1^p, x_1^p)) - \varepsilon$. But then consider the sequence $p^n \to p$, $\hat{x}^n(p^n) \to \hat{x}^p$. Since π is continuous in p, and $(\hat{x}_1^n, r^n(\hat{x}_1^n \backslash x_1^{(p^n)})) \to R(p, \hat{x}_1^p, x_1^p)$, for n large enough p^n would gain by deviating to x^p in the game μ^n. Q.E.D.

To illustrate the import of Theorem 5.2, we examine the example of Section 4, in the case $\lambda < \frac{1}{6}$. In this case the unique second-period non-atomic distributional equilibrium is

$$\bar{x}_2 = r(\bar{x}_1) = \lambda \bar{x}_1 / (1 - \lambda).$$

It follows directly that there is a unique closed-loop equilibrium in which $\bar{x}_1 = \bar{x}_2 = 0$, and which is also open-loop. From Theorem 5.2, it follows that for n large all open- and closed-loop equilibria have $\bar{x}_1$ and $\bar{x}_2$ close to zero, and consequently the two are approximately the same.

6. The Differential Approach to Lower Hemi-continuity

We now wish to consider whether equilibria are lower hemi-continuous; that is, whether a given nonatomic open- or closed-loop equilibrium has a close approximation in a similar game with many small atomic players. If this is true for both open- and closed-loop equilibria, then we can conclude that every open-loop equilibrium in the large finite game has a closed-loop equilibrium nearby.

There are two approaches to lower hemi-continuity. One, explored in the next section, is to weaken the definition of equilibrium by considering ε-equilibria with $\varepsilon > 0$. Alternatively, we can make various differentiabiltity assumptions and use a version of the implicit function theorem. This latter program is complicated and is carried out in another paper (Fudenberg and Levine [3]). Here we summarize how the results of that paper apply to the example of the previous section with $\lambda = 1$.

The relevant equilibrium to which the theorem applies is the limiting open-loop equilibrium at (1, 5/2). This satisfies the first-order conditions

$$\begin{aligned} \frac{\partial \pi^p}{\partial x_1^p} &= x_1^p - 1 + \left[\frac{1}{n} x_2^p\right] = 0 \\ \frac{\partial \pi^p}{\partial x_2^p} &= -3x_2^p + 4 + (\bar{x}_2 + \bar{x}_1) + \left[\frac{1}{n}(x_2^p - 2\bar{x}_2)\right] = 0, \end{aligned} \tag{6.1}$$

which we can write as $\phi(\bar{x}) + (1/n)\,\psi(\bar{x}) = 0$. Since

$$D\phi = \begin{bmatrix} 1 & 0 \\ 1 & -2 \end{bmatrix}$$

is nonsingular, and the perturbing term $\psi(\bar{x})$ is uniformly C^1 bounded, for n sufficiently large the implicit function theorem guarantees a solution to (6.1) near $\phi(\bar{x}) = 0$. That is, for n large there are open-loop equilibria near to the open-loop equilibrium (1, 5/2) of the nonatomic game.

The case of closed-loop equilibrium is similar, but more complicated. The first-period condition is

$$\phi_1(\bar{x}) + \left[\frac{1}{n}(x_2^p - 2\bar{x}_2)\right] + \left\{\frac{1}{1+2n}(x_2^p - 2\bar{x}_2)\right\} = 0, \tag{6.2}$$

where the final term arises from muliplying $\partial \pi^p / \partial \bar{x}_2^p$ times $\partial r_n / \partial x_1$. Again as $n \to \infty$, (6.2) approaches $\phi_1(\bar{x}) = 0$ and the implicit function theorem applies. In Fudenberg and Levine [3] we generalize this line of argument to a larger class of nonsymmetric games.

Finally, consider the atomic closed-loop equilibrium of the limit game.

Here the argument fails because $\partial r^n/\partial x_1$ does not vanish as $n \to \infty$. Moreover, lower hemi-continuity fails as well. In the finite games, all the equilibria with $\bar{x}_1$ approaching 0 are symmetric; in the limit, there are also many nonsymmetric equilibria.

7. The ε-Equilibrium Approach to Lower Hemi-continuity

Hereafter we shall suppose that μ is finite. We shall accordingly drop the nonatomicity assumption.

Obviously it is unreasonable to expect that open-loop and closed-loop equilibria are exactly the same with finite players: in general it *will be* optimal to respond in period two to deviations in period one. Consider instead the assertion that "open-loop equilibria are almost closed-loop equilibria in large games." There are two things we might mean by this. We might mean that every open-loop equilibrium has a closed-loop equilibrium nearby. This is not generally true, as we have seen. Alternatively we might weaken the closed-loop concept and assert that open-loop play is "almost" optimal. It might be plausible to suppose that players will not deviate provided the gain from doing so is very small. Thus we might be interested in not only equilibria, but all ε-equilibria with ε a very small number. The assertion we shall now state and prove is that every open-loop equilibrium is an ε-closed-loop equilibrium where the smaller is the interaction between players, the smaller is ε.

With this in mind we define

Definition 7.1. The *atomicity* of a game μ is

$$\text{atom}(\mu) = \sup_{p \neq q, \hat{x}, x^q} |\pi^p[\hat{x}, \mu] - \pi^p[\hat{x} \backslash x^q, \mu]|.$$

This is a measure of the largest effect any player can have on another. In a nonatomic game continuity implies $\text{atom}(\mu) = 0$.

Theorem 7.1. *Every open-loop equilibrium is a closed-loop ε-equilibrium where* $\varepsilon = 2\,\text{atom}(\mu)$.

Proof. Let $\hat{x}$ be an open-loop equilibrium. Define r as in the proof of Theorem 3.1 so that p reacts to himself by responding optimally, but no one else responds at all. We need only show that following any $\hat{x}_1 \backslash x_1^q$, $\hat{x}_2 \backslash x_2^q$ is ε-optimal for $p \neq q$ with $\varepsilon = 2\,\text{atom}(\mu)$. If not, there is player p and action x_2^p such that

$$\pi^p[\hat{x}_1 \backslash x_1^q, \hat{x}_2 \backslash x_2^p \backslash x_2^q, \mu] - \pi^p[\hat{x} \backslash x^q, \mu] > 2\,\text{atom}(\mu).$$

We know $\pi^p[\hat{x}, \mu] - \pi^p[\hat{x}_1, \hat{x}_2\backslash x_2^p, \mu] \leqslant 0$, because $\hat{x}_2$ is a second-period equilibrium given $\hat{x}_1$. We also know that $|\pi^p[\hat{x}, \mu] - \pi^p[\hat{x}\backslash x^q, \mu]| <$ atom(μ) and that $|\pi^p[\hat{x}_1, \hat{x}_2\backslash x_2^p, \mu] - \pi^p[\hat{x}_1\backslash x_1^q, \hat{x}_2\backslash x_2^p\backslash x_2^q, \mu]| <$ atom(μ). Combining these inequalities, we see from the inequality triangle that 2 atom$(\mu) <$ 2 atom(μ), a contradiction. Q.E.D.

Notice, incidentally, that atom(μ) small only says that the effect of a single player deviating is small; the effect of all rivals simultaneously deviating is potentially as large as $(n-1)$ atom(μ), where n is the number of players. In large games the *aggregate* effect of rivals can be quite large. Indeed if we have a sequence of n-player games μ^n we can clearly have $\lim_{N \to \infty}$ atom$(\mu^n) \to 0$ and $\lim_{N \to \infty} (n-1)$ atom$(\mu^n) \to \infty$ so that individual effects vanish, but aggregate effects become large.

Now we consider whether a closed-loop equilibrium is an ε-open-loop equilibrium. We can say little about this without very strong assumptions. First, there is no reason subsequent play should be relatively insensitive to deviations by just one player. Although such deviations change each rival's decision problem only a little, a small perturbation of a game may shift the *equilibrium* quite a lot, even if each player's payoff is "very concave" in his own actions. Second, even if subsequent play only moves a little, this response is by *all* players and each player moving a little can have a large aggregate effect on the first-period decision problem. Moreover, in contrast to the infinite case, assuming that the second-period equilibrium is unique does not solve any of these problems.

8. The Many-Period Case

While we have presented only a two-period model, the results of this section extend to games with any finite number of periods. A closed-loop equilibrium with many periods is (as before) the same as a subgame-perfect equilibrium except that we only require the reaction functions to be defined for unilateral deviations from the equilibrium path. An open-loop equilibrium is still a closed-loop equilibrium if players are nonatomic, and every closed-loop equilibrium with nonatomic reaction functions is an open-loop equilibrium. The analog to Theorem 3.2, which provides sufficient conditions for reaction functions to be nonatomic, requires the assumption that there is a unique Nash equilibrium in every subgame.

How do the results on ε-equilibrium change in games with more than two periods? The formal answer is that they do not. The atomicity of a game is the maximum effect one player has on any other (moving in all periods at the same time), and every open-loop equilibrium is a closed-loop equilibrium with $\varepsilon = 2$ atom(μ). Thus increasing the number of periods in a "stationary" game does not change the required epsilon if the payoffs are normalized so that the "total payoff" does not increase with the period length. One example of this is a sequence of T-period games with an L-period lag structure, that is,

$$\pi^p[x, \mu] = \frac{1}{T-L} \sum_{t=L+1}^{T} u^p[x_{t-1}, ..., x_{t-L}, \mu],$$

as is the case, for example, in Maskin and Tirole [8]. Of course in non-stationary games epsilon cannot be expected to be independent of the number of periods. It is also essential that payoffs be normalized by $1/(T-L)$, so that the overall size of payoffs does not increase with the length of the game.

The interpretation of Theorem 7.1 does change with the number of periods, because the set of closed-loop ε-equilibria typically grows with the horizon. Indeed, Fudenberg and Levine [1] show that as the horizon grows, the set of closed-loop ε-equilibria of this game approaches the set of infinite-horizon closed-loop equilibria. Thus, knowing that the open-loop equilibrium is a closed-loop equilibrium is "less informative" in games with more periods, because they have more closed-loop equilibria. Why does this matter? A story one might tell to justify open-loop equilibria using the ε-equilibrium approach is: First, one accepts ε-equilibrium as a *descriptive* model of behavior, arguing that players do not bother to attain small gains. Then one argues that in the set of (possibly quite complicated) closed-loop ε-equilibria, the open-loop equilibrium is a natural focal point because of its simplicity. This focal point argument may become less compelling as the set of equilibria increases. In the prisoner's dilemma, "efficient" outcomes become ε-equilibria once the horizon is long enough and these may have an equal claim to being "focal."

9. Related Work

The literature most closely related to our paper is that on the limit points of monopolistic competition. Monopolistic competition can be viewed as a two-stage game in which firms choose quantities in the first period, and an exchange equilibrium occurs in the second. With this interpretation, monopolistic competition is a closed-loop equilibrium and price-taking is open-loop. Consequently, Roberts' [10] results on the limit of monopolistically competitive equilibria can be viewed as a special case of ours.

Green [4] also studies the relationship between price-taking and monopolistically competitive equilibria with many firms, but differs in treating "anonymous" equilibria of infinitely repeated games. Finally, papers by Green [5] and Walker [11] provide sufficient conditions for the (open-loop) Nash equilibrium correspondence to be upper hemicontinuous in the number of players.

References

1. D. Fudenberg and D. Levine, Subgame perfect equilibrium of finite- and infinite-horizon games, *J. Econ. Theory* **31** (1983), 251–268.
2. D. Fudenberg and D. Levine, Open- and closed-loop equilibria in dynamic games with many players (January 1985).
3. D. Fudenberg and D. Levine, Finite player approximations to a continuum of players (March 1986).
4. E. Green, Noncooperative price taking in large dynamic markets, *J. Econ. Theory* **22** (1980), 155–182.
5. E. Green, Continuum and finite-player non-cooperative models of competition, *Econometrica* **52** (1984), 975–994.
6. J. Kelley, "General Topology," Springer–Verlag, New York/Berlin, 1955.
7. A. Mas-Colell, The Cournotion foundations of Walrasian equilibrium theory: An exposition of recent theory, *in* "Advances in Economic Theory" (W. Hildenbrand, Ed.), Chap. 6, Cambridge Univ. Press, New York, 1982.
8. E. Maskin and J. Tirole, "A Theory of Dynamic Oligopoly. I. Overview and Quantity Competition with Large Fixed Costs," MIT Working Paper, MIT, Cambridge, MA, 1982.
9. W. Novshek and H. Sonnenschein, Small efficient scale as a foundation for Walrasian equilibrium, *J. Econ. Theory* **22** (1980), 243–256.
10. K. Roberts, The limit points of monopolistic competition, *J. Econ. Theory* **22** (1980), 256–278.
11. M. Walker, A generalization of the maximum theorem, *Int. Econ. Rev.* (February 1979).

Printed by Catherine Press, Ltd., Tempelhof 41, B-8000 Brugge, Belgium