DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

NASH EQUILIBRIUM SEARCH FOR THE BEST ALTERNATIVE

Jennifer F. Reinganum

SOCIAL SCIENCE WORKING PAPER 413

January 1982

ABSTRACT

In a recent paper, Weitzman [1979] described a policy of ''optimal search for the best alternative.'' The present paper is concerned with the development and characterization of a policy of ''Nash equilibrium search for the best alternative.'' Specifically, it is shown that, under certain monotonicity assumptions, and under the assumption that firms have incomplete information regarding the results of rivals' search behavior, a Nash equilibrium search policy exists and has the same form as Weitzman's optimal search policy.

NASH EQUILIBRIUM SEARCH FOR THE BEST ALTERNATIVE

Jennifer F. Reinganum[1]

I. INTRODUCTION

In a recent paper, Weitzman [1979] described a policy of ''optimal search for the best alternative.''[2] The present paper is concerned with the development and characterization of a policy of ''Nash equilibrium search for the best alternative.'' Specifically, it is shown that, under certain monotonicity assumptions, and under the assumption that firms have incomplete information regarding the results of rivals' search behavior, a Nash equilibrium search policy exists and has the same form as Weitzman's optimal search policy.

This paper generalizes the results of an earlier paper [Reinganum 1980], and illustrates another application of strategic search theory. The model is applied to the problem of strategically developing drilling or mining sites for the purpose of subsequent resource extraction.

II. THE MODEL

The model is developed in the context of two distinct phases — a site evaluation phase followed by an extraction phase. Suppose that each of M firms holds title to a geographical region believed to contain deposits of a natural resource. This region may be divided into a number of potential drill or excavation sites. Each firm's deposit may be irregularly distributed beneath the ground, so that the amount of resource which is recoverable from any given site may vary from site to site. It is assumed to be prohibitively expensive to sink more than one shaft to extract the resource reserves of the region. However, less dramatic procedures such as core samples or surface examination of sites can be done at more moderate expense. This process of ascertaining, for each site, the amount of the resource deposit which is recoverable if access is from that site, is termed site evaluation. Site evaluation is assumed to occur essentially instantaneously before extraction begins, and determines the initial stocks of recoverable reserves for the extraction phase. Each firm must evaluate and develop its own resource deposit in ignorance of the exploratory behavior of its rivals. However, each firm can foresee the manner in which the resource stocks will be depleted once extraction has begun. Specifically, we assume that extraction occurs in a Nash equilibrium fashion over a known time horizon (possibly infinite), given the initial stocks determined in the previous phase. Let $(y_1,\ldots,y_M)$ denote the M firms' eventual recoverable resource stocks.

Assumption 1. The value of y_m to firm m under Nash equilibrium extraction is denoted $u_m(y_1,\ldots,y_M)$, where u_m is nonnegative, jointly continuous in all its arguments and twice differentiable with $\partial u_m/\partial y_m > 0$, and $\partial^2 u_m/\partial y_m \partial y_k < 0$, for $k,m = 1,2,\ldots,M$.

That is, we assume that an increase in m's own resource stock increases the present value of m's profits; moreover, an increase in m's own stock or in any rival's stock decreases the marginal value to m of an increase in y_m.[3] This valuation is assumed to include the costs of sinking the access shaft via which the resource will be extracted.

Each firm m holds a geographical region containing a known amount of resource, denoted $\bar{x}_m$. This region is assumed to be divided into finitely many potential drill or excavation sites. The number of sites for firm m is denoted n_m and m's sites are indexed by $i \varepsilon I_m = \{1,2,\ldots,n_m\}$. An evaluation of firm m's ith site consists of a surface examination or core sample, which results in the observation of x_{mi}, which represents the amount of the resource deposit which is recoverable if access to the deposit is gained from site i. Since the firm doesn't know precisely which sites are associated with the highest levels of recoverable reserves, the quantities x_{mi} are random variables. Let $F_{mi}(x)$ denote the probability that $x_{mi} \leq x$, for $x \varepsilon [\underline{x}_m, \bar{x}_m]$. Let $F_{mi}(x)$ be continuous and strictly increasing in x with a continuous density $f_{mi}(x)$. Note that x_{mi} and $x_{mi'}$ are independent random variables for all i, i' ε I_m, and for all m. A cost of c_{mi} is required to ascertain the value of x_{mi}.

Each firm must decide the order in which to explore its sites and when to cease exploration and settle upon an extraction site. Moreover, we assume that firm m cannot observe the exploratory

behavior of its rivals or the results of its rivals' site evaluations (i.e., the realizations of rival firms' levels of recoverable reserves from each of their sites). Thus the firm's information at any stage consists of its own remaining sites, its most promising site observed to date and the level of reserves recoverable from that site. Partition I_m into two sets. Define $S_m = \{i \varepsilon I_m \mid i \text{ has been evaluated}\}$, and $\bar{S}_m = I_m - S_m$, the unexamined sites. Since all x_{mi} are independent, the values of x_{mi} are unimportant; it is only y_m, the maximum of those x_{mi}'s observed to date, which matters. Thus the firm's state variables are $(\bar{S}_m, y_m)$. Firm m must choose a selection rule describing which site to explore next, and a stopping rule which specifies whether to continue or cease exploration if the current state is $(\bar{S}_m, y_m)$. Let P^m denote the set $\{\bar{S}_m \mid \bar{S}_m \subseteq I_m\}$.

Definition 1. A selection rule for firm m is a mapping $\rho_m : P^m \times [\underline{x}_m, \bar{x}_m] \to [0,1]^{n_m}$ where $\rho_{mi}(\bar{S}_m, y_m)$ is the probability that firm m will examine site i ε I_m next if $(\bar{S}_m, y_m)$ is the current state. Thus $0 \leq \rho_{mi}(\cdot,\cdot) \leq 1$ with $\sum_{i \varepsilon \bar{S}_m} \rho_{mi}(\bar{S}_m, y_m) = 1$ for all $(\bar{S}_m, y_m) \varepsilon P^m \times [\underline{x}_m, \bar{x}_m]$.

Definition 2. A stopping rule for firm m is a mapping $\pi_m : P^m \times [\underline{x}_m, \bar{x}_m] \to [0,1]$, where $\pi_m(\bar{S}_m, y_m)$ denotes the probability that firm m will examine another site when the current state is $(\bar{S}_m, y_m)$. Again, $0 \leq \pi_m(\bar{S}_m, y_m) \leq 1$ for all $(\bar{S}_m, y_m) \varepsilon P^m \times [\underline{x}_m, \bar{x}_m]$.

Definition 3. A <u>strategy</u> for firm m consists of a pair (ρ_m, π_m) of a selection rule and a stopping rule for m.

For compactness of notation, let $\sigma_m = (\rho_m, \pi_m)$ and $\tilde{\sigma}_m = (\sigma_1, \ldots, \sigma_{m-1}, \sigma_{m+1}, \ldots, \sigma_M)$. Thus $(\sigma_m, \tilde{\sigma}_m)$ is a complete list of strategies for the M firms, while $\tilde{\sigma}_m$ lists only those of m's rivals.

III. NASH EQUILIBRIUM

Since firm m cannot observe the outcomes of its rivals' site examinations and evaluations, and is assumed to take the rival strategy vector $\tilde{\sigma}_m$ as given, we can define $U_m(y_m, \tilde{\sigma}_m) = E_{\tilde{\sigma}_m}[u_m(y_m, \tilde{y}_m)]$, where the expectation is taken with respect to the distribution of $\tilde{y}_m$ which is induced by the strategies $\tilde{\sigma}_m$. Note that $U_m(y_m, \tilde{\sigma}_m)$ inherits the property that $\partial U_m / \partial y_m > 0$ from the assumption that $\partial u_m / \partial y_m > 0$.

Define $\Omega_m(\bar{S}_m, y_m, \tilde{\sigma}_m)$ to be the expected value to m of continuing optimally from $(\bar{S}_m, y_m)$ on, <u>when m's rivals play the strategies</u> $\tilde{\sigma}_m$. For each $(\bar{S}_m, y_m) \varepsilon P^m \times [\underline{x}_m, \bar{x}_m]$, Ω_m must satisfy

$$\Omega_m(\bar{S}_m, y_m, \tilde{\sigma}_m) = \max\{U_m(y_m, \tilde{\sigma}_m), \max_{i \varepsilon \bar{S}_m} \{-c_{mi} + \Omega_m(\bar{S}_m - \{i\}, y_m, \tilde{\sigma}_m) F_{mi}(y_m)$$

$$+ \int_{y_m}^{\bar{x}_m} \Omega_m(\bar{S}_m - \{i\}, x_{mi}, \tilde{\sigma}_m) f_{mi}(x_{mi}) dx_{mi}\}\} \qquad (1)$$

where $\Omega_m(\emptyset, y_m, \tilde{\sigma}_m) \equiv U_m(y_m, \tilde{\sigma}_m)$.

Definition 4. A strategy σ_m for which equation (1) holds is a <u>best response</u> for m to the rivals' strategies $\tilde{\sigma}_m$.

Note that m's selection rule is determined by the maximization inside the outer brackets, while m's stopping rule is determined by the maximization outside the outer brackets.

Assumption 2. Suppose that for each $\tilde{\sigma}_m$ there exists some $i \varepsilon I_m$ such that

$$\int_{\underline{x}_m}^{\bar{x}_m} U_m(x_{mi}, \tilde{\sigma}_m) f_{mi}(x_{mi}) dx_{mi} - c_{mi} > 0$$

$m = 1, 2, \ldots, M$.[4]

Assumption 2 states that before any site evaluation is conducted, the expected marginal contribution to firm m's profits due to examining a site exceeds the cost of examining that site, regardless of the resource stocks held by the other firms, for at least one of m's sites; moreover, this is assumed to hold for all firms m. This assumption rules out the possibility that no site examination and evaluation is a best response to some strategy $\tilde{\sigma}_m$; at least one site will be examined by each firm, and each firm will exploit its resource deposit.

Definition 5. Define firm m's <u>contingent reservation stock for site i</u>, $\phi_{mi}(\tilde{\sigma}_m) = \min\{\phi \varepsilon [\underline{x}_m, \bar{x}_m] \mid H_{mi}(\phi, \tilde{\sigma}_m) \leq 0\}$, where

$$H_{mi}(\phi,\tilde{\sigma}_m) = \int_{\phi}^{\bar{x}_m} [U_m(x_{mi},\tilde{\sigma}_m) - U_m(\phi,\tilde{\sigma}_m)] f_{mi}(x_{mi}) dx_{mi} - c_{mi}. \quad (2)$$

To see that $\phi_{mi}(\tilde{\sigma}_m)$ exists and is unique, note that $H_{mi}(\bar{x}_m,\tilde{\sigma}_m) < 0$ for all $\tilde{\sigma}_m$), and $\partial H_{mi}(\phi,\tilde{\sigma}_m)/\partial\phi < 0$ by the fact that U_m is increasing in y_m. Thus either $H_{mi}(\underline{x}_m,\tilde{\sigma}_m) \leq 0$, in which case $\phi_{mi}(\tilde{\sigma}_m) = \underline{x}_m$, or $H_{mi}(\underline{x}_m,\tilde{\sigma}_m) > 0$, in which case $\phi_{mi}(\tilde{\sigma}_m)$ is uniquely defined by $H_{mi}(\phi,\tilde{\sigma}_m) = 0$.

Proposition 1. A best response for firm m to the rival strategy vector $\tilde{\sigma}_m$ exists and has the following form.

$$\rho_{mi}(\bar{S}_m,y_m) = \begin{cases} 1 & \text{if } i\varepsilon\bar{S}_m \text{ and } \phi_{mi}(\tilde{\sigma}_m) > \phi_{mj}(\tilde{\sigma}_m) \text{ for all } j\varepsilon\bar{S}_m \\ 1 - \sum_{i'\varepsilon\bar{S}_m'} \rho_{mi'}(\bar{S}_m,y_m) & \text{if } i \ \varepsilon\ \bar{S}_m' \\ 0 & \text{otherwise} \end{cases}$$

where

$$\bar{S}_m' = \{i',i''\varepsilon\bar{S}_m \mid \phi_{mi'}(\tilde{\sigma}_m) = \phi_{mi''}(\tilde{\sigma}_m) > \phi_{mj}(\tilde{\sigma}_m) \text{ for all } j\ \varepsilon\ \bar{S}_m,\ j \neq i',i''\},$$

and

$$\pi_m(\bar{S}_m,y_m) = \begin{cases} 1 & \text{if } y_m < \phi_{mi}(\tilde{\sigma}_m) \text{ for some } i\ \varepsilon\ \bar{S}_m \\ [0,1] & \text{if } y_m = \phi_{mi}(\tilde{\sigma}_m) \text{ for some } i\ \varepsilon\ \bar{S}_m. \\ 0 & \text{otherwise} \end{cases}$$

Proof: Application of Weitzman's [1979] results to the problem or finding a best response for m to $\tilde{\sigma}_m$ implies that the strategy above is a best response for m.

Q.E.D.

We can characterize the best response rule verbally as follows.

Selection Rule: Given a strategy vector for m's rivals, $\tilde{\sigma}_m$, firm m should examine next that site which has the highest contingent reservation stock; if there is a tie, firm m may select any randomization over the tying sites.

Stopping Rule: Firm m should cease exploration as soon as current reserves exceed the contingent reservation stock of the next site selected (and consequently, of every remaining site).

Assumption 3. Suppose that for each m and all $i,i' \ \varepsilon\ I_m$, the following statements are true: if $i > i'$, then

(a) $F_{mi}(x) \geq F_{mi'}(x)$ for all x; and

(b) $c_{mi} \geq c_{mi'}$.

That is, sites which are more costly to examine and evaluate are also more likely to yield lower levels of recoverable resources. This is consistent with the interpretation that more costly sites are less accessible; lower accessibility would also tend to make extraction of the deposit through that site more difficult. This monotonicity assumption allows us to further simplify the selection rule.

Since $H_{mi}(\phi,\tilde{\sigma}_m)$ is non-decreasing in i under Assumption 3, firm m's contingent reservation stocks can be ordered as follows: if $i > i'$, then $\phi_{mi'}(\tilde{\sigma}_m) \geq \phi_{mi}(\tilde{\sigma}_m)$; moreover, this is true for all $\tilde{\sigma}_m$. Thus firm m simply examines its sites in the order $1,2,\ldots,n_m$, regardless of the rivals' strategies.

Thus we have shown that under assumptions 1-3, if a Nash equilibrium exists, it is in strategies of the following form:

firm m selects n_m (non-contingent) reservation stocks

$\xi_m = (\xi_{m1},\ldots,\xi_{mn_m}) \in D_m = \{\xi_m \in [\underline{x}_m,\bar{x}_m]^{n_m} \mid \bar{x}_m \geq \xi_{m1} \geq \ldots \geq \xi_{mn_m} \geq \underline{x}_m\}$,

and stops exploring after N_m sites, where

$N_m = \min\{i \in I_m \mid x_{mi} \geq \xi_{mi}\}$.

Thus firm m's strategy can be fully characterized by $\xi_m \in D_m$.

Definition 6. Redefine a strategy for firm m to be a reservation cost vector $\xi_m \in D_m$.

Denote $(\xi_1,\ldots,\xi_{m-1},\xi_{m+1},\ldots,\xi_M)$ by $\tilde{\xi}_m$ as before.

Definition 7. Redefine a best response for firm m to $\tilde{\xi}_m$ to be $\phi_m(\tilde{\xi}_m)$, where $\phi_{mi}(\tilde{\xi}_m) = \min\{\phi \in [\underline{x}_m,\bar{x}_m] \mid H_{mi}(\phi,\tilde{\xi}_m) \leq 0\}$, where

$$H_{mi}(\phi,\tilde{\xi}_m) = \int_{\phi}^{\bar{x}_m} [U_m(x_{mi},\tilde{\xi}_m) - U_m(\phi,\tilde{\xi}_m)] f_{mi}(x_{mi})\,dx_{mi} - c_{mi},$$

and $U_m(y_m,\tilde{\xi}_m) = E_{\tilde{\xi}_m}[u_m(y_m,\tilde{y}_m)]$, where the expectation is taken with respect to the distribution of $\tilde{y}_m$ which is induced by the reservation stocks $\tilde{\xi}_m$. (These distributions are calculated explicitly in the Appendix).

The following result is proven in the Appendix.

Proposition 2. The functions $U_m(y_m,\tilde{\xi}_m)$, $H_{mi}(\phi,\tilde{\xi}_m)$ and $\phi_{mi}(\tilde{\xi}_m)$ are jointly continuous in all their arguments.

Proposition 3. If $\phi_{mi}(\tilde{\xi}_m) > \underline{x}_m$, then $\phi_{mi}(\tilde{\xi}_m)$ is a decreasing function of $\tilde{\xi}_m$.

Proof: Note that if $\tilde{\xi}'_m \geq \tilde{\xi}_m$ with strict inequality in at least one component, then $\tilde{y}_m$ is stochastically greater under $\tilde{\xi}'_m$ than under $\tilde{\xi}_m$. Thus $U_m(y_m,\tilde{\xi}'_m) = E_{\tilde{\xi}'_m}[u_m(y_m,\tilde{y}_m)] < E_{\tilde{\xi}_m}[u_m(y_m,\tilde{y}_m)] = U_m(y_m,\tilde{\xi}'_m)$. Moreover, $\partial U_m(y_m,\tilde{\xi}'_m)/\partial y_m < \partial U_m(y_m,\tilde{\xi}_m)/\partial y_m$ because $\partial^2 u_m/\partial y_m \partial y_k < 0$ and y_k is stochastically greater for some k under $\tilde{\xi}'_m$ than under $\tilde{\xi}_m$. This in turn implies that H_{mi} is decreasing in $\tilde{\xi}_m$. Recall that H_{mi} is also decreasing in ϕ. To maintain the equality $H_{mi}(\phi_{mi}(\tilde{\xi}_m),\tilde{\xi}_m) = 0$ in the face of an increase in $\tilde{\xi}_m$, it must be that $\phi_{mi}(\tilde{\xi}_m)$ decreases. If $\phi_{mi}(\tilde{\xi}_m) = \underline{x}_m$, then it is unaffected by a marginal change in $\tilde{\xi}_m$.

Q.E.D.

Definition 8. An M-vector of strategies $\xi^* = (\xi_1^*,\ldots,\xi_M^*)$ is a Nash equilibrium if $\xi_m^* = \phi_m(\tilde{\xi}_m^*)$, $m = 1,2,\ldots,M$.

Proposition 4. There exists a Nash equilibrium $\xi^* \in D_1 \times D_2 \times \cdots \times D_M$.

Proof: Notice that $\phi_m(\tilde{\xi}_m) \in D_m$ due to Assumption 3. Define the vector mapping $\Phi: D_1 \times D_2 \times \cdots \times D_M \rightarrow D_1 \times D_2 \times \cdots \times D_M$ by $\Phi(\xi) = \{(\phi_{mi}(\tilde{\xi}_m))_{i \in I_m}\}_{m=1}^{M}$. Φ is continuous in ξ while $D_1 \times D_2 \times \cdots \times D_M$ is compact, convex and nonempty. Therefore Brouwer's fixed-point theorem assures us that Φ has a fixed point ξ^* $\in D_1 \times D_2 \times \cdots \times D_M$. Such a point has the property that $\xi_m^* = \phi_m(\tilde{\xi}_m^*)$, $m = 1,2,\ldots,M$, and thus constitutes a Nash equilibrium.

Q.E.D.

IV. SOME LIMITATIONS

Weitzman [1979] discusses a list of perturbations of the basic model to which his optimal search rule is not robust. In addition to this list, the equilibrium model developed above is sensitive to several other aspects of the problem which are not encountered in the single-agent model. For instance, if we relax Assumption 3, the monotonicity assumption, this causes no difficulty in the single-agent case. However, in the game described above, the following difficulty arises. Let $\tilde{\xi}_m \in \tilde{D}_m$ denote m's rivals' strategies. Then $\phi_m(\tilde{\xi}_m)$ is well-defined but not continuous in $\tilde{\xi}_m$. To see this, notice that a small change in $\tilde{\xi}_m$ may result in a different sampling order for some firm $k \neq m$, since the order of sampling is dependent on the ranking of $(\xi_{ki})_{i \in I_k}$. At the point at which k's ordering switches, although firm k is indifferent between the orderings (since both have the same reservations stocks at the switchpoint), firm m is not indifferent to k's sampling order (unless F_{ki} is the same for all $i \in I_k$ for all firms k). Thus unless the firms' sites differ only in their examination costs and not in their distribution functions, the function $U_m(\phi,\tilde{\xi}_m)$ is sensitive to the rivals' sampling orders and thus experiences a discontinuity at any point $\tilde{\xi}_m$ at which a rival changes the order in which it evaluates its sites. It may be possible to prove the existence of an equilibrium of this sort if we expand the strategy space for m to be the set of distribution functions over $[\underline{x}_m,\bar{x}_m]^{n_m}$, so that a random set of reservation stocks is drawn; this in turn implies an ordering of the sites. One may be able to prove that a Nash equilibrium in these distribution functions exists, but the existence of such an equilibrium would be considerably more difficult to establish.

Finally, the result that the Nash equilibrium rule is of the same form as Weitzman's optimal search rule is extremely sensitive to the assumption that the firms have incomplete information. That is, it is essential that the firms act in ignorance of their opponent's realizations and with full knowledge of their own. In the context of this application, this informational assumption seems quite reasonable. However, it is not reasonable in all contexts and this result should not be expected to hold up under alternative informational assumptions.

APPENDIX

The proof of Proposition 2 will rely upon the following results from analysis. Define a measure space (X, β, μ). For our purposes, X is a compact, convex, nonempty subset of R^{M-1} and μ is Lebesgue product measure.

Claim 1. If $f_n \to f$ in $L^1(\mu)$ and $\langle g_n \rangle$ is a sequence of measurable functions such that $|g_n| \leq M$, all n, and $g_n \to g$ almost everywhere (μ), then $g_n f_n \to gf$ in $L^1(\mu)$.

Claim 2. Suppose $f(t,x)$ is defined on a compact subset $TxX \subset R^M$. If f is a measurable function of x for each t and a continuous function of t for each x, then $h(t) = \int f(t,x)\,d\mu(x)$ is a continuous function of t.

The statements of these results for $X \subset R^1$ can be found in Royden [1968, p. 119 and 91, respectively]; the extensions to $X \subset R^{M-1}$ are straightforward.

Proof of Proposition 2. Define

$G_{ki}(a_k;\xi_k) = \Pr\{x_{kN_k} \leq a_k, N_k = i \mid \xi_k\}$. Direct computation and simplification of these probabilities yields

$$G_{k1}(a_k;\xi_k) = \begin{cases} 0 & a_k \leq \xi_{k1} \\ F_{k1}(a_k) - F_{k1}(\xi_{k1}) & a_k \geq \xi_{k1} \end{cases}$$

and, for $i \geq 2$,

$$G_{ki}(a_k;\xi_k) = \begin{cases} 0 & a_k \leq \xi_{ki} \\ \prod_{j=1}^{i} F_{kj}(a_k) - \prod_{j=1}^{i} F_{kj}(\xi_{ki}) & \xi_{ki} \leq a_k \leq \xi_{k(i-1)} \\ \prod_{j=1}^{i-1} F_{kj}(\xi_{k(i-1)}) F_{ki}(a_k) - \prod_{j=1}^{i} F_{kj}(\xi_{ki}) & a_k \geq \xi_{k(i-1)} \end{cases}$$

Let $g_{ki}(a_k;\xi_k)$ denote the density function of $G_{ki}(a_k;\xi_k)$.

$$g_{ki}(a_k;\xi_k) = \begin{cases} 0 & a_k < \xi_{ki} \\ \sum_{j=1}^{i} \prod_{\ell \neq j}^{i} f_{kj}(a_k) F_{k\ell}(a_k) & \xi_{ki} \leq a_k \leq \xi_{k(i-1)} \\ \prod_{j=1}^{i-1} F_{kj}(\xi_{ki}) f_{ki}(a_k) & a_k > \xi_{k(i-1)} \end{cases}$$

Let ξ_k^n be a sequence of vectors from D_k, which converges to $\xi_k \in D_k$. Define

$$g_{ki}^n = g_{ki}(a_k;\xi_k^n) = \sum_{j=1}^{i} \prod_{\ell \neq j}^{i} f_{kj}(a_k) F_{k\ell}(a_k) I_{[\xi_{ki}^n, \xi_{k(i-1)}^n]}(a_k)$$

$$+ \prod_{j=1}^{i-1} F_{kj}(\xi_{ki}^n) f_{ki}(a_k) I_{(\xi_{k(i-1)}^n, \bar{x}_k]}(a_k)$$

where $I(\cdot)$ is an indicator function. Note that g_{ki}^n is uniformly bounded since $|g_{ki}(a_k,\xi_{ki})| \leq M_k = \sum_{i \in I_k} \max_{a_k \in [\underline{x}_k, \bar{x}_k]} f_{ki}(a_k)$ for all $(a_k,\xi_k) \in [\underline{x}_k, \bar{x}_k] \times D_k$ because the densities $f_{ki}(\cdot)$ are continuous on compact domains and are thus bounded by their maximum values.

Moreover, g^n_{ki} converges pointwise to g_{ki} since

$$\lim_{n\to\infty} g^n_{ki} = \left[\lim_{n\to\infty}\sum_{j=1}^{i}\prod_{\ell\neq j}^{i} f_{kj}(a_k)F_{k\ell}(a_k)\right]\left[\lim_{n\to\infty} I_{[\xi^n_{ki},\xi^n_{k(i-1)}]}(a_k)\right]$$
$$+\left[\lim_{n\to\infty} F_{k1}(\xi^n_{ki})\cdots F_{k(i-1)}(\xi^n_{ki})f_{ki}(a_k)\right]\left[\lim_{n\to\infty} I_{(\xi^n_{k(i-1)},\bar{x}_k]}(a_k)\right] = g_{ki}(a_k;\xi_k)$$

Recall that $\tilde{a}_m = (a_1,\ldots,a_{m-1},a_{m+1},\ldots,a_M)$; $\tilde{\xi}_m$ is similarly defined. Define $G_k(a_k;\xi_k) = \sum_{i\in I_k} G_{ki}(a_k;\xi_k)$, and $\bar{G}_m(\tilde{a}_m;\tilde{\xi}_m) = \prod_{k\neq m} G_k(a_k;\xi_k)$. This is the joint distribution of $\tilde{y}_m$ by the independence of all the random variables. Denote the joint density by $\bar{g}_m(\tilde{a}_m;\tilde{\xi}_m)$. Then $\bar{g}^n_m = \bar{g}_m(\tilde{a}_m;\tilde{\xi}^n_m)$ is uniformly bounded by $\prod_{k\neq m} M_k$, and $\bar{g}^n_m$ converges pointwise to $\bar{g}_m$ as $n \to \infty$ (i.e., as $\tilde{\xi}^n_m \to \tilde{\xi}_m$).

Since $u_m(y_m,\tilde{a}_m)$ is continuous in y_m for each $\tilde{y}_m$, $\int u_m(y^n_m,\tilde{a}_m)d\mu(\tilde{a}_m) \to \int u_m(y_m,\tilde{a}_m)d\mu(\tilde{a}_m)$ as $n \to \infty$ [Claim 2]. Thus $u^n_m = u_m(y^n_m,\tilde{a}_m) \to u_m(y_m,\tilde{a}_m) = u_m$ in $L^1(\mu)$.

Now associate our sequences $\langle \bar{g}^n_m\rangle$ and $\langle u^n_m\rangle$ with Claim 1's $\langle g_n\rangle$ and $\langle f_n\rangle$. Claim 1 then implies that $\bar{g}^n_m u^n_m \to \bar{g}_m u_m$ in $L^1(\mu)$. That is,

$$\int u_m(y^n_m,\tilde{a}_m)\bar{g}_m(\tilde{a}_m;\tilde{\xi}^n_m)d\mu(\tilde{a}_m) \to \int u_m(y_m,\tilde{a}_m)\bar{g}_m(\tilde{a}_m;\tilde{\xi}_m)d\mu(\tilde{a}_m)$$

as $n \to \infty$. But $\int u_m(y_m,\tilde{a}_m)\bar{g}_m(\tilde{a}_m;\tilde{\xi}_m)d\mu(\tilde{a}_m)$ is what was previously defined as $U_m(y_m,\tilde{\xi}_m)$. Thus

$$U^n_m = U_m(y^n_m,\tilde{\xi}^n_m) \to U_m(y_m,\tilde{\xi}_m) = U_m \text{ as } n \to \infty.$$

Alternatively put, $U_m(y_m,\tilde{\xi}_m)$ is jointly continuous in $(y_m,\tilde{\xi}_m)$. The function H_{mi} inherits joint continuity from U_m and, finally, $\phi_{mi}(\tilde{\xi}_m)$ inherits continuity from H_{mi}.

Q.E.D.

FOOTNOTES

1. California Institute of Technology. I would like to thank Kim Border and Ed Green for their assistance. The financial support of the National Science Foundation is gratefully acknowledged.

2. Notation is consistent with Weitzman's wherever possible so as to make apparent the applicability of his results to the problem at hand. Research on problems related to those studied by Weitzman has been done by Bellman [1975], Kadane [1969], De Groot [1970], Stone [1975], Kadane and Simon [1977] and Spulber [1979].

3. Models of oligopolistic resource extraction can be found in Lewis and Schmalensee [1980a,b], Loury [1980], and Salant [1979]. No results pertaining to Assumption 1 are available from these studies. In another model, Stokey [1981] computes the Nash equilibrium in a model with a constant elasticity demand function and no extraction costs. In this case, the Nash equilibrium profit function for m, given the initial stocks $(y_1,\ldots,y_M)$, is easily shown to possess the properties discussed in Assumption 1.

4. A sufficient condition for Assumption 2 to hold is that there exist some $i \in I_m$ for which

$$\int_{\underline{x}_m}^{\bar{x}_m} u_m(y_1,\ldots,y_{m-1},x_{mi},y_{m+1},\ldots,y_M)f_{mi}(x_{mi})dx_{mi} - c_{mi} > 0$$

for all $\tilde{y}_m$. In an earlier model (Reinganum [1982]), which is simpler than this one in most respects, Nash equilibrium is shown to exist without assuming the equivalent of Assumption 2. Without Assumption 2, the distributions of rival firms' ultimate stock levels may be spikes at zero, rather than smooth distributions. This introduces a discontinuity in the best response functions.

REFERENCES

Bellman, Richard. Dynamic Programming. Princeton: Princeton University Press, 1957.

De Groot, Morris H. Optimal Statistical Decisions New York: McGraw-Hill, 1970.

Kadane, J. B. ''Quiz Show Problems.'' Journal of Mathematical Analysis and Applications 27 (1969), 609-623.

Kadane, J. B. and Herbert A. Simon. ''Optimal Strategies for a Class of Constrained Sequential Problems.'' The Annals of Statistics (March 1977), 237-255.

Kohn, Meir G. and Steven Shavell. ''The Theory of Search.'' Journal of Economic Theory 9 (1974), 93-123.

Lewis, Tracy R. and Richard Schmalensee. ''Cartel and Oligopoly Pricing of Nonreplenishable Natural Resources.'' In Dynamic Optimization and Mathematical Economics. New York: Plenum, 1980, 133-156. (a)

Lewis, Tracy R. and Richard Schmalensee. ''On Oligopolistic Markets for Nonrenewable Natural Resources.'' Quarterly Journal of Economics (November 1980), 475-491. (b)

Lippman, Steven A. and John J. McCall. ''The Economics of Job Search: A Survey.'' Economic Inquiry 14 (June 1976), 155-189.

Loury, Glenn C. ''A Theory of 'Oil'igopoly: Cournot Equilibrium In Exhaustible Resource Markets With Fixed Supplies.'' manuscript, 1980.

Reinganum, Jennifer F. ''Strategic Search Theory.'' International Economic Review February, 1982 (forthcoming).

Review of Economic Studies Symposium on the Economics of Information, 1977.

Royden, H. L. Real Analysis (2nd edition). New York: Macmillan, 1968.

Salant, Stephen W. ''Imperfect Competition In the International Energy Market: A Computerized Nash-Cournot Model.'' Operations Research (forthcoming).

Spulber, Daniel F. ''Optimal Search Across Labor Markets.'' Brown University Department of Economics Working Paper No. 79-5 (May 1979).

Stokey, Nancy L. ''Open-Loop and Closed-Loop Equilibria of An Oligopoly Model of Resource Extraction.'' manuscript (July 1981).

Stone, L. D. Theory of Optimal Search. New York: Academic Press, 1975.

Weitzman, Martin L. ''Optimal Search For the Best Alternative.'' Econometrica 47 (May 1979), 693-654.