Congestion Games with Player-Specific Payoff Functions*

Igal Milchtaich

Department of Mathematics,

The Hebrew University of Jerusalem,

Jerusalem 91904, Israel

March 1996

Abstract

A class of noncooperative games in which the players share a common set of strategies is described. The payoff a player receives for playing a particular strategy depends only on the total number of players playing the same strategy and decreases with that number in a manner which is specific to the particular player. It is shown that each game in this class possesses at least one Nash equilibrium in pure strategies. Best-reply paths in which players, one at a time, shift to best-reply strategies may be cyclic. But there is always at least one such path that connects an arbitrary initial point to an equilibrium.

1 Introduction

Rosenthal (1973) introduced a class of games in which each player chooses a particular combination of factors out of a common set of primary factors. The

*The present work is part of a Ph.D. dissertation written under the supervision of Professor B. Peleg, Professor U. Motro, and Professor S. Hart. I thank them for many helpful discussions. I thank a referee for raising the question of the properties of weighted congestion games with equal payoff functions and for the remark concerning best-reply cycles and multiplicity of best-reply strategies.

payoff associated with each primary factor is a function of the number of players who include it in their choice. The payoff a player receives is the sum of the payoffs associated with the primary factors included in his choice. Each game in this class possesses at least one pure-strategy Nash equilibrium. This result follows from the existence of a potential (Monderer and Shapley, 1991)—a real-valued function over the set of pure strategy-tuples having the property that the gain (or loss) of a player shifting to a new strategy is equal to the corresponding increment of the potential function.

The present report is concerned with games in which the payoff function associated with each primary factor is not universal but player-specific. This generalization is accompanied, on the other hand, by assuming these two limiting assumptions: that each player chooses only one primary factor and that the payoff received actually decreases (not necessarily strictly so) with the number of other players selecting the same primary factor. These congestion games, while not generally admitting a potential, nevertheless always possess a Nash equilibrium in pure strategies.

Such congestion games may have certain realizations in such fields as economics, traffic flows, and ecology. Milinsky (1979) simulated two different drift food patches in a stream by feeding six sticklebacks from two ends of a tank. On average, the fish distributed themselves between the two halves of the tank in the ratio of the food supply rates. Thus no individual could achieve a higher feeding rate by moving to the other patch. This Nash equilibrium is an example of an evolutionary stable strategy (ESS) (Maynard Smith, 1982, p. 63) or more specifically of an ideal free distribution (IFD) (Fretwell and Lucas, 1969). Where suitability of food patches or habitats decreases with density, and where individuals are free to enter any patch on an equal basis with residents, an IFD is said to occur when each individual settles in the patch most suited for its survival or reproduction (for a review of the effect of competition for resources on an individual's choice of patch see Milinsky and Parker, 1991). Does such an equilibrium always exist? The results of the present work suggest that a Nash equilibrium in pure strategies, that is, an equilibrium assignment of individuals to patches, should exist for any number of individuals and any number of patches, provided that, within each patch, individuals have equal feeding rates (Theorem 2). This result holds even if patches differ in the kind of food being offered in them and if individuals differ in their food preferences and in the additional value they attach to patches (taking, for example, predation risks into consideration). If, however, individuals also differ in their relative feeding rates, and if there are more than two patches, then a pure-strategy Nash equilibrium may not exist. Individual differences in competitive ability or in dominance were observed, under broadly similar experimental conditions, in sticklebacks (Milinsky, 1984), in cichlid fish (Godin and Keenleyside, 1984), in mallards (Harper, 1982), and in goldfish (Sutherland et al., 1988). This case of differential individual effect of players upon the payoff of others, which goes beyond the basic model portrayed above, is modeled in Section 8. Such "weighted" congestion games, in contrast with the "unweighted" congestion games considered above, do not always possess a pure-strategy Nash equilibrium.

The players of a game may reach an equilibrium by some sort of an adaptation process (see, for example, Kandori et al., 1993, and Young, 1993). Perhaps most simply, "myopic" players may react to the strategies played by the other players by deviating to best-reply strategies. Considering the underlying time axis to be a continuum, such deviations may be assumed to take place one at a time. Does such a process always converge? For the congestion games under consideration, the process always converges when there are only two common strategies to choose from (Theorem 1), or when players have equal payoff functions. But in the case of general "unweighted" congestion games counterexamples can be found. Assuming, however, a stochastic order of deviators, a convergence almost surely occurs (Theorem 3).

2 The model

In the present report, noncooperative games satisfying the following condition are referred to as (unweighted) congestion games: The n players share a common set of r strategies; the payoff the ith player receives for playing the jth strategy is a monotonically nonincreasing function S_{ij} of the total number n_j of players playing the jth strategy. Denoting the strategy played by the ith player by σ_i, the *strategy-tuple* $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ is a Nash equilibrium iff each σ_i is a best-reply strategy:

$$S_{i\sigma_i}(n_{\sigma_i}) \geq S_{ij}(n_j + 1) \qquad \text{for all } i \text{ and } j.$$

Here $n_j = \#\{1 \leq i \leq n \mid \sigma_i = j\}$. $(n_1, n_2, \ldots, n_r)$ is called the *congestion vector* corresponding to σ.

3 The symmetric case

A congestion game is symmetric if and only if all players share the same set of payoff functions, denoted by $S_1, S_2, \ldots, S_r$. Rosenthal (1973) defined for such symmetric games the exact potential function

$$P(\sigma) = \sum_{k=1}^{r} \sum_{m=1}^{n_k} S_k(m).$$

When only the ith player shifts to a new strategy, the jth one, the potential changes by

$$\Delta P = S_j(n_j + 1) - S_{\sigma_i}(n_{\sigma_i}),$$

which is equal to what the ith player gains (or loses). Any "local" maximum of P, a strategy-tuple where changing one coordinate cannot result in a greater value of P, hence corresponds to a pure-strategy Nash equilibrium.

The existence of an exact potential function further implies the *finite improvement property* (FIP) (Monderer and Shapley, 1991): Any sequence of strategy-tuples in which each strategy-tuple differs from the preceding one in only one coordinate (such a sequence is called a *path*), and the unique deviator in each step strictly increases the payoff he receives (an *improvement* path), is finite. The first strategy-tuple of a path is called the *initial point*; the last one is called the *terminal point*. Obviously, any *maximal* improvement path, an improvement path that cannot be extended, is terminated by an equilibrium.

For later reference we note that the fact that congestion games with equal payoff functions possess the FIP can also be proved without invoking a potential function. If there exists an infinite improvement path for such a game then there exists an improvement path $\sigma(0), \sigma(1), \ldots, \sigma(M)$ ($M > 1$), where $\sigma(0) = \sigma(M)$. Let $(n_1(k), n_2(k), \ldots, n_r(k))$ be the congestion vector corresponding to $\sigma(k)$ ($0 \leq k \leq M$). We can rearrange the indices in such a way that $S_1(n_1(1)) \leq \min_k S_j(n_j(k))$ holds for every strategy j for which $n_j(k)$ ($0 \leq k \leq M$) is not constant and in such a way that $n_1(1) > n_1(0)$. The latter assumption implies that for the unique deviator at the first step, player i, $\sigma_i(1) = 1$ but $\sigma_i(0) = j \neq 1$. This implies $S_1(n_1(1)) > S_j(n_j(0))$, which contradicts the above minimality assumption.

4 The two-strategy case

Nonsymmetric congestion games do not generally admit an exact potential function. Nevertheless, in the special case $r = 2$ we have

THEOREM 1. *Congestion games involving only two strategies possess the finite improvement property.*

Proof. Suppose the contrary, that there exists an infinite improvement path $\sigma(0), \sigma(1), \ldots$ for some two-strategy congestion game. It may be assumed that, for some $M > 1$, $\sigma(0) = \sigma(M)$. Let $(n_1(k), n_2(k))$ be the congestion vector corresponding to $\sigma(k)$ $(0 \leq k \leq M)$. Without loss of generality, it may be assumed that $n_2(1) = \max_k n_2(k)$, and therefore $n_1(1) = n - n_2(1) = \min_k n_1(k)$. This implies that the unique deviator in the first step, player i, deviate from 1 to 2; hence $S_{i2}(n_2(1)) > S_{i1}(n_1(1)+1)$. By the monotonicity of the payoff functions, $S_{i2}(n_2(k)) > S_{i1}(n_1(k)+1)$ holds for all $0 \leq k \leq M$. Hence player i does not deviate back to strategy 1 in steps $2, 3, \ldots, M$. This contradicts the assumption that $\sigma_i(M) = \sigma_i(0) = 1$. $\square$

5 Games without the finite improvement property

The finite improvement property is equivalent to the existence of a *generalized ordinal potential* for the game under consideration—a real-valued function over the set of pure strategy-tuples that strictly increases along any improvement path (Monderer and Shapley, 1991). Indeed, for a game that possesses the FIP the integer-valued function that assigns to a strategy-tuple σ the number of strategy-tuples which are the initial point of an improvement path with the terminal point σ is easily seen to be a generalized ordinal potential.

Here is an example of a two-player congestion game which does not possess the finite improvement property. Three strategies (a minimal number, by Theorem 1) are involved, numbered 1, 2, and 3. Assuming that $S_{11}(1) > S_{12}(n_2) > S_{13}(n_3) > S_{11}(2)$ and $S_{22}(1) > S_{21}(n_1) > S_{23}(n_3) > S_{22}(2)$ hold for all n_1, n_2, and n_3, the path $(3,2)$, $(2,2)$, $(2,3)$, $(1,3)$, $(1,1)$, $(3,1)$, and back to $(3,2)$ is an improvement path, involving six strategy-tuples. The existence of such a cycle demonstrates that the game under consideration does not admit even a generalized ordinal potential. Nevertheless, pure-strategy Nash equilibria do exist: these are the strategy-tuples $(1,2)$ and $(2,1)$.

Paths in which in each step the unique deviator shifts to a strategy which is a best reply against the strategies played by the other players are called *best-reply paths*. A best-reply strategy need not be unique. If players deviate only when the strategy they are currently playing is not a best-reply strategy then the path is a *best-reply improvement path*. Clearly, the finite improvement property implies the corresponding property for best-reply improvement paths, the *finite best-reply property* (FBRP), but the converse is not true.

The following argument shows that infinite best-reply improvement paths involve at least three players. At each step only one player, i, is changing his strategy, by deviating to the jth strategy, say. Therefore only one coordinate of the congestion vector, the jth one, is increased. Hence a second player is negatively affected—his payoff is reduced—only if he too plays the jth strategy. Occasionally, such reductions must take place: no player's payoff can increase indefinitely. If only two players are involved then it is this second player who makes the next move, changing from j to a strategy which is a best reply against the strategy (j) played by i. As this can only result in a smaller n_j (and greater n_k, for some $k \neq j$) the jth strategy remains a best reply for i, and thus an equilibrium is reached.

The path $(2,1,1)$, $(3,1,1)$, $(3,3,1)$, $(3,3,2)$, $(2,3,2)$, $(2,1,2)$, and back to $(2,1,1)$, where 1, 2, and 3 are three distinct strategies, is a best-reply improvement path in a three-player congestion game where the inequalities $S_{13}(1) > S_{12}(1)$, $S_{23}(2) > S_{21}(2)$, $S_{32}(1) > S_{31}(1)$, $S_{12}(2) > S_{13}(2)$, $S_{21}(1) > S_{23}(1)$, and $S_{31}(2) > S_{32}(2)$ hold, and S_{ij}'s not listed here are minimal. This path is shown graphically in Fig. 1. Thus, a three-player congestion game for which these conditions hold does not possess the finite best-reply property and therefore does not admit a generalized ordinal potential. It does, however, possess two pure-strategy Nash equilibria: $(3,1,2)$ and $(2,3,1)$. Note that this example involves a *generic* game, a game where different strategies yield different payoffs. The existence of best-reply cycles does not thus depend on multiplicity of best-reply strategies.

It can be shown (by extending the earlier argument concerning best-reply paths involving two players) that for a generic three-player congestion game the above inequalities are *necessary* conditions for an infinite best-reply path to occur: *every* infinite best-reply path consists of a finite path followed by an endless repetition of a six strategy-tuple cycle having the form indicated above.

6 The existence of a pure-strategy Nash equilibrium

The above examples illustrate a general property of the class of games defined in Section 2:

THEOREM 2. *Every (unweighted) congestion game possesses a Nash equilibrium in pure strategies.*

Before proving Theorem 2, we prove a lemma. Part (a) of the lemma is concerned with paths where each deviator moves to the next deviator's present position. Part (b) is concerned with paths where each deviator takes the last deviator's previous position.

LEMMA. (a) *If $j(0), j(1), \ldots, j(M)$ is a sequence of strategies, $\sigma(0), \sigma(1), \ldots, \sigma(M)$ is a best-reply improvement path, and $\sigma(k)$ results from the deviation of one player from $j(k-1)$ (the strategy which he played in $\sigma(k-1)$) to $j(k)$ ($k = 1, 2, \ldots, M$), then $M \leq n$ (n is the number of players).*

(b) *Similarly, if the deviation in the kth step is from $j(k)$ to $j(k-1)$ ($k = 1, 2, \ldots, M$) then $M \leq n \cdot (r-1)$ (r is the number of strategies).*

Proof. (a) Let $(n_1(k), n_2(k), \ldots, n_r(k))$ be the congestion vector corresponding to $\sigma(k)$ ($0 \leq k \leq M$), and set $(n_j)_{\min} = \min_k n_j(k)$ ($1 \leq j \leq r$). Clearly $(n_j)_{\min} \leq n_j(k) \leq (n_j)_{\min} + 1$ holds for all j and k. Equality on the right holds for $j = j(k)$; equality on the left holds for $j \neq j(k)$. Hence by deviating to $j(k)$ the unique deviator in the kth step brings $n_{j(k)}$ to its maximum and all other n_j's to their minimum. Therefore, by the monotonicity of the payoff functions, $j(k)$ remains a best reply for that player in all subsequent steps. Thus each player deviates at most once.

(b) Here too $(n_j)_{\min} \leq n_j(k) \leq (n_j)_{\min} + 1$, but equality on the left holds for $j = j(k)$. By deviating from $j(k)$, the unique deviator in the kth step thus brings $n_{j(k)}$ to its minimum. This implies that his payoff in $\sigma(k)$ is not only greater than in $\sigma(k-1)$ (which is the case by definition of a best-reply improvement path) but also greater than his payoff when he deviated to $j(k)$, if he did, or the payoff he will get by deviating to $j(k)$ at some later stage. Therefore a player will not return to a strategy he deviated from; each player thus deviates at most $r-1$ times. □

Proof of Theorem 2. The proof proceeds by induction on the number n of players. For $n = 1$ the proof is trivial. We assume that the theorem holds true for all $(n-1)$-player congestion games and prove it for n-player games. A given n-player congestion game Γ can be reduced into an $(n-1)$-player game $\overline{\Gamma}$ by "deleting" the last player. The reduced game is also a congestion game; the payoff functions $\overline{S}_{ij}$ in this game are defined by

$$\overline{S}_{ij}(\overline{n}_j) = S_{ij}(\overline{n}_j) \qquad \text{for } 1 \leq i \leq n-1 \text{ and all } j,$$

where $\overline{n}_j = \#\{1 \leq i \leq n-1 \mid \sigma_i = j\}$. By the induction hypothesis, there exists in $\overline{\Gamma}$ a pure-strategy Nash equilibrium $\overline{\sigma} = (\sigma_1(0), \sigma_2(0), \ldots, \sigma_{n-1}(0))$. Let $(\overline{n}_1, \overline{n}_2, \ldots, \overline{n}_r)$ be the congestion vector corresponding to $\overline{\sigma}$. Going back to Γ, let

$\sigma_n(0)$ be a best reply of player n against $\overline{\sigma}$. Note that $S_{i\sigma_i(0)}(\overline{n}_{\sigma_i(0)}) \geq S_{ij}(\overline{n}_j+1)$ holds for all i and j. Starting with $j(0)=\sigma_n(0)$ and with $\sigma(0)=(\sigma_1(0),\sigma_2(0),\ldots,\sigma_{n-1}(0),\sigma_n(0))$, we can find a sequence $j(0),j(1),\ldots,j(M)$ of strategies, and a best-reply improvement path $\sigma(0),\sigma(1),\ldots,\sigma(M)$ connected to it as in part (a) of the lemma, such that M (≥ 0) is maximal. We claim that $\sigma(M)=(\sigma_1(M),\sigma_2(M),\ldots,\sigma_n(M))$ is an equilibrium. For every player i who has deviated from the strategy he played in $\sigma(0)$, the strategy $\sigma_i(M)$ is a best reply against $\sigma(M)$—this is shown in the proof of the lemma. It remain to show that $\sigma_i(M)=\sigma_i(0)$ is a best-reply strategy for every player i who has not deviated. If $\sigma_i(M)=j(M)$ and $j(M)$ is not a best-reply strategy for player i then by deviating from $j(M)$ to a best-reply strategy $j(M+1)$ player i changes $\sigma(M)$ into a new strategy-tuple $\sigma(M+1)$, which may be appended to the above best-reply improvement path and thus contradict the assumed maximality of M. If, on the other hand, $\sigma_i(M)\neq j(M)$ then the number of players playing $\sigma_i(M)=\sigma_i(0)$ is the same in $\sigma(M)$ and in $\overline{\sigma}$; all other strategies are being played by at least as many players (for $n_j(M)\geq \overline{n}_j$ holds for all j, and equality holds for $j\neq j(M)$). As remarked above, $S_{i\sigma_i(0)}(\overline{n}_{\sigma_i(0)}) \geq S_{ij}(\overline{n}_j+1)$ holds for all j. In the case under consideration these inequalities imply $S_{i\sigma_i(M)}(n_{\sigma_i(M)}(M)) \geq S_{ij}(n_j(M)+1)$, for all j, and thus $\sigma_i(M)$ is a best reply for i against $\sigma(M)$. $\square$

7 Convergence to an equilibrium

The proof of Theorem 2 is a constructive one: an algorithm is given for finding an equilibrium in a given n-player congestion game—by adding one player after the other—in at most $\binom{n+1}{2}$ steps. The question arises, can an equilibrium be reached in the given game itself, when the constant presence of all n players is being taken into consideration? The next theorem gives an affirmative answer to this question.

THEOREM 3. *Given an arbitrary strategy-tuple $\sigma(0)$ in a congestion game Γ, there exists a best-reply improvement path $\sigma(0),\sigma(1),\ldots,\sigma(L)$ such that $\sigma(L)$ is an equilibrium and $L\leq r\binom{n+1}{2}$.*

Proof. Suppose first that $\sigma(0)=(\sigma_1(0),\sigma_2(0),\ldots,\sigma_n(0))$ is "almost" an equilibrium: $\sigma_i(0)$ is a best-reply strategy against $\sigma(0)$ for all $1\leq i\leq n-1$, but not necessarily for $i=n$. Starting with $j(0)=\sigma_n(0)$ and with $\sigma(0)$, we can find a sequence $j(0),j(1),\ldots,j(M)$ of strategies, and a best-reply improvement path $\sigma(0),\sigma(1),\ldots,\sigma(M)$ connected to it as in part (a) of the lemma, such that M is maximal. Clearly the first deviator in this path is the nth player. If $j(M) \neq \sigma_n(0)$ then, starting with $\tilde{j}(M) = \sigma_n(0)$ and with $\sigma(M)$, we can find a sequence $\tilde{j}(M), \tilde{j}(M+1), \ldots, \tilde{j}(N)$ ($N \geq M$) of strategies, and a best-reply improvement path $\sigma(M), \sigma(M+1), \ldots, \sigma(N)$ connected to it as in part (b) of the lemma, such that N is maximal. If $j(M) = \sigma_n(0)$ then we set $N = M$. We claim that $\sigma(N) = (\sigma_1(N), \sigma_2(N), \ldots, \sigma_n(N))$ is an equilibrium. Suppose the contrary, that $\sigma_i(N)$ is not a best reply against $\sigma(N)$ for some player i. Suppose that j is a best-reply strategy for that player. If $\sigma_i(N) = \sigma_i(N-1) = \cdots = \sigma_i(k)$ and k is minimal (i.e., $k = 0$ or $\sigma_i(k-1) \neq \sigma_i(k)$), then, by construction of the above best-reply improvement path, $\sigma_i(k)$ is a best reply for i against $\sigma(k)$. There can be two reasons why j, but not $\sigma_i(N) = \sigma_i(k)$, is a best reply against $\sigma(N)$. Either (i) $n_{\sigma_i(N)}(N) > n_{\sigma_i(N)}(k)$, (ii) $n_j(N) < n_j(k)$, or both. By construction, (ii) can hold only if $j = \tilde{j}(N)$. But this contradicts the maximality of N: by deviating from $\sigma_i(N)$ to $\tilde{j}(N)$, player i changes $\sigma(N)$ into a new strategy-tuple, $\sigma(N+1)$, which may be appended to the above best-reply improvement path. The other possibility, (i), can hold only if $\sigma_i(N) = j(M)$ and $k = 0$, i.e., if $\sigma_i(0) = \sigma_i(1) = \cdots = \sigma_i(N) = j(M)$. By the maximality of M, $\sigma_i(N) = \sigma_i(M) = j(M)$ must then be a best reply strategy for i against $\sigma(M)$ (the argument is the same as in the proof of Theorem 2). Returning again to the above two possibilities, the assumption that this strategy is not a best reply for i against $\sigma(N)$ implies that either $n_{\sigma_i(N)}(N) > n_{\sigma_i(N)}(M)$, which is impossible by construction, or else $n_j(N) < n_j(M)$, which is possible only if $j = \tilde{j}(N)$. But, again, the latter possibility contradicts the maximality of N. Thus $\sigma_i(N)$ must be a best-reply strategy for i against $\sigma(N)$.

The theorem is evidently true for one-player games. To complete the proof, by induction on the number n of players, we assume that the theorem holds true for all $(n-1)$-player congestion games and show that it must then hold for all n-player congestion games. Let Γ be an n-player congestion game and let $\sigma(0) = (\sigma_1(0), \sigma_2(0), \ldots, \sigma_n(0))$ be given. The game $\tilde{\Gamma}$ derived from Γ by restricting the strategy set of the nth player to the single strategy $\sigma_n(0)$ is effectively an $(n-1)$-player congestion game. It therefore follows from the induction hypothesis that there exists in $\tilde{\Gamma}$ a best-reply improvement path $\sigma(0), \sigma(1), \ldots, \sigma(L)$, where the terminal point $\sigma(L)$ is an equilibrium of $\tilde{\Gamma}$. Clearly in Γ too this path is a best-reply improvement path and $\sigma(L)$ is "almost" an equilibrium of Γ. As shown above, this path may be extended to reach an equilibrium. By the lemma, the extension need not be more than rn steps long. This estimate gives $r\binom{n+1}{2}$ as an upper bound to the length of the shortest best-reply improvement path that connects an arbitrary initial point to an equilibrium. $\square$

Games in which every strategy-tuple is connected to some pure-strategy Nash equilibrium by a best-reply path are called *weakly acyclic* (WA). (This definition is slightly more general than the definition given in Young (1993), where the equilibrium reached is also required to be strict. The definitions coincide for generic games.) Assuming that the number of strategies is finite, WA games have the property that if the order of deviators is decided more-or-less randomly, and if players do not deviate simultaneously, then a best-reply path almost surely reaches an equilibrium. More precisely, suppose that the process of forming a best-reply path, starting from a fixed initial point, is a stochastic process in which each player who is not currently playing a best-reply strategy has a positive probability—which may depend on history—of being the next deviator and that these probabilities are bounded away from zero by some positive constant ε. If there are several best-reply strategies for a player then each one is played with a probability of at least ε. If each strategy-tuple is connected to some equilibrium by a best-reply path of length L or less then the probability that at least one of the strategy-tuples $\sigma(k), \sigma(k+1), \sigma(k+2), \ldots, \sigma(k+L)$ is an equilibrium, given $\sigma(0), \sigma(1), \ldots, \sigma(k)$, is at least ε^L, for all k and for all histories $\sigma(0), \sigma(1), \ldots, \sigma(k)$. It follows that the probability that an equilibrium is *not* reached within the first mL ($m \geq 1$) steps of a best-reply path is no more than $(1-\varepsilon^L)^m$ and thus tends to zero as m goes to infinity. In particular, it follows from Theorem 3 that a best-reply improvement path in an (unweighted) congestion game converges to an equilibrium with probability one. For example, an infinite path in the game shown in Fig. 1 can persist only if the order of the deviators is exactly as shown. Any deviation from that order would result in reaching an equilibrium. Similarly, Young (1993) showed that, in WA games, an equilibrium is almost surely reached when simultaneous deviations of several players are allowed, and the strategy played by each player is a best-reply strategy against some k strategy-tuples out of the most recent m strategy-tuples played, provided that these k strategy-tuples are randomly chosen and the sampling by each player is sufficiently incomplete (in the precise sense that $k/m \leq 1/(L+2)$).

If players occasionally make mistakes, deviating to strategies which are not best reply strategies, then the concept of an equilibrium strategy-tuple should be replaced by that of a stationary distribution of strategy-tuples. Such mistakes can be made at random (Kandori et al., 1993; Young, 1993) or may result from the players' lack of information. According to the Bayesian approach applied by Cézilly and Boy (1991) players start with some *a priori* estimates concerning the payoffs associated with each strategy, these estimates are later modified to best fit the actual gains, and the modified *a posteriori* estimates are then used for deciding whether and to what strategy the player should deviate. Cézilly and Boy (1991) simulated a situation similar to that experimentally studied by Milinsky (1979, 1984). Their model is based on a six-player, two-strategy game with two types of players. The two types differ in competitive ability: players of one type have a relative feeding rate which is twice that of the other players and thus have a double effect on the congestion vector. There is a unique congestion vector corresponding to the pure-strategy Nash equilibria of this game, namely the one where the congestion in a patch is proportional to the food supply rate (Parker and Sutherland, 1986). The mean congestion vector, based on 500 independent computer simulations, apparently quickly converged to this equilibrium congestion vector. The variance of the congestion vector apparently converged as well, suggesting a convergence in distribution.

In the following section the model introduced in Section 2 is generalized in order to allow for such player-specific contributions to the congestion.

8 Weighted congestion games

In the model considered thus far all players have a similar influence upon the congestion. This model may be generalized by introducing *weights*, which are positive constants $\beta_1, \beta_2, \ldots, \beta_n$, and modifying the definition of the congestion vector by setting

$$n_j = \sum_{\substack{i \\ \sigma_i = j}} \beta_i, \qquad j = 1, 2, \ldots, r.$$

Inspection of the argument in Section 5, the proof of Theorem 1, and the remark at the end of Section 3 shows these three to hold good, *mutatis mutandis*, for the case in hand. Thus *weighted congestion games* involving only two players, involving only two strategies, or where the players have equal payoff functions possess the finite improvement property or (at least) the finite best-reply property and therefore possess a Nash equilibrium in pure strategies, which may be reached by constructing a maximal best-reply improvement path.

This, however, is not the case in general. Even a three-player, three-strategy weighted congestion game may not possess a pure-strategy Nash equilibrium. For an example, refer to the game in Fig. 2. For each player in this game there are effectively only two strategies (the third strategy invariably yields a minimal payoff). Let the first of these strategies be called the left strategy of the player under consideration and the second one the right strategy. Referring to the inequalities in Fig. 2, it is readily verified that it is always optimal for a player to play the strategy which is diametrically opposite to the strategy played by the player who precedes him (in the sense that the second player precedes the third one, the first player precedes the second one, and the third player precedes the first one). For example, if the third player plays left (first strategy) then right (third strategy) is a unique best-reply strategy for the first player, and if he plays right (third strategy) then left (second strategy) is the best reply. As the number of players is odd, an equilibrium clearly does not exist.

Weighted congestion games are compared with unweighted congestion games (equal weights) in the table below. Recall that FIP (the finite improvement property) implies FBRP (the finite best-reply property), which implies WA (the game is weakly acyclic). No property implies the preceding one. All three properties imply the existence of a pure-strategy Nash equilibrium. In each case, the strongest property is given.

	Unweighted congestion games	Weighted congestion games
Equal sets of payoff functions	FIP	FIP
Two strategies	FIP	FIP
Two players	FBRP	FBRP
General case	WA	—

References

[1] Cézilly, F., and Boy, V. (1991). "Ideal Free Distribution and Individual Decision Rules: A Bayesian Approach," *Acta Œcologia* **12**, 403–410.

[2] Fretwell, S. D., and Lucas, H. L., Jr (1969). "On Territorial Behavior and Other Factors Influencing Habitat Distribution in Birds," *Acta Biotheor.* **19**, 16–36.

[3] Godin, J-G. J., and Keenleyside, M. H. A. (1984). "Foraging on Patchily Distributed Prey by a Cichlid Fish (Teleostei, Cichlidae): A Test of the Ideal Free Distribution Theory," *Anim. Behav.* **32**, 120–131.

[4] Harper, D. G. C. (1982). "Competitive Foraging in Mallards: 'Ideal Free' Ducks," *Anim. Behav.* **30**, 575–584.

[5] Kandori, M., Mailath, G. J., and Rob, R. (1993). “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica* **61**, 29–56.

[6] Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge, UK: Cambridge Univ. Press.

[7] Milinsky, M. (1979). “An Evolutionary Stable Feeding Strategy in Sticklebacks,” *Z. Tierpsychol.* **51**, 36–40.

[8] Milinsky, M. (1984). “Competitive Resource Sharing: An Experimental Test of a Learning Rule for ESSs,” *Anim. Behav.* **32**, 233–242.

[9] Milinsky, M., and Parker, G. A. (1991). “Competition for Resources,” in *Behavioural Ecology: An Evolutionary Approach* (J. R. Krebs and N.B. Davies, Eds.), 3rd ed. Oxford: Blackwell.

[10] Monderer, D., and Shapley, L. S. (1996). “Potential Games,” *Games Econ. Behav.* **14**, 124–143.

[11] Parker, G. A., and Sutherland, W. J. (1986). “Ideal Free Distributions when Individuals Differ in Competitive Ability: Phenotype-Limited Ideal Free models,” *Anim. Behav.* **34**, 1222–1242.

[12] Rosenthal, R. W. (1973). “A Class of Games Possessing Pure-Strategy Nash Equilibrium,” *Int. J. Game Theory* **2**, 65–67.

[13] Sutherland, W. G., Townsend, C. R., and Patmore, J. M. (1988). “A Test of the Ideal Free Distribution with Unequal Competitors,” *Behav. Ecol. Sociobiol.* **23**, 51–53.

[14] Young, H. P. (1993). “The Evolution of Conventions,” *Econometrica* **61**, 57–84.

Strategy	*First*	*Second*	*Third*
Deviator *Player 1* $S_{12}(1) < S_{13}(1)$	2 3	1	
Player 2 $S_{21}(2) < S_{23}(2)$	2 3		1
Player 3 $S_{31}(1) < S_{32}(1)$	3		2 1
Player 1 $S_{12}(2) > S_{13}(2)$		3	2 1
Player 2 $S_{21}(1) > S_{23}(1)$		1 3	2
Player 3 $S_{31}(2) > S_{32}(2)$	2	1 3	

Figure 1: An infinite best-reply improvement path in a three-player, three-strategy unweighted congestion game. The path is generated by the six strategy profiles shown, endlessly repeated. The payoff functions satisfy $S_{13}(1) > S_{12}(1) > S_{12}(2)$ $> S_{13}(2) > S_{11}(n_1)$ (first player), $S_{21}(1) > S_{23}(1) > S_{23}(2) > S_{21}(2) > S_{22}(n_2)$ (second player), and $S_{32}(1) > S_{31}(1) > S_{31}(2) > S_{32}(2) > S_{33}(n_3)$ (third player), for all n_1, n_2, and n_3. The inequality relevant to each step, the one that guarantees that the unique deviator strictly increases the payoff he receives, is shown on the left. The strategy-tuples $(3,1,2)$ and $(2,3,1)$ are equilibria of this game. (Numerical example: $S_{ij}(n_j) = ((j-i) \bmod 3) \cdot (1/n_j - 2/3)$ for $i \neq j$; $S_{ii}(n_i) = -10$ for all i.)

Strategy	*First*	*Second*	*Third*
Deviator *Player 1* $S_{12}(3) > S_{13}(4)$		2	1 3
Player 2 $S_{22}(3) < S_{23}(5)$		1 2	3
Player 3 $S_{31}(3) > S_{33}(5)$		1	2 3
Player 1 $S_{12}(1) < S_{13}(3)$	3	1	2
Player 2 $S_{22}(2) > S_{23}(3)$	3		1 2
Player 3 $S_{31}(3) < S_{33}(4)$	3	2	1

Figure 2: An infinite best-reply improvement path in a three-player, three-strategy weighted congestion game with weights $\beta_1 = 1$, $\beta_2 = 2$, and $\beta_3 = 3$. Assuming that the inequalities on the left hold and that $S_{11}(n_1)$, $S_{21}(n_1)$, and $S_{32}(n_2)$ are minimal, pure-strategy Nash equilibria do not exist. (Numerical example: $S_{12}(n_2) = 1/15n_2 + 1/4$, $S_{22}(n_2) = 1/n_2 - 3/20$, $S_{31}(n_1) = 2/3n_1$, $S_{i3}(n_3) = 1/n_3$ $(i = 1,2,3)$, $S_{ij}(n_j) = -10$ otherwise.)