Minmax Payoffs of a Location Game

Shuchi Chawla[1] Uday Rajan[2] R. Ravi[3]
Amitabh Sinha[4]

May 21, 2003
CMU-CS-03-141

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

We consider a two-player, sequential location game, with n stages. At each stage, players 1 and 2 choose locations from a feasible set in sequence. After all moves are made, consumers each purchase one unit of the good from the closest location. Since player 1 has a natural first-mover disadvantage here (player 2 can obtain a payoff of $\frac{1}{2}$ just by replicating player 1's moves), we examine her minmax payoff. When the number of stages is known to both players we show that (i) if the feasible locations form a finite set in $\mathbf{R}^d$, player 1 must obtain at least $\frac{1}{d+1}$ in the single-move game (ii) in the original Hotelling game (uniformly distributed consumers on the unit interval), player 1 obtains $\frac{1}{2}$ even in the multiple stage game. However, player 1's minmax payoff suffers if she does not know the number of moves, but player 2 does. In the Hotelling game, where the number of stages is either 1 or 2, player 1's payoff falls to $\frac{5}{12}$. If she has no information at all about n, we provide a lower bound for her minmax payoff: it must at least equal half the payoff of the single-stage game.

[1]Computer Science Department, E-mail: shuchi@cs.cmu.edu, Tel: 412-268-3564
[2]GSIA, E-mail: urajan@andrew.cmu.edu, Tel: 412-268-5744
[3]GSIA, E-mail: ravi@cmu.edu, Tel: 412-268-3694
[4]GSIA, E-mail: asinha@andrew.cmu.edu, Tel: 412-268-6895

Keywords: Game Theory, Location Games, Hotelling Game, Min-Max Payoffs, Online Games

1 Introduction

Starting with the classic Hotelling model (Hotelling, 1929), there is an extensive literature on location games. These games have been applied in several different contexts, including firms competing in a market (Gabscewicz and Thisse, 1992, provide a survey), political competition among parties or candidates (see Shepsle, 1991, for a survey), and facility location (surveyed by Eiselt, Laporte, and Thisse, 1993).

In this paper, we consider minmax payoffs in a sequential location game with two players. After players have chosen their locations, each consumer buys one unit of the product from the closest player, breaking ties uniformly at random. We consider the game without prices: Each player maximizes its market share. Since we allow players to locate at previously occupied locations, it is immediate that player 1 has a first-mover disadvantage in this game. By replicating the moves of player 1, the second player obtains a payoff no worse than $\frac{1}{2}$. Hence, we focus on the minmax payoff of player 1.

We consider the location game without prices. This version is commonly applied to, e.g., political contests and the facility location problem. As Osborne and Pitchik (1987) show, the (simultaneous-move) game with prices may not possess a pure strategy equilibrium. With mixed strategy equilibria, the range of possible outcomes may be large. Further, characterizing the set of mixed strategy equilibria can be difficult. For a similar reason, we consider the sequential rather than simultaneous location game.[1]

We first examine a class of games in which the set of feasible locations is finite, and contained in $\mathbf{R}^d$. Without loss of generality, consumers are distributed over $\mathbf{R}^d$ (so there are d attributes of the product a consumer cares about). In the single-stage game (with each player choosing just one location), we characterize completely the set of feasible minmax payoffs for player 1 over all choices of consumer distribution and location set. In this case, the minmax payoff of player 1 is equivalent to her payoff in a subgame-perfect equilibrium.

We show that there exists such a location game in $\mathbf{R}^d$, such that observed market shares are a result of a subgame-perfect equilibrium of this game if and only if the share of the first mover is between $\frac{1}{d+1}$ and $\frac{1}{2}$, and the shares of the players sum to 1. That is, over all location games in d-dimensional Euclidean space, the minimum payoff to player 1 in a subgame-perfect equilibrium is $\frac{1}{d+1}$, and the maximum is $\frac{1}{2}$. Further, for any $y \in [\frac{1}{d+1}, \frac{1}{2}]$, there exist instances of the game such that $r_1 = y$. With a location set in $\mathbf{R}^2$, player 1 must obtain at least $\frac{1}{3}$ of the payoff.

This model is, therefore, testable in the terminology of Brown and Matzkin (1996), who use the Tarski-Seidenberg theorem to provide testable restrictions on the equilibrium manifold of an exchange economy. The Tarski-Seidenberg theorem holds for a finite system of polynomial inequalities, so cannot be used in our setting, since the consumer distribution may be continuous. We therefore provide a direct proof of our result.

This result provides an upper bound for the size of the first-mover disadvantage in such a game. Entry timing games are often characterized by a trade-off between factors that imply a first-mover advantage (for example, in the political context, an early entrant has more time to raise money) and those that lead to a disadvantage. Our result implies that, ceteris paribus, if the payoff increase as a result of a first-mover advantage exceeds $\frac{d-1}{2(d+1)}$ (so that total payoff exceeds $\frac{1}{2}$), players should seek immediate entry in the single-stage game.

We then consider a multi-stage game in which the two players move sequentially at each stage, with player 1 picking a location first, followed by player 2. General results on multi-stage games

[1] Prescott and Vischer (1977) show that the outcomes of a sequential location game can differ significantly from those that obtain in a simultaneous move game.

may not be feasible. In particular, player 1's payoff need not be monotone in the number of stages. We provide two examples to demonstrate this. In one, we construct a game, in which, player 1 obtains $\frac{1}{2}$ in a subgame-perfect equilibrium of the single-stage, but only $\frac{1}{3}$ in the two-stage game. Conversely, we exhibit a game in which player 1's payoff converges to $\frac{1}{2}$ as the number of stages grows.

In the original Hotelling game (with the location set being the unit interval, and consumers uniformly distributed over this interval), we show that in the n-move game, for any n, the minmax payoff of player 1 is $\frac{1}{2}$. In fact, we demonstrate a set of locations such that, if firm 1 occupies each location in this set, it obtains a payoff of at least $\frac{1}{2}$.

Such games have also been studied in computational geometry, under the label "Voronoi games." In these games, the location set is continuous, and the consumers are assumed to be uniformly distributed over some compact set. Co-location of players is not permitted. Cheong et al. (2002), show that when the Voronoi game is played on a square with uniform demand, with a large enough number of moves, and the second player locates all her points after observing all of player 1's moves, player 2 obtains a payoff of at least $\frac{1}{2} + \alpha$ for a fixed constant α. Some of the results we obtain here are cited as open questions by Cheong et al. In particular, we characterize the value of the sequential game and the corresponding optimal strategies, when played in a high dimensional space. For the Voronoi game on the uniform line and uniform circle, Ahn et al. (2001) show that player 1 has a strategy which guarantees her a payoff of strictly more than $\frac{1}{2}$, while player 2 can get a payoff arbitrarily close to $\frac{1}{2}$ without actually getting $\frac{1}{2}$. Variations of the original single-move Hotelling game with multiple players, have also been considered under the name of "competitive facility location." Eiselt et al. (1993) is an excellent survey of some of this work.

We next consider an "online" game, in the same "adversarial" spirit as the online algorithms literature in computer science (Borodin and El Yaniv, 1998). Much of this literature examines single-player decision problems, with nature being an adversary that chooses the input to minimize the player's payoff (or maximize her cost). The player must therefore make decisions that are "robust" with respect to future inputs. Single-player online games studied previously include location games where demand arrives over time (Mettu and Plaxton, 2000) and auctions (Bar-Yossef et al., 2002).

To extend this framework to our two-player game, we assume that player 2 knows exactly the number of stages, but player 1 knows only that the number of stages is in some feasible set. In this case, the minmax payoff of player 1 contains an additional minimization over the set of stages. Hence, this minmax payoff is no longer interpretable as occurring in a subgame-perfect equilibrium.

Suppose player 1 knows that there are one or two stages to the game, whereas player 2 knows the actual number of stages. Then, even in the original Hotelling game, player 1 can no longer guarantee a payoff of $\frac{1}{2}$; in fact, we show that her minmax payoff is $\frac{5}{12}$. Finally, suppose player 1 has no information about the number of stages (i.e., she believes that this can be any positive integer). By replicating the previous moves of player 2, player 1 obtains a payoff no worse than half the payoff it gets in an equilibrium of the single-stage game. This provides a lower bound for player 1's minmax payoff.

The rest of this paper is organized as follows. We begin by describing the model and definitions in Section 2. In Section 3 we study the multiple move game when both players know the number of moves. In Section 4, we extend these results to the online game, where player 1 does not know n. We conclude in Section 5.

2 Preliminaries

Consider $\mathbf{R}^d$ with $d \geq 1$, endowed with the Euclidean distance function, δ. Consumers are distributed on $\mathbf{R}^d$, with distribution $F(\cdot)$ defined over the Borel $\sigma-$algebra on $\mathbf{R}^d$. Without loss of generality, the total mass of consumers is normalized to 1.

There are two players. $L \subset \mathbf{R}^d$ denotes a compact set of points at which players may locate.[2] The game has n stages. At each stage, the players move in sequence. First, player 1 chooses a location in L, and then player 2 responds. At any stage, either player is allowed to choose a location already occupied by either of the players. The game is therefore represented as a 4-tuple, (n, d, L, F).

Let s_i denote the location chosen by player 1 at stage i, and t_i the location chosen by player 2. Let S_i and T_i denote the first i moves of players 1 and 2 respectively, with $S_0 = T_0 = \emptyset$. A pure behavior strategy for player 1 at stage i is a map $a_i : S_{i-1} \times T_{i-1} \to L$. Similarly, a pure strategy for player 2 at stage i is a map $b_i : S_i \times T_{i-1} \to L$. A pure strategy for player 1 in the game as a whole is denoted $A = (a_1, \ldots, a_n)$ and similarly for player 2.

After each player has chosen its n locations, each consumer buys 1 unit of the good from the closest location. If the closest location is not unique, the consumer randomizes with equal probability over the set of closest locations.

Given a multiset Y of locations and some point v in $\mathbf{R}^d$, define $\delta(v, Y) = \min_{y \in Y} \delta(v, y)$ as the distance between v and the point in Y closest to v. Let $\kappa_Y(v) = |\{y \in Y : \delta(v, y) = \delta(v, Y)\}|$ be the number of points in Y which are at minimum distance from v. The demand gathered by a point $y \in Y$ is defined as $r(y, Y \setminus \{y\}) = \int_{v \in \mathbf{R}^d : \delta(v,y) = \delta(v,Y)} \frac{1}{\kappa_Y(v)} dF(v)$. Now let S and T be the locations chosen by player 1 and player 2 respectively. Then, player 1's payoff is given by $r(S, T) = \sum_{s \in S} r(s, S \cup T \setminus \{s\})$. Player 2's payoff is $r(T, S) = 1 - r(S, T)$. Note that by definition, for any move x and set of moves Y, we have $r(x, Y) \leq r(x, y)\ \forall y \in Y$.

The strategy choices of the two players, a and b, imply chosen locations, $S(a, b)$ and $T(a, b)$ respectively. Notationally, for convenience, we often suppress the dependence of S, T on a, b. The minmax payoff of player is defined as $\underline{r}_1 = \max_a \min_b\ r(S(a, b), T(a, b))$.

Since this is a constant-sum game, a strategy of player 2 that minimizes the payoff of player 1 must maximize the payoff of player 2. Hence, when n is known to both players, the strategies $(\hat{a}, \hat{b})$ that lead to player 1 earning its minmax payoff constitute a subgame-perfect equilibrium of the game.

3 Known number of stages

In this section, we examine the game when the number of stages is known to both players. First, suppose there is a single stage in the game, so that each player moves only once. The single-stage game is particularly suited to model political competition, where a location is interpreted as choosing a political platform (see Shepsle, 1991, for a survey). Whereas much of the spatial literature on political competition has considered the simultaneous-move game, Osborne (1993) notes that a sequential game may be more appropriate. Osborne considers a multi-agent game in which agents can choose their platforms at any time, but platforms once chosen cannot be

[2] Without loss of generality, we assume that L spans $\mathbf{R}^d$ (in other words, the convex hull of L is d-dimensional). Otherwise, we can project the d-dimensional space orthogonally to the subspace spanned by L. The orthogonal projection π has the property that for any two location points $l_1, l_2 \in L$ and a demand point $x \in \mathbf{R}^d$, $\delta(l_1, x) \leq \delta(l_2, x) \Leftrightarrow \delta(l_1, \pi(x)) \leq \delta(l_2, \pi(x))$. Thus payoffs and equilibrium strategies in the game remain unaffected.

changed. In focusing on the minmax payoff to player 1, we essentially bound the size of the first mover disadvantage in this model.

We first consider the case of a finite location set.[3] Let $G_d(1) = (1, d, L, F)$ denote an instance of the single-stage location game in d-dimensional Euclidean space, where L is a finite location set and d is also finite. Let $\mathcal{G}_d$ denote the set of such games.

It is clear that $\underline{r}_1 \leq \frac{1}{2}$, since player 2 can ensure $r_2 = \frac{1}{2}$ via the strategy $b = a$, which replicates each move of player 1. How low can the minmax payoff of player 1 be? The following example shows that, when the location set is in $\mathbf{R}^2$, player 1's payoff can be as low as $\frac{1}{3}$.

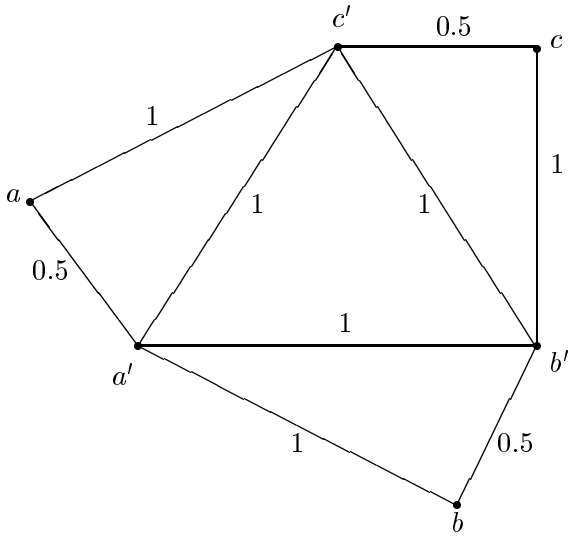

Figure 1: A location game in the Euclidean plane. Points a, b, and c have demands x, $\frac{1}{2}(1-x)$ and $\frac{1}{2}(1-x)$ respectively, and, $L = \{a', b', c'\}$. Lines are labeled by the Euclidean distance between their endpoints.

Example 1 Consider the game given by Figure 1, with $L = \{a', b', c'\}$, and $f(a) = f(b) = f(c) = \frac{1}{3}$, where $f(v)$ denotes the density of demand at v. Player 2's best response is as follows: If Player 1 chooses a', player 2 chooses b'; if player 1 chooses b', player 2 chooses c'; otherwise, player 2 chooses a'. Given this, player 1 is indifferent over $\{a', b', c'\}$. Whichever location it chooses, it obtains a payoff of $\frac{1}{3}$, with player 2 obtaining $\frac{2}{3}$.

In fact, we show that this game represents the worst case for player 1 over all such location games in $\mathbf{R}^2$. That is, there does not exist a demand distribution and a finite location set in $\mathbf{R}^2$, such that player 1 obtains a subgame-perfect equilibrium payoff strictly less than $\frac{1}{3}$ in this single-move location game. The result extends more generally: in $\mathbf{R}^d$, player 1 must obtain at least $\frac{1}{d+1}$, and there exists a game in which it obtains exactly $\frac{1}{d+1}$ (so the bound is tight).

[3] Finiteness of the location set is necessary to prove Theorem 1 below. The demand distribution $F(\cdot)$ may be continuous.

Recall that, when the number of stages is known to both players, the minmax payoff of player 1 is identical to its payoff in a subgame-perfect equilibrium. We therefore state our result in terms of subgame-perfect equilibrium payoffs.

Theorem 1 *There exists a location game $G_d(1) \in \mathcal{G}_d$ such that r_1, r_2 are payoffs in a subgame-perfect equilibrium of $G_d(1)$ if and only if $r_1 \in [\frac{1}{d+1}, \frac{1}{2}]$ and $r_2 = 1 - r_1$.*

Proof of Theorem:

It is immediate that, in any equilibrium, $r_1 + r_2 = 1$. Hence, we prove that $r_1 \in [\frac{1}{d+1}, \frac{1}{2}]$.

"If" part:

Given a value $x \in [\frac{1}{d+1}, \frac{1}{2}]$, we construct a game $G_d(1)$ for which $r_1 = x$. This essentially reconstructs Example 1 in d dimensions. We first construct the game in the $(d+1)$-dimensional Euclidean space (for ease of exposition), then project it down to the d-dimensional Euclidean space.

Define $\hat{d} = 1 + \sum_{i=1}^{d} i^2 = 1 + \frac{d(d+1)(2d+1)}{6}$. Fix $\epsilon > 0$ such that $\epsilon \ll 1$.

The set of location points is $L = \{l_1, l_2, \ldots, l_{d+1}\}$. The i^{th} co-ordinate of point l_i is -1, and all other co-ordinates of point l_i are 0.

There are also $d+1$ demand points v_i. Let f represent the density of demand. Set $f(v_1) = x \in [\frac{1}{d+1}, \frac{1}{2}]$ and $f(v_i) = \frac{1}{d}(1-x) \, \forall i > 1$. Demand point v_i has i^{th} co-ordinate $1 - \epsilon$, and for $j \neq i$, the j^{th} co-ordinate is $\epsilon[(j-i) \bmod d]$.

This induces the following distance function between demand points and location points:

$$\delta^2(l_i, v_j) = \begin{cases} 2 - 2\epsilon[1 + (j-i) \bmod d] + \hat{d}\epsilon^2 & : i \neq j \\ \hat{d}\epsilon^2 & : i = j. \end{cases}$$

For any demand point v_j, we can define a precedence relation $\prec_j$ as $l_i \prec_j l_{i'}$ if $\delta(l_i, v_j) < \delta(l_{i'}, v_j)$. It follows that for every j, we have $l_j \prec_j l_{(j+1) \bmod d} \prec_j l_{(j+2) \bmod d} \prec_j \ldots \prec_j l_{(j-1) \bmod d}$. This precedence relation is identical to that induced by a Condorcet voting paradox (Condorcet, 1785) instance with $d+1$ voters and $d+1$ choices.

It is now immediate that $r_1(l_i, l_{(i-1) \bmod d}) = x$ for $i = 1$, and $r_1(l_i, l_{(i-1) \bmod d}) = \frac{1}{d}(1-x)$ for $i > 1$. For $x \in [\frac{1}{d+1}, \frac{1}{2}]$, we have $x \geq \frac{1}{d}(1-x)$. Player 1's equilibrium strategy, therefore, is to choose l_1, and the resulting payoff is $r_1 = x$.

Finally, we obtain our d-dimensional instance by orthogonally projecting D to the d-dimensional hyperplane formed by the points in L. Such a projection reduces each $\delta^2(l_i, v_j)$ by the same amount, and hence preserves the precedence relation $\prec_j$.

"Only if" part:

Note first that, for any $G_d(1) \in \mathcal{G}_d$, we have $r_1 \leq \frac{1}{2}$ in any subgame-perfect equilibrium. By choosing $t = s$, Player 2 earns $r_2 = \frac{1}{2}$, and so can do no worse in equilibrium. Hence, $r_1 \leq \frac{1}{2}$.

We show that for any $G_d(1) \in \mathcal{G}_d$, we have $r_1 \geq \frac{1}{d+1}$ in any subgame-perfect equilibrium. Define a precedence relation on L, as follows: $l \prec l'$ if and only if $r_1(l, l') < \frac{1}{d+1}$. We need to show that there exists a point $l \in L$ such that there is no $l' \in L$ with $l \prec l'$. First, we prove that $\prec$ is *acyclic* on L; that is, there is no sequence of elements $(l_1, l_2, \ldots, l_m)$ in L with $l_1 \prec l_2 \prec l_3 \prec \ldots \prec l_m \prec l_1$.

For any two location points $x, y \in L$, define $H_{xy} = \{v \in \mathbf{R}^d : \delta(x, v) \geq \delta(y, v)\}$ to be the half-space containing points in $\mathbf{R}^d$ which are at least as close to y as to x. For any subset S of $\mathbf{R}^d$, let $F(S) = \int_{v \in S} dF(v)$ represent the total demand of all points in S. It follows that if $x \prec y$, then $F(H_{xy}) > \frac{d}{d+1}$.

Suppose $\prec$ induces a cycle in L; let this cycle be $l_1 \prec \ldots \prec l_m \prec l_1$. For ease of notation, define $l_{m+1} = l_1$, and $L' = \{l_1, \ldots, l_m\}$. Consider the half-spaces $H_{l_i, l_{i+1}}$ for $1 \leq i \leq m$. Given a set X,

let $\overline{X}$ denote the complement of X. Since $F(H_{l_i,l_{i+1}}) > \frac{d}{d+1}$, if must be that $F(\overline{H}_{l_i,l_{i+1}}) < \frac{1}{d+1}$. Hence, for any set K of $d+1$ or fewer indices, $F(\cup_{k\in K}\overline{H}_{l_k,l_{k+1}}) < 1$, so $F(\cap_{k\in K}H_{l_k,l_{k+1}}) > 0$. That is, any $d+1$ or fewer half-spaces must have an intersection with strictly positive demand.

We now appeal to the following theorem due to Helly, 1921 (see also Eckhoff, 1993).

Helly's Theorem: *Let $\mathcal{C}$ be a family of convex sets in $\mathbf{R}^d$, and suppose $\mathcal{C}$ is finite or each member of $\mathcal{C}$ is compact. If every $d+1$ or fewer members of $\mathcal{C}$ have a non-empty intersection, then all members of $\mathcal{C}$ have a non-empty intersection.*

We can now define $\mathcal{C} = \{H_{l_i,l_{i+1}} : 1 \le i \le m\}$ and use Helly's Theorem, since each $H_{l_i,l_{i+1}}$ is convex. It follows that $\cap_{i=1}^m H_{l_i,l_{i+1}}$ is non-empty. Choose any point $o \in \cap_{i=1}^m H_{l_i,l_{i+1}}$, and renormalize all points in L relative to o as the origin. Note that by definition of $H_{l_i,l_{i+1}}$, we have $\delta(l_i, o) \ge \delta(l_{i+1}, o)$, for all i, $1 \le i \le m$. Then, $l_{m+1} = l_1$ implies that all the points l_i are equidistant from the origin and the origin lies on the hyperplanes defined by $\{v \in \mathbf{R}^d : \delta(v, l_i) = \delta(v, l_{i+1})\}$, for $i \in \{1, \ldots, m\}$.

Now let ν_i be the vector $l_{i+1} - l_i$. Then, each half-space is given by $H_{l_i,l_{i+1}} = \{u \in \mathbf{R}^d : u \cdot \nu_i \ge 0\}$. We refer to l_{i+1} and l_i as the location points defining ν_i. Since L' is a cycle, we also have $\sum_{i=1}^m \nu_i = 0$. That is, a positive linear combination of these vectors ν_i sums to zero.

Now, Carathéodory's theorem (see Eckhoff, 1993) implies that there exists a positive combination of at most $d+1$ of the vectors ν_i that sums to zero. Without loss of generality, let these be $\nu_1, \ldots, \nu_{d'}$ $(d' \le d+1)$, and let α_i be positive reals such that $\sum_{i=1}^{d'} \alpha_i \nu_i = 0$. Let the corresponding half-spaces be called H_i.

Next consider any point x lying in $X = \cap_{i=1}^{d'} H_i$. Since $x \cdot \nu_i \ge 0$ for all i, $1 \le i \le d'$, we must have $0 \le x \cdot \alpha_i \nu_i = -\sum_{j \le d', j \ne i} x \cdot \alpha_j \nu_j \le 0$, implying that $x \cdot \nu_i = 0$ for all $i \le d'$. In other words, x is equidistant from the location points defining each vector ν_i, so the entire set X consists only of points equidistant from the location points of the vectors ν_i.

For each half-space H_i, we have $F(H_i) \ge \frac{d}{d+1} + \frac{F(X)}{2}$. Moreover, $F(X) > 0$, since $X \supseteq \cap_{i=1}^m H_{l_i,l_{i+1}}$, and $F(\cap_{i=1}^m H_{l_i,l_{i+1}}) > 0$. So we have $F(\cup_{i<d'} \overline{H_i} \cup X) < \frac{d'-1}{d+1} - \frac{d'-1}{2} F(X) + F(X) \le \frac{d}{d+1} + \frac{F(X)}{2}$. Therefore, $F(\cap_{i<d'} H_i \setminus X) > \frac{1}{d+1} - \frac{F(X)}{2}$. But $\cap_{i<d'} H_i \setminus X$ is disjoint from $H_{d'}$, which has demand greater than $\frac{d}{d+1} + \frac{F(X)}{2}$. This contradicts the fact that the total demand is 1.

We have shown that the relation $\prec$ is acyclic. An acyclic relation on a finite set must contain a point l_0 which is not preceded by any other point $l' \in L$. Such a point can be found by starting at any point $l \in L$, and moving to any point $l' \in L$ such that $l' \prec l$. Since $\prec$ is acyclic and L is finite, this process must terminate at an l_0 such that there is no point $l' \in L$ with $l' \prec l$. This completes the proof of the "only if" part. ∎

Remarks:

1. As part of the proof of Theorem 1, we show that for any game $G_d(1) \in \mathcal{G}_d$, there exists a point $l \in L$, such that a half space defined by l and any other point $l' \in L$ contains at most $\frac{d}{d+1}$ demand. This concept is similar to the concept of *centerpoints*, used widely in Computational Geometry (see Edelsbrunner, 1987). Given a demand distribution with finite support, a centerpoint always exists; however, it may not lie in the set L. We can overcome this problem by picking a point $l \in L$ which is closest to the centerpoint. If such a point l is unique (for example, when L is convex), then the theorem holds. On the other hand, if it is not unique, picking any one of the closest points only implies $r_1 \ge \frac{1}{2(d+1)}$, necessitating a proof based on first principles, as given above.

2. Finiteness of the location set, L, is used in the "if" part of the theorem. The following

example, a variant of the largest number game, indicates that there is no extension to a countably infinite set. Consider the unit interval, $[0,1]$. Let $f(0) = 1$ (so that all demand is at the point 0). Let $L = \{\frac{1}{n}\}_{n \in \mathbf{Z}_+}$, where $\mathbf{Z}_+$ is the set of positive integers. For any point l_1 chosen by player 1, player 2 can find a point closer to 0, and obtain a payoff of 1.

3. In the game in $\mathcal{G}_d$ constructed in the "If" part of Theorem 1, with $x = \frac{1}{d+1}$, consider the payoff of player 1 as the number of moves n increases (with both players knowing n). While the number of moves is less than $d+1$, player 1 can weakly increase her payoff by locating at each stage at a location where she has not located yet. When the number of moves is $d+1$ or more, the strategy of first locating at all points in L and then replicating player 2's previous move guarantees a payoff which converges from below to $\frac{1}{2}$ as n increases.

Given the last remark above, one might conjecture that, in the multi-stage game, the minmax payoff of player 1 is weakly increasing in the number of moves, n. However, the following example demonstrates that this is not always true.

Example 2 Consider two replicas of the game in Example 1, with location sets $L_i = \{a_i', b_i', c_i'\}$ for $i = 1, 2$. The demand density is $\frac{1}{6}$ at each of the points in $D_i = \{a_i, b_i, c_i\}$, for $i = 1, 2$. Further, let a_j' be the closest location point in L_j to the demand points D_i, for $i = 1, 2$ and $j \neq i$. Let $\delta(a_i, a_j') > 2$ for $i = 1, 2$ and $j \neq i$, so that the points in L_j are sufficiently far from the points in D_i.

Suppose $n = 1$, so that each player moves just once. Player 1's optimal action is to choose either a_1' or a_2'. If player 1 chooses a_1', player 2's best response is to choose any of $\{a_1', a_2', b_2', c_2'\}$, with a corresponding best response set if player 1 chooses a_1. In either case, player 1 obtains a payoff of $\frac{1}{2}$.

Now, suppose $n = 2$. Without loss of generality, suppose player 1 chooses a location in L_1 with her first move. Conditional on choosing a point in L_1, locating at a_1' is an optimal action for player 1. Now, player 2 responds by locating at b_1'. Consider player 1's best response. If she chooses any point in L_2, player 2 will choose the corresponding point in L_2 such that it obtains $\frac{2}{3}$ of the demand closest to each of L_1 and L_2, and hence captures a payoff of $\frac{2}{3}$ in the game. If she chooses any point in L_1, player 2 will then choose a_2', obtaining all of the demand closest to L_2, and at worst $\frac{1}{3}$ of the demand closest to L_1, for an overall payoff no worse than $\frac{2}{3}$. Hence, player 1 can obtain no more than $\frac{1}{3}$ in the 2-move game.

Example 2 and Remark 3 above suggest that there is no general result on the equilibrium payoffs as n increases. Since results on the general n-move game are difficult to obtain, we next study the game in Hotelling's original setting, where the demand is distributed uniformly over $[0,1]$, and $L = [0,1]$. Let $H(n) = (n, 1, [0,1], U[0,1])$ denote the Hotelling game with n rounds, $L = [0,1]$, and $F(x) = x$ for $x \in [0,1]$. We first show that there is no second-mover advantage in $H(n)$. In particular, for any fixed n, there exists a set of location points S that player 1 can choose which implies that its payoff is at least $\frac{1}{2}$, regardless of the strategy of player 2.

Theorem 2 *For the game $H(n)$, we have $\underline{r}_1 = \frac{1}{2}$.*

Proof: Consider $S = (s_1, s_2, \ldots, s_n)$, where $s_i = \frac{1}{2n} + \frac{(i-1)}{n}$. This divides the unit line into $n+1$ intervals—the two border intervals are of length $\frac{1}{2n}$, while the internal intervals are of length $\frac{1}{n}$.

Let the second player's chosen location points be given by $T = (t_1, \ldots, t_n)$. We will show that each point t_i gets payoff at most $\frac{1}{2n}$. This implies that $r_1 \geq \frac{1}{2}$. As observed earlier, player 2 can obtain a payoff of $\frac{1}{2}$ by simply replicating each of player 1's moves (i.e. set $t_i = s_i$ for each i). It

is easy to see that, even in the absence of any points t_i, the total demand captured by point s_i in this case is at most $\frac{1}{n}$ for any i.

Consider the point t_i. First consider the case when $t_i = s_j$ for some j. Clearly, the market share of point t_i is at most $\frac{1}{2n}$ from our observation above. Next consider the case that t_i lies in one of the border intervals. Again, since the length of these intervals is $\frac{1}{2n}$, the market share of t_i is at most $\frac{1}{2n}$.

Finally, consider the case when t_i lies in some interval (s_j, s_{j+1}). If there is at least one other point t_k in this interval, t_i and t_k may share the total demand in that interval, each getting at most $\frac{1}{2n}$. If t_i is the only point in this interval, then, it gets $\frac{1}{2}(s_{j+1} - t_i)$ demand from the left and $\frac{1}{2}(t_i - s_j)$ demand from the right. Combining the two, we have that t_i gets at most $\frac{1}{2n}$ of the demand. Thus player 2 obtains a payoff no greater than $\frac{1}{2}$. ■

A similar result was obtained independently by Ahn et al. (2001), in the context of Voronoi games, which differ from our location games in that co-location is not allowed in Voronoi games.

Note that player 1's strategy in Theorem 2 is independent of player 2's strategy T. Thus, player 1's strategy guarantees her a payoff of at least $\frac{1}{2}$ even when both players move simultaneously at each round, or indeed, even if the order of moves is completely arbitrary.

4 Player 1 does not know the number of stages

Next, we consider an "online" version of the location game. In this game, the number of stages, n, is known to player 2 but not to player 1. Instead, player 1 merely knows that $n \in N$, where N is some feasible set for the number of stages.

In terms of minmax payoffs, this changes the flavor of the game completely. The minmax payoff of player 1 now contains an additional uncertain element, the number of stages in the game. As a result, the minmax payoffs in the game can no longer be thought of as equilibrium payoffs. Given location sets S, T for the two players, and a known number of stages n, let $r_1(S, T, n) = r(S_n, T_n)$ denote player 1's payoff in the game. Then, when player 1 does not know the number of stages, her minmax payoff is given by $\underline{r}_1(N) = \max_a \min_b \min_{n \in N} r_1(S(a, b), T(a, b), n)$.

To illustrate the nature of the difficulty in analyzing this case, suppose first that $N = \{1, 2\}$, that is, player 1 knows that the number of stages is either 1 or 2. In contrast with Theorem 2, the following theorem shows that, in the set-up of the original Hotelling game H, player 1 can no longer ensure a payoff of $\frac{1}{2}$ across all possible outcomes.

Theorem 3 *Suppose player 1 knows that $n \in N = \{1, 2\}$, and player 2 knows n. Then, in the game $H(N)$, we have $\underline{r}_1(N) = \frac{5}{12}$.*

Proof: We first show that $\underline{r}_1(N) \geq \frac{5}{12}$. Consider the following strategy for player 1. It first locates at $s_1 = \frac{1}{2}$. If $n = 1$, player 2 will also choose $t_1 = \frac{1}{2}$, so player 1 earns exactly $\frac{1}{2}$ (that is, $\underline{r}_1(1) = \frac{1}{2}$).

Suppose $n = 2$. Without loss of generality (w.l.o.g.), player 2's first move is to $t_1 \leq s_1$. If $t_1 > \frac{1}{3}$, player 1 then chooses $s_2 = t_1 - \epsilon$, for some small $\epsilon > 0$. Now, regardless of player 2's second move, player 2 obtains a payoff at most $\frac{1}{2} + (\frac{1}{2} - t_1)/2 \leq \frac{7}{12}$. By locating at $\frac{1}{2} + \epsilon$, for some small $\epsilon > 0$, player 2 obtains a payoff that approximates (but is strictly less than) $\frac{7}{12}$.

On the other hand, if player 2 first locates at $t_1 \leq \frac{1}{3}$, then player 1 then chooses $s_2 = \frac{5}{6}$. Now, if player 2 chooses $t_2 > s_1$, it earns a payoff at most $\frac{7}{12}$. If $t_2 = s_1$, its payoff is at most $\frac{13}{24}$. For any other point $t_2 < s_1$, its payoff is at most $\frac{1}{2}$. In any case, $r_2 \leq \frac{7}{12}$, so $\underline{r}_1(2) \geq \frac{5}{12}$. Hence, $\underline{r}_1(N) \geq \frac{5}{12}$.

Next we show that $\underline{r}_1(N) \leq \frac{5}{12}$. Suppose not. Then, player 1's first move must be to some point in $(\frac{5}{12}, \frac{7}{12})$ (else $\underline{r}_1(1) \leq \frac{5}{12}$). W.l.o.g, suppose player 1's first move is to $s_1 \in (\frac{5}{12}, \frac{1}{2}]$. Suppose

$n = 2$, and consider the following sequence of play. Player 2 chooses $t_1 = \frac{2}{3}(1 - s_1) < s_1$. At the second stage, if player 1 moves to $s_2 < s_1$, then player 2 makes its second move to $t_2 = s_1 + \epsilon$ for some small $\epsilon > 0$. Otherwise, player 2 moves to some $t_2 > s_1$ that obtains maximum payoff. The latter payoff is at least $\frac{1}{3}(1 - s_1)$. A simple calculation again shows that in either of these cases, player 2 earns a payoff of at least $\frac{2}{3} - \frac{s_1}{6} \geq \frac{7}{12}$. ■

The above theorem shows that if H is played with the number of stages restricted to being no more than 2, then player 1's minmax payoff is lower than $\frac{1}{2}$. What if player 1 has no information at all about the number of stages? The techniques used for the above theorem do not extend easily to larger n, since the number of cases increases rapidly as n increases. However, we show below that a simple strategy guarantees a payoff of $\frac{1}{4}$ to player 1 irrespective of the number of rounds in the game.

We first show a more general theorem that applies to all sequential two-player location games, including H and those in $\mathcal{G}_d$. The theorem shows that in a multi-stage game, player 1 must obtain at least $\frac{1}{2}$ of its payoff in the single-stage game, even when it has no knowledge of the number of stages (that is, the set of feasible stages, N, is the set of positive integers). We prove the theorem by exhibiting a particular strategy that earns this payoff: locate at the single-stage equilibrium location, then replicate each move of player 2.

Theorem 4 *Suppose that, in a subgame-perfect equilibrium of a single stage location game, player 1 earns $r_1 = \rho$. Consider the multiple-stage game in which player 1 knows that $n \in Z_+$, but player 2 knows n. In this game, $\underline{r}_1(Z_+) \geq \frac{\rho}{2}$.*

Proof: Consider the following strategy for player 1. At stage 1, she chooses a location s_1 that yields the payoff of a single-stage equilibrium, ρ. For $i > 1$, player 1 replicates player 2's previous move, so that $s_i = t_{i-1}$. For any location $y \in S \cup T$, we have $r(y, S \cup T \setminus \{y\}) \leq r(y, s_1) \leq 1 - \rho$.

Now, $r_1(S, T, n) \geq \sum_{i=2}^{n} r(s_i, S \cup T \setminus \{s_i\}) = \sum_{i=1}^{n-1} r(t_i, S \cup T \setminus \{t_i\}) = r(T, S) - r(t_n, S \cup T \setminus \{t_n\})$. This implies $2r_1(S, T, n) \geq 1 - r(t_n, S \cup T \setminus \{t_n\}) \geq \rho$. Thus, $\underline{r}_1(Z_+) \geq \min_n r_1(S, T, n) \geq \frac{\rho}{2}$. ■

An immediate implication is that player 1 can obtain at least $\frac{1}{2(d+1)}$ in any game in $\mathcal{G}_d$, and at least $\frac{1}{4}$ in the game H, when she does not know the number of stages.

Corollary 5 *Suppose player 1 has no information about n, but player 2 knows n.*
(i) for any location game $G_d(Z_+) \in \mathcal{G}_d$, we have $\underline{r}_1 \in [\frac{1}{2(d+1)}, \frac{1}{2}]$.
(ii) for the game $H(Z_+)$, we have $\underline{r}_1 \geq \frac{1}{4}$.

5 Conclusion

We have shown that in a one move location game in $\mathbf{R}^d$, player 1 can always guarantee at least $\frac{1}{d+1}$ of the total payoff. If player 1 earns a payoff strictly less that $\frac{1}{d+1}$, this payoff could not have emerged from a subgame-perfect equilibrium of the location game in d-dimensional Euclidean space. Conversely, for every $x \in [\frac{1}{d+1}, \frac{1}{2}]$, there exists a location game such that player 1 obtains a market share exactly x in equilibrium. Hence, the model is testable.

In the multiple-move game on a unit line, when both players know the number of moves, both obtain a payoff of $\frac{1}{2}$ in a subgame-perfect equilibrium. It would be interesting to generalize this result to games in higher dimensions.

The situation changes when player 1 does not know the number of moves. Even if the number of moves is 1 or 2, in the game on a unit line, player 1 obtains a payoff strictly less than $\frac{1}{2}$. However,

we demonstrate a strategy for player 1, using which she can obtain at least half the payoff of the single-move game in a subgame perfect equilibrium. An interesting open problem is to completely characterize this minmax payoff.

References

[1] H. Ahn, S. Cheng, O. Cheong, M. Golin and R. van Oostrum (2001), "Competitive facility location along a highway," in *Proceedings of the* 7^{th} *Annual International Computing and Combinatorics Conference*, LNCS 2108:237–246.

[2] Z. Bar-Yossef, K. Hildrum and F. Wu (2002), "Incentive-compatible online auctions for digital goods," *ACM-SIAM Symposium On Discrete Algorithms, 2002*: 964–970.

[3] A. Borodin and R. El-Yaniv (1998), *Online Computation and Competitive Analysis*, Cambridge, United Kingdom.

[4] D.J. Brown and R.L. Matzkin (1996), "Testable restrictions on the equilibrium manifold," *Econometrica* 64(6):1249-1262.

[5] O. Cheong, S. Har-Peled, N. Linial and J. Matousek (2002), "The one-round Voronoi game," in *Proceedings of the* 18^{th} *Annual Symposium on Computational Geometry*, 97-101.

[6] Marquis de Condorcet (1785), *Essais sur l'application de l'analyse a la probabilite des decisions redues a la pluralite des voix*, Paris.

[7] H. Edelsbrunner (1987), *Algorithms in Combinatorial Geometry*, Springer-Verlag, Berlin.

[8] H.A. Eiselt, G. Laporte and J.-F. Thisse (1993), "Competitive location models: A framework and bibliography," *Transportation Science*, 27:44-54.

[9] J.J. Gabscewicz and J.-F. Thisse (1992), "Location," in R.J. Aumann and S. Hart (eds.) *Handbook of Game Theory with Economic Applications*, North Holland.

[10] J. Eckhoff (1993), "Helly, Radon, and Carathéodory type theorems," in P.M. Gruber and J.M. Wills (eds.) *Handbook of Convex Geometry*, North Holland.

[11] E. Helly (1921), "Uber systeme linearer gleichungen mit unendlich vielen unbekannten," *Monatschr. Math. Phys.*, 31:60–91.

[12] H. Hotelling (1929), "Stability in competition," *Economic Journal*, 39(1):41–57.

[13] R. Mettu and C. Plaxton (2000), "The Online Median Problem," *Foundations Of Computer Science, 2000*: 339–348.

[14] E.C. Prescott and M. Vischer (1977), "Sequential location among firms with foresight," *Bell Journal of Economics*, 8:378–393.

[15] M. Osborne (1993), "Candidate Positioning and Entry in a Political Competition," *Games and Economic Behavior* 5: 133–151.

[16] M. Osborne and C. Pitchik (1987), "Equilibrium in Hotelling's model of spatial competition," *Econometrica*, 55(4): 911–922.

[17] K. Shepsle (1991), *Models of Multiparty Competition*, Harwood Academic Publishers.