OR Spektrum (1998) 20: 101–107

OR Spektrum

Reachability of locational Nash equilibria

H. A. Eiselt, J. Bhadury

Faculty of Administration, University of New Brunswick, P.O. Box 4400, Fredericton, NB, E3B 5A3, Canada

Received: 12 February 1997 / Accepted: 26 August 1997

Abstract. This paper examines the location of duopolists on a tree. Given parametric prices, we first delineate necessary and sufficient conditions for locational Nash equilibria on trees. Given these conditions, we then show that Nash equilibria, provided they exist, can be reached in a repeated sequential relocation process in which both facilities follow short-term profit maximization objectives.

Zusammenfassung. In der Arbeit werden die Standorte von Duopolisten in einem Baum untersucht. Unter der Annahme festgesetzter Preise werden notwendige und hinreichende Bedingungen für Nash Gleichgewichte für Standorte auf Bäumen hergeleitet. Unter Verwendung dieser Bedingungen wird dann gezeigt, daß – angenommen Nash Gleichgewichte existieren – diese in einem wiederholt angewandten sequentiellen Standortfindungsprozeß, in dem beide Duopolisten als Zielfunktion kurzfristige Gewinnmaximierung haben, auch erreicht werden.

"Equilibrium is a place in heaven, but how do we get there from here?"

Key words: Competitive location model, Nash equilibria, stability, reachability

Schlüsselwörter: Wettbewerbsmodelle in der Standorttheorie, Nash Gleichgewicht, Stabilität, Erreichbarkeit

1 Introduction

Competitive location models were introduced by Hotelling (1929) who studied the price-setting and locational behavior of two duopolists, who compete for a common market in the shape of a line segment. Results of the model and their implications were of immediate interest to economists, geographers, political scientists, marketing researchers, and, more recently, operations researchers. Applications of competitive location models range from brand positioning problems and political positioning to models that examine locations of competing facilities within an industry, such as fast-food chains. The main interest in these models is based on their explanatory power. A survey of competitive location models is provided by Friesz et al. (1988) and a framework and taxonomy is found in Eiselt et al. (1993).

Hakimi (1983) was the first to consider competitive location models on a network and hence bring them to the attention to management scientists. The simplest scenario involves two decision makers who locate a fixed, but not necessarily equal, number of facilities each in some given space. One of the solution concepts for this problem is the Nash equilibrium; a locational arrangement in which neither decision maker has an incentive to unilaterally relocate any of his facilities. Nash equilibria are well studied in economic game theory. In the locational context, some aspects of Nash equilibria were investigated by Labbé and Hakimi (1991). Two major questions arise when considering Nash equilibria. The first concerns their existence and their uniqueness. In general, the existence of locational Nash equilibria depends highly on the specific model under consideration, thus reaffirming the well-known sensitivity of Hotelling models. As an example, given fixed and equal prices, a locational Nash equilibrium exists on trees, but may not exist on general graphs. Assuming that at least one Nash equilibrium exist, the second question is then: given arbitrary initial locations of the competitors and a set of objectives followed by the duopolists who relocate according to a given set of rules, will they ever reach any one of the equilibria? If this question can be answered in the affirmative for at least on pair of equilibrium locations, we will refer to the equilibria as *reachable*. Some progress on reachability has been made in the context of voting theory (see, e.g., Tovey 1993). Hakimi (1990) shows that reachability is guaranteed on trees in the case of fixed and equal prices, given a specific demand allocation rule. Using the same rule, Hakimi shows that in a general graph where facilities are permitted to locate any-

Correspondence to: H. A. Eiselt

where, equilibria may not be reached even if each decision maker locates only one facility and the graph is as simple as a cycle with three vertices. Similarly, he shows that an equilibrium may not even be reachable on a tree network in case one decision maker locates one and the other two facilities. In our research, we investigate the simple case of duopolists who each locate a single facility at a vertex of a tree network, and each facility charges a fixed price. This paper will answer the question whether or not equilibria exist and, if so, if they are reachable.

This paper is organized as follows. In the next section we state some useful results concerning Nash equilibria along with some basic concepts that are used later in the paper. In Section 3, we then investigate the reachability of equilibria given two facilities that charge fixed, but unequal, prices.

2 Nash equilibria on trees

In this section we first introduce our basic model. We then restate some useful results on trees and locational Nash equilibria as well as their relations to medians. Consider a tree $T=(V, E)$, where $V=\{v_1, v_2, \dots, v_n\}$ is the set of vertices and $E=\{e_{ij}: v_i, v_j \in V\}$ symbolizes the set of undirected edges. By d_{ij} we denote the length of edge e_{ij} in case v_i and v_j are connected by an edge, and the length of the (unique) path between v_i and v_j in case these two vertices are not connected by an edge. Customers are assumed to be located at the vertices of the tree. Their demands are $w_i > 0 \; \forall \; v_i \in V$ which are satisfied by the facility that offers the lowest full price, i.e., mill price plus transport cost. The reason for requiring positive rather than nonnegative weights at the vertices is that allowing zero weights renders the analysis quite messy without providing further insight. Furthermore, without loss of realism we require that $|V|>2$. As in many competitive location models, transport costs are assumed to be linear with unit transport cost normalized to one, so that it costs d_{ij} dollars to ship one unit from v_i to v_j. Two facilities A and B are assumed to supply customers. The facilities are restricted to locate at vertices of the tree and their current locations are denoted by v_A and v_B, respectively. The mill prices charged at facilities A and B are p_A and p_B; since Hakimi (1990) has already dealt with the case of equal prices, we restrict ourselves to unequal prices p_A and p_B. Without loss of generality let $p_A < p_B$. Whenever customers at a vertex v_i purchase from facility A, we will say that facility A "captures" v_i, a concept first introduced by Stackelberg (1943) and later rediscovered by ReVelle (1986). Formally, facility A captures all demand at vertex v_i, if $p_A + d_{Ai} \leq p_B + d_{Bi}$. Thus A's market area is $M(A)=\{v_i: p_A + d_{Ai} \leq p_B + d_{Bi}\}$ and $M(B)=\{v_i: p_B + d_{Bi} < p_A + d_{Ai}\}$. Note that we have assumed that ties are broken in favor of the less expensive facility A. This tie breaking rule is somewhat arbitrary; we have modeled it after the "incumbent advantage" rule that has been used, yet stirred controversies, for a long time (for a short survey see, e.g., Tirole 1995). The demands (or sales) captured by the duopolists can then be expressed as $S(A) = \sum_{i: v_i \in M(A)} w_i$ and similar for $S(B)$, and the profits are $\wp(A) = p_A S(A)$ and $\wp(B) = p_B S(B)$, respectively. As prices in this paper are parametric, maximization of $\wp(A)$ and $\wp(B)$ is equivalent to maximizing $S(A)$ and $S(B)$.

In this paper, we investigate a process, in which two decision makers sequentially relocate their facilities so as to maximize $S(A)$ and $S(B)$, respectively. This process is repeated until it converges. In particular, we employ three rules in the individual optimization process:

(1) A facility, when given the option to relocate, will do so by maximizing its profit *given* its opponent's current location. The result is referred to as a $(1|1)$ medianoid (see, e.g., Hakimi 1983). Note that this is a short-term view as it pays no attention to the potential reactions of the competing facility.
(2) The facilities move in a sequential manner. The idea is that it takes a certain amount of time for a facility to react to its opponent's action, so that a facility planner has an opportunity to maximize its profit now.
(3) Location at the same vertex is prohibited. While this rule sounds somewhat restrictive, there are good reasons for employing it. One reason is that in the case of fixed and unequal prices, if location at the same vertex were allowed, then the less expensive facility could always completely annihilate its opponent by locating at the same vertex. Moreover, due to the discrete nature of space available for location, it is not realistic to assume that facilities can actually co-locate.

In the following we restate a few definitions and lemmas regarding trees which are used in this paper.

Definition 1. The weight of a subtree $T^k=(V^k, E^k)$ with $V^k \subseteq V$ and $E^k=\{e_{ij}: e_{ij} \in E \text{ and } v_i, v_j \in V^k\}$ is defined as the sum of weights of its vertices, i.e., $w(T^k) = \sum_{i: v_i \in V^k} w_i$. The weight of the entire tree is $w(T)$.

Definition 2. Given some vertex v_k, the subtrees spanned by v_k are obtained by deleting from T the vertex v_k as well as all edges incident to it. The subtrees can then be numbered as $T_1^k, T_2^k, \dots$ with $w(T_1^k) \geq w(T_2^k) \geq \dots$ where ties are broken arbitrarily. A subtree T_i^k is called *heavier* than a subtree T_j^k, $i<j$.

It is also useful to restate the classical.

Lemma 3 (Goldman 1971). A vertex v_q is a median, if and only if $w(T_1^q) \leq \frac{1}{2} w(T)$, i.e., the largest subtree spanned by v_q has a demand that is no more than half of the total market's demand.

Lemma 3 implies immediately

Corollary 4 (Median Location Corollary). For any vertex $v_k \in V$ with $w(T_1^k) > \frac{1}{2} w(T)$ the median $v_q \in T_1^k$.

We are now able to formally define Nash equilibria.

Definition 5. A *locational Nash equilibrium* is a pair of locations $(v_A, v_B) = (v^*, v^{**})$, such that

$$\wp(A: v_A = v^*, v_B = v^{**}) \geq \wp(A: v_A \neq v^*, v_B = v^{**})$$

and

$$\wp(B\colon v_A = v^*, v_B = v^{**}) \geq \wp(B\colon v_A = v^*, v_B \neq v^{**}) .$$

In other words, a locational Nash equilibrium is a locational arrangement in which neither facility can gain an advantage by relocating unilaterally. In some cases, it is easy to find locational Nash equilibria. For instance, the case of equal prices has an easy solution (see, e.g., Wendell and McKelvey 1981, Hakimi 1990).

Consider now the case of unequal but fixed, i.e., parametric, prices. Here, it will be useful to consider for any pair of locations $v_A \neq v_B$ the subdivision of the given tree into three subsets: the *hinterland of A*, the *hinterland of B*, and the *competitive region*. Facility A's hinterland consists of vertices which can be reached from v_A only via v_B. The vertices v_A and v_B are assumed to be included in their own respective hinterlands. All vertices that are not included in either of the two hinterlands are said to be in the competitive region. For the analysis of this case, we need to restate a definition introduced by Eiselt (1992).

Definition 6. Two vertices v_A and v_B are said to be *sufficiently spatially separated (SSS)*, if $d_{AB} > SSS = |p_A - p_B|$. Furthermore, a vertex v_i is said to be *edge-protected*, if $d_{ij} > SSS \; \forall \, e_{ij} \in E$.

In essence, both, the concept of sufficient spatial separation and that of edge-protection guarantee, given our tie breaking rule, that the weaker, i.e., more expensive, facility B is protected from the stronger, i.e., less expensive, facility A. Here, protection refers to the fact that a facility is located so as to enable it to capture its own hinterland. This is somewhat reminiscent of predator – prey models with A preying on B, where B is temporarily protected if it locates SSS away from A (who, in its next move, will attempt to cut out B), and B is permanently protected (at least as long as it does not relocate) if it locates at an edge-protected vertex. The concept of protection is one of the major differences between Hotelling's linear market and models on trees with facilities allowed to locate only at vertices, as linear markets allow facilities to locate arbitrarily close to each other and thus provide no protection whatsoever.

We are now ready to delineate conditions for locational Nash equilibria on trees given differential parametric prices. For convenience, define now a "circle" around v_A with radius SSS by $C_A = \{v_i\colon d_{iA} \leq SSS\}$ and similar circles C_B, C_*, and C_{**} around vertices v_B, v^*, and v^{**}, respectively. These circles are designed, so that if facility B were to locate anywhere inside C_A then it is cut out and A captures the entire market. Similarly, if A locates in C_B facility B is again cut out. Observe that generally, $C_B \subseteq M(B)$ and similar for C_A and $M(A)$.

In the following we characterize locational Nash equilibria. Lemmas 7 and 9 provide necessary and sufficient conditions for a Nash equilibrium to exist. The proofs of both lemmas are provided in the appendix.

Lemma 7. Let at least one edge-protected vertex exist in T. Then a pair of locations (v^*, v^{**}) is a locational Nash equilibrium, if and only if

(a) v^{**} is edge-protected
(b) v^{**} has the heaviest hinterland of any vertex outside C_*.
(c) $M(A) = T_1^{**}$.

For the discussion in the next section, it is useful to prove the following

Lemma 8. Assume that there exists at least one edge-protected vertex in T, but none of the medians v_q is edge-protected. Furthermore, let C_q, resp. C_*, be a circle around v_q, resp. v^*, with radius SSS. Then

(a) $v^* \in C_q$
(b) $v^{**} \notin C_q$
(c) $v^{**} \notin C_*$.

Consider now the case in which there exist no edge-protected vertices in the given tree. We can then prove

Lemma 9. Suppose that there exists no edge-protected vertex in T. Then a (degenerate) locational Nash equilibrium exists if and only if there exists at least one vertex v_i such that $C_i = V$.

Proof. Sufficiency of the condition is easily proved. Let $v_A = v_i$, then the condition implies that regardless of B's location, $S(A) = w(T)$ and $S(B) = 0$. Neither facility can improve, implying that $v^* = v_i$ and v^{**} anywhere else is a locational Nash equilibrium.

To prove necessity, assume that there exists no vertex v_i satisfying the above condition. Then either $S(A) = w(T)$ and $S(B) = 0$ in which case facility B can locate at some $v_i \notin C_A$ and obtain at least $w_i > 0$ thus improving its current sales; or $S(A) < w(T)$ and $S(B) > 0$ in which case A can relocate adjacent to B. As v_B cannot be edge-protected as, by assumption, no edge-protected vertices exist, A cuts out B and captures the entire market. This contradicts the existence of a locational Nash equilibrium. □

It is worthwhile to point out that at equilibrium, the two facilities may *not* be adjacent. A counterexample is given in Fig. 1, where the double-digit weights and the single-digit distances are shown next to vertices and edges, respectively. Furthermore, $SSS = 6$, and $v^* = v_3$, $v^{**} = v_7$ is the unique Nash equilibrium.

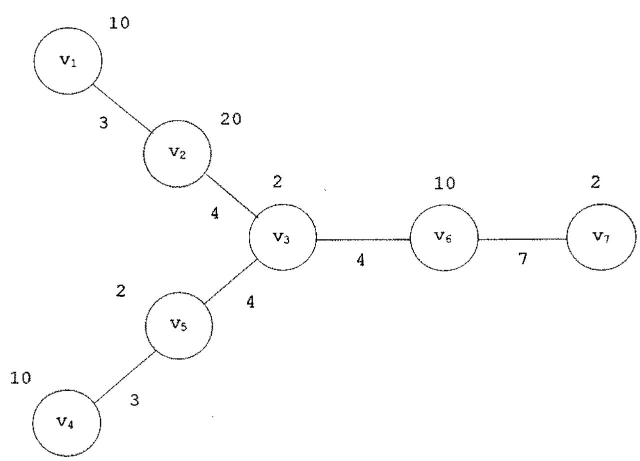

Fig. 1

There are two special cases that can be dealt with easily. For this purpose, define v_q^1 as the vertex adjacent to a median v_q that is located in the heaviest subtree spanned by v_q. In other words, $v_q^1 \in T_1^q$ and the edge $(v_q, v_q^1) \in E$. Then the two special cases are summarized in

Corollary 10. If one of the medians v_q is edge-protected, then $v^* = v_q^1$ and $v^{**} = v_q$ is a locational Nash equilibrium. Similarly, if v_q^1 is edge-protected, then $v^* = v_q$ and $v^{**} = v_q^1$ is a locational Nash equilibrium.

The proof of Corollary 10 is done by checking the conditions of Lemma 7 for the two special cases and is omitted here. Note that in both cases v^* and v^{**} are generally not interchangeable.

The problem investigated in the next two sections can now formally be stated as follows. Given two facilities A and B that are located arbitrarily at v_A, and v_B, respectively. An equilibrium exists at vertices v^* and v^{**}, respectively. The question is then: given that the decision makers of both facilities use short-term profit maximizing objectives and relocate in a sequential manner, will A and B relocate so that after a finite number of steps $v_A = v^*$ and $v_B = v^{**}$?

3 Reachability of equilibria in case of unequal prices

In order to stimulate the discussion, we will first provide a numerical example in which both competitors behave as assumed above, but a locational Nash equilibrium may *not* be reached despite its existence.

For convenience, we represent the profits or payoffs of the two competitors in a matrix. For that purpose, denote by p_{ij}^A the profit of facility A given that $v_A = v_i$ and $v_B = v_j$, and similar for p_{ij}^B. Then the matrix $\mathbf{P} = [(p_{ij}^A, p_{ij}^B)]$ is the payoff matrix of the associated bimatrix game. In our case of unequal, but fixed, prices the game is not a constant-sum game, but there is a constant-sum game with payoffs equaling market capture, which is equivalent to the revenue-maximizing original game in the sense that a Nash equilibrium in one game is also a Nash equilibrium in the other.

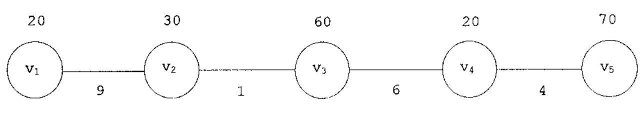

Fig. 2

With prices $p_A = 2$ and $p_B = 3$, the payoff matrix $\mathbf{P}$ for the problem in Fig. 2 is:

	v_1	v_2	v_3	v_4	v_5	min
v_1	–	40; 540	40; 540	40; 540	220; 270	40
v_2	360; 60	–	400; 0	220; 270	220; 270	220
v_3	360; 60	400; 0	–	220; 270	220; 270	220
v_4	360; 60	180; 330	180; 330	–	260; 210	180
v_5	300; 150	180; 330	180; 330	140; 390	–	140
min	60	0	0	270	210	

The row and column minima have been added to indicate the players' worst-case scenarios. Here, a saddle point exists: the location pairs (v_2, v_4) and (v_3, v_4) are indeed Nash equilibria as neither facility has an incentive to move out of its current location.

Suppose that the current locational configuration is (v_4, v_2) and it is A's turn to relocate. A will move to (v_3, v_2), where B now has the choice between (v_3, v_4) and (v_3, v_5). Assume that B chooses to relocate to (v_3, v_5). At that point, A has no choice but to move to (v_4, v_5). Here, B may either relocate to (v_4, v_2) or (v_4, v_3), and we assume that B moves to (v_4, v_2). Now the two facilities have moved in cyclical fashion, and as long as B does not change his strategy to pick a new location in case of a tie, they will continue on this cycle without ever reaching of the equilibria at (v_2, v_4) and (v_3, v_4). Inspection reveals that in this example, randomization among the best strategies will lead to an equilibrium, a property similar to absorbing states in Markov chains. The question is whether or not convergence to an equilibrium can be proved in general.

Tackling the problem as a general bimatrix game suggests that equilibria, if they exist, may not necessarily be reached. As an example, consider the following payoff matrix of a constant-sum game, where the strategies of the two players are shown as s_1^A, s_2^A, and s_3^A for player A and s_1^B, s_2^B, and s_3^B for player B.

	s_1^B	s_2^B	s_3^B
s_1^A	(½, ½)	(½, ½)	(½, ½)
s_2^A	(1, 0)	(0, 1)	(½, ½)
s_3^A	(0, 1)	(1, 0)	(½, ½)

The element (s_1^A, s_3^B) is the unique Nash equilibrium. However, suppose that the initial strategy mix is (s_3^A, s_1^B). Now facility A will change strategy to (s_2^A, s_1^B), then B changes to (s_2^A, s_2^B), A subsequently changes to (s_3^A, s_2^B) and then B changes to (s_3^A, s_1^B) and we have reached the initial strategy mix. Here, randomization in case of a tie does not help and the above example illustrates that in general bimatrix games Nash equilibria may not be reached at all, not only by virtue of the "wrong" choice of tie breaking rule.

In order to analyze the reachability problem with sequential moves in our location model, it is useful to first examine the basic moves of the two facilities. Suppose that facility B is located at some vertex v_B. Then we must distinguish between two cases.

Case 1. Vertex v_B is edge-protected. Then A cannot cut out B and will locate in T_1^B. One of A's optimal strategies is to locate at v_1^B although other choices may exist. One such example is shown in Figure 1, given $SSS < 7$ and $v_B = v_7$. Here, A may locate at v_6 or v_3, capturing $w(T) - w_7$ in both cases.

Case 2. Vertex v_B is not edge-protected. Now A can cut out B and will do so by locating at some vertex v_A so that $d_{AB} \leq SSS$. As by assumption v_B is not edge-protected, such a vertex v_A always exists. Facility A then captures the entire market $w(T)$. This leads to

Relocation Rule RR1. Given the location of a vertex v_B, facility A will always locate at a vertex v_A, so that if v_B is edge-protected, A locates so as to capture its own hinterland as well as the entire competitive region; otherwise v_A is chosen, so that $d_{AB} \leq SSS$.

Suppose now that facility A's location is temporarily fixed at some vertex v_A. As defined in Sect. 2, a circle C_A around v_A consists of all vertices that are no farther than SSS from v_A. If B were to locate at any vertex in C_A, it would be cut out and its market share would drop to zero. Clearly, if $C_A = V$, facility B must accept defeat and live with zero sales. It cannot do better, and neither can A, so that a Nash equilibrium is reached. Suppose now that there exists at least one vertex $v_i \notin C_A$. Facility B will now locate at one such vertex, so as to maximize its market share. Note that B can always do so by locating at a vertex v_B, so that all vertices on the path $[v_A, \ldots, v_B[$ are in C_A. Otherwise, B could simply move closer to A without losing any demand in its hinterland and possibly gain in the competitive region. This implies

Relocation Rule RR2. Given that facility A locates at a vertex v_A facility B will always locate just outside C_A at a vertex v_B that maximizes the demand in its own hinterland.

In order to prove reachability, we first deal with the easy case in which no edge-protected vertex exists and a locational Nash equilibrium, if it exists, must be degenerate. As per Lemma 9, there must exist a vertex v^*, such that $C_* = V$. Suppose that at present, B is located optimally with respect to A and it is now A's turn to relocate. If $v_A = v^*$, and equilibrium has been reached. Let now $v_A \neq v^*$. If B is located at some v_B, so that $M(B) = 0$, then neither facility can improve its market area and an equilibrium has been reached. Suppose now that B captures $M(B) > 0$. As by assumption no edge-protected vertices exist, A can relocate so as to cut out B. In particular, if currently $v_B \neq v^*$, then A can locate at v^*, thus reaching an equilibrium. If $v_B = v^*$, then A can locate at some $v_{\hat{A}}$ and cut out B; in the next step B moves out of v^* (if it cannot improve, we have again reached an equilibrium), so that A can then locate at v^* when it is its turn. This demonstrates the reachability of a Nash equilibrium in case no edge-protected vertex exists.

Consider now the case in which at least one edge-protected vertex exists and assume that $v_A, v_B \notin C_q$. Let A be temporarily fixed at v_A and B relocates to some vertex $v_{\hat{B}}$. We can then prove

Lemma 11. If $v_A, v_B \notin C_q$ then $v_{\hat{B}} \in p_{Aq} =]v_A, \ldots, v_q]$, such that $d_{A\hat{B}} > SSS$ and there exists no other vertex $v_k \in p_{Aq}$ with $d_{Ak} < d_{A\hat{B}}$, i.e., facility B locates on the path p_{Aq} just outside C_A.

Proof. The fact that $v_A \notin C_q$ implies $v_q \notin C_A$ and thus $v_{\hat{B}}$ must be located in A's hinterland or in the region between v_q and v_A. Now corollary 4 implies that $C_q \subseteq T_1^A$ and RR2 then indicates that B will not locate in A's hinterland. Furthermore, if there were a vertex v_k, such that $v_k \in p_{A\hat{B}}$ and $v_k \notin C_A$ then B would capture at least $S(\hat{B}) + w_k$ by locating at v_k which is no less than if B had located at $v_{\hat{B}}$. Hence we can conclude that B locates on p_{Aq} just outside C_A. □

Consider now facility A's move. From $v_{\hat{B}} \notin C_A$ follows $v_A \notin C_{\hat{B}}$. If $v_{\hat{B}}$ is edge-protected, then $C_{\hat{B}} = \{v_{\hat{B}}\}$ and A's optimal location (or at least one of its optimal locations) is at $v_{\hat{A}} = v_{\hat{B}}^1$. Then corollary 4 implies that $v_q \in T_1^{\hat{B}}$ and the distance between the median and A's new location is

$$d_{\hat{A}q} = d_{Aq} - d_{A\hat{B}} - d_{\hat{A}\hat{B}} < d_{Aq}$$

implying that facility A has moved closer to the median.

If, on the other hand, vertex $v_{\hat{B}}$ is not edge-protected, facility A will locate at some vertex $v_{\hat{A}}$ so that $d_{\hat{A}\hat{B}} \leq SSS$ and B is cut out. Then

$$d_{\hat{A}q} = d_{Aq} - d_{A\hat{B}} \pm d_{\hat{A}\hat{B}} < d_{Aq} \quad (\text{as } d_{A\hat{B}} > SSS \geq d_{\hat{A}\hat{B}})$$

and again, A is located closer to the median v_q than in the previous iteration. Defining a *relocation round* as two successive moves by the competitors, one by A and one by B, then the above discussion implies.

Lemma 12. After at most n relocation rounds, facility A or B will locate in C_q.

By virtue of Lemma 12, we can now assume that facility A or B (or both) is (are) located in C_q. Consider first the special case of an edge-protected median, i.e., $C_q = \{v_q\}$. As per Lemma 12, A or B will eventually locate at v_q. Suppose $v_A = v_q$. Then B will locate at $v_B = v_q^1$ where it is not cut out. If v_q^1 is edge-protected, the conditions of Lemma 7 are satisfied and an equilibrium has been reached. If v_q^1 is not edge-protected, then A can cut out B by relocating at some vertex $v_{\hat{A}} = v_B^k$, $k \neq 1$, so that $d_{\hat{A}B} \leq SSS$. Facility B's optimal reaction is then to relocate at $v_{\hat{B}} = v_q$, a move that protects B and leads A to locate at v_q^1. The pair of locations $(v_A, v_B) = (v_q^1, v_q)$ satisfies the conditions of Lemma 7 and thus constitutes a Nash equilibrium. On the other hand, if $v_B = v_q$ occurs instead of $v_A = v_q$, then facility A responds by locating at v_q^1 and again, an equilibrium is reached.

Another special case occurs if v_q is not edge-protected, but v_q^1 is. Clearly, $v_q^1 \notin C_q$. Again invoke Lemma 12. If $v_A \in C_q$ then facility B's best response is to locate at v_q^1 as there it captures $w(T_1^q)$ whereas elsewhere at a distance of at least SSS from A, it can never get more than $w(T_k^q)$ for some k. One of facility A's optimal responses is, in turn, to locate at v_q. According to Lemma 7, this pair of locations is a Nash equilibrium. On the other hand, if $v_B \in C_q$, and as, by assumption, v_q is not edge-protected and $|C_q| \geq 2$, A can locate at one of the vertices in C_q and cut out B. Wherever A locates in C_q, B's best response is to locate at $v_B = v_q^1$ which, again by virtue of Lemma 7, then constitutes a pair of Nash equilibria. This proves

Corollary 13. If the median v_q or its adjacent vertex v_q^1 is edge-protected, then a Nash equilibrium will be reached in a sequential relocation process with short-term profit maximization objective.

In the following we assume that neither v_q nor v_q^1 is edge-protected. According to Lemma 12, one of the two facilities will locate inside C_q after no more than n relocation moves. We can now distinguish between two cases.

Case 1. $v_B \in C_q$, i.e., facility B is the first to locate in C_q. As by assumption v_q is not edge-protected, $|C_q| \geq 2$ and C_q cannot include any edge-protected vertices, so that v_B is currently not edge-protected. This leaves two subcases for consideration.

Case 1a. $v_B = v_q$. Then A in its next move can cut out B by locating at v^* as $v^* \in C_q$ by virtue of Lemma 8a. Now facility B's best response is to locate at v^{**} which then satisfies the conditions of Lemma 7b and a locational Nash equilibrium has been reached.

Case 1b. $v_q \neq v_B \in C_q$. One of facility A's optimal moves is to cut out B by locating at v_q. Now, $v_A = v_q$ and it is B's turn to relocate at some vertex $v_{\hat{B}}$. Clearly, B will not locate in C_q as it would be cut out (and, by assumption, at least one edge-protected vertex exists, so that $C_q \subsetneqq V$).

(i) $v_{\hat{B}} \in C_*$. In this case, A can cut out B by locating at $v_{\hat{A}} = v^*$, and in the next step, B will respond by relocating to v^{**} as per RR2. This is a Nash equilibrium.
(ii) $v_{\hat{B}} \notin C_*$. By virtue of Lemmas 8b, 8c, and 7b, v^{**} is B's best option to locate. Subsequently, A will move to v^* and again, an equilibrium has been reached.

Case 2. $v_A \in C_q$, i.e., facility A is the first to locate at some vertex $v_A \in C_q$. In order to obtain a positive market share, facility B must move to a vertex $v_{\hat{B}} \in V \backslash C_q$. This is the same situation as that discussed in subcase (ii) under 1b above which was already shown to lead to a Nash equilibrium.

The above discussion implies

Theorem 14. Given two facilities that locate at different vertices of a tree and the facilities charge fixed, but different mill prices. Then sequential relocation on the basis of short-term profit maximization, with rules as specified in cases 1, 2, and 3 above, will lead to Nash equilibrium locations.

4 Conclusions

In this paper we have examined the problem of convergence of duopolists towards their respective equilibrium locations, given that decision makers at both facilities apply a short-term maximization objective. It was shown that, given some additional rules in case of ties, the duopolists will reach equilibrium locations in the process.

The results presented in this paper open up a number of new avenues. One such possibility is the examination of "forbidden regions", a subject that was first studied by Katz and Cooper (1981). In our model, such forbidden zones could be vertices at which the competitors may not locate. Restrictions of this nature could be due to local zoning laws. It is apparent that the introduction of forbidden zones in a problem may dramatically change the results. Take, for example, the graph in Fig. 1 and assume that location at vertex v_7 is prohibited. In this case, given $SSS = 6$, the only edge-protected vertex is now no longer available for location and a locational Nash equilibrium no longer exists. (In contrast, if in the original Hotelling model with variable prices and location and a linear market that stretches from 0 to L, we were to introduce a forbidden zone between ¼L and ¾L, then the model would indeed have a subgame perfect Nash equilibrium.) Results concerning competitive location models and forbidden zones have not been reported in the literature.

A number of other research directions is also possible. For instance, one may examine scenarios in which one or both of the planners at the facilities make their decisions with foresight. Another possibility is to investigate models in which customers do not patronize the facility with the lowest delivered price, but use other (deterministic) utility functions, such as proportional models, similar to that used by Bauer et al. (1993).

Acknowledgements. This research was in part supported by grants from the Natural Sciences and Engineering Research Council of Canada under grant numbers OGP0009160 and OGP0121689. This support is gratefully acknowledged. We would also like to thank two anonymous referees whose careful reading of the paper resulted in a number of clarifications.

Appendix

Proof of Lemma 7. (i) We first prove necessity of each of the three conditions by contradiction. Suppose that v^{**} is not edge-protected. Then either (i) $S(B) = 0$ and B could increase its sales by locating at any edge-protected vertex (if A were to be located at the only edge-protected vertex, then B would capture its own vertex and $S(B) > 0$). Alternatively, (ii) $S(B) > 0$ in which case A could relocate so as to cut out B; hence the current arrangement is not an equilibrium.

Suppose now that v^{**} does not satisfy condition (b). Then B could relocate to the vertex with the heaviest hinterland outside C_* and capture a larger share of the market.

Finally, let $M(A) \neq T_1^{**}$. Then either (i) $v^* \in T_1^{**}$ but there exists in the competitive region a vertex $v_i \in T_1^{**} \cap M(B)$, in which case A could move towards B and capture v_i as well, thus increasing its market area, or (ii) $v^* \notin T_1^{**}$, in which case A either cuts out B, which is not possible as v^{**} is edge-protected, as per condition (a), or A could improve by locating, so that $M(A) = T_1^{**}$. This demonstrates that each of the three conditions is necessary for (v^*, v^{**}) to be an equilibrium.

(ii) We now prove sufficiency of conditions (a)–(c). Conditions (a) and (c) ensure that facility A has no incentive to move, as per (a) facility A cannot cut out B, and by way of (c) A captures all of the market except B's hinterland. Condition (c) also implies that A's market share is maximal, given B. Furthermore, given that A locates at $v_A = v^*$, B's market share is maximal by way of condition (b). □

Proof of Lemma 8. (a) We first prove that $v_q \in C_*$. Suppose not. Then by Lemma 7b, v^{**} must be located just outside C_* on the path $p_{*q} =]v^*, \ldots, v_q]$. By virtue of Lemma 7a, v^{**} is edge-protected, but by the above assumption v_q is not, hence $v^{**} \neq v_q$, so that $p_{*q} = [v^*, \ldots, v^{**}, \ldots, v_q]$. Thus facility A located at v^* captures strictly less than $\frac{1}{2}w(T)$ (*strictly* less, as A's market share cannot exceed $w(T_1^q) \leq \frac{1}{2}w(T)$ by Lemma 3, but as it does not capture $w(v^{**}) > 0$, the inequality is strict). On the other hand, Lemma 3 also implies that if facility A were to locate at v_q it would capture at least $\frac{1}{2}w(T)$ which contradicts the assumption that $v_q \notin C_*$. Hence $v_q \in C_*$ which, in turn, implies that $v^* \in C_q$.

(b) By definition, the circle C_q cannot include edge-protected vertices. However, Lemma 7a states that v^{**} is edge-protected; which is a contradiction.

(c) Vertex v^{**} is edge-protected as per Lemma 7a, thus it cannot be in C_*. □

References

Anderson S (1987) Spatial competition and price leadership. International Journal of Industrial Organization 5:369–398

Bauer A, Domschke W, Pesch E (1993) Competitive location on a network. European Journal of Operational Research 66:372–391

Eiselt HA (1992) Hotelling's duopoly on a tree. Annals of Operations Research 40:195–207

Eiselt HA, Laporte G (1991) Locational equilibrium of two facilities on a tree. Operations Research/Recherche opérationnelle 25:5–18

Eiselt HA, Laporte G, Thisse J-F (1993) Competitive location models: a framework and bibliography. Transportation Science 27:44–54

Friesz TL, Miller TC, Tobin RL (1988) Competitive network location models: a survey. Papers of the Regional Science Association 65:47–57

Goldman AJ (1971) Optimal center location in simple networks. Transportation Science 5:212–221

Hakimi SL (1983) On locating new facilities in a competitive environment. European Journal of Operational Research 12:29–35

Hakimi SL (1990) Locations with spatial interactions: competitive locations and games. In: Mirchandani PB, Francis RL (eds) Discrete location theory. Wiley Interscience, New York, pp 439–478

Hotelling H (1929) Stability in competition. Economic Journal 39:41–57

Katz IN, Cooper L (1981) Facility location in the presence of forbidden regions. 1. Formulation and the case of Euclidean distances with one forbidden circle. European Journal of Operational Research 6:166–173

Labbé M, Hakimi SL (1991) Market and locational equilibrium for two competitors. Operations Research 39:749–756

ReVelle C (1986) The maximum capture or 'sphere of influence' location problem: Hotelling revisited on a network. Journal of Regional Science 26:343–358

Stackelberg H von (1943) Grundlagen der theoretischen Volkswirtschaftslehre. Translated by: Peacock AT, The Theory of the Market Economy (1952). Hodge, London Edinburgh Glasgow

Tirole J (1995) The theory of industrial organization. The MIT Press, Cambridge, MA

Tovey CA (1993) Dynamical convergence of majority rules to ε-cores in Euclidean spaces. Working paper, School of Industrial and Systems Engineering, Georgia Tech, Atlanta, GA

Wendell RE, McKelvey RD (1981) New perspectives in competitive location theory. European Journal of Operational Research 6:174–182