

R-1538-PR

July 1974

A Note on the Lemke-Howson Algorithm

Lloyd S. Shapley

A Report prepared for
UNITED STATES AIR FORCE PROJECT RAND

Rand
SANTA MONICA, CA. 90406

The research described in this Report was sponsored by the United States Air Force under Contract No. F44620-73-C-0011 — Monitored by the Directorate of Operational Requirements and Development Plans, Deputy Chief of Staff, Research and Development, Hq USAF. Reports of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

R-1538-PR

July 1974

A Note on the Lemke-Howson Algorithm

Lloyd S. Shapley

A Report prepared for
UNITED STATES AIR FORCE PROJECT RAND



PREFACE

This Report relates to the computation of noncooperative equilibrium strategies in certain nonzero-sum two-person games. Such games can arise both in models of tactical and strategic operations and in models of such things as force posture and arms-limitation negotiations. The Report contains a graphical description of the basic "path-following" algorithm due to C. E. Lemke and J. T. Howson, Jr., and presents some theorems concerning the algebraic and topological properties both of the algorithm and of the equilibrium points themselves.

This work was done as part of supporting research for USAF Project RAND.

SUMMARY

The Lemke-Howson algorithm for bimatrix games provides both an elementary proof of the existence of equilibrium points and an efficient computational method for finding at least one equilibrium point. The first half of this Report presents a geometrical view of the algorithm that makes its operation especially easy to visualize. Several illustrations are given, including Wilson's example of "inaccessible" equilibrium points. The second half presents an orientation theory for the equilibrium points of (nondegenerate) bimatrix games and the Lemke-Howson paths that interconnect them; in particular, it is shown that there is always one more "negative" than "positive" equilibrium point.

CONTENTS

PREFACE	iii
SUMMARY	v
Section	
1. INTRODUCTION	1
2. ALMOST-COMpletely-LABELED PATHS	2
3. EXAMPLES	10
4. INDEX THEORY	16
REFERENCES	26

1. INTRODUCTION

The first half of this paper is frankly expository; it contains little of substance that is not covered in Lemke and Howson's original work [2]. But rather than skim over this material as rapidly as possible, we have taken this occasion to set down the details of a geometric labeling system for bimatrix games, which we have found very effective for explaining, at least to "lay" audiences, the workings of the Lemke-Howson method.*

The second half presents an index theory for bimatrix games, in which we show that each equilibrium point (in the nondegenerate case) has an intrinsic orientation, and each almost-complementary path or loop an intrinsic sense of direction. The possibility of such a theory will hardly surprise those who are familiar with the "strong" form of Sperner's lemma [1, p. 133]; it is analogous to (and perhaps reducible to) Ky Fan's theory for abstract orientable pseudo-manifolds given in his 1967 paper [2]. In any case, our concern here is with the concrete realization of the theory, as we shall give explicit definitions for the various indices in terms of the signs of suitably constructed determinants derived from the payoff matrices.

*As presented, e.g., at the Second World Congress of the Econometric Society in Cambridge, England, September 1970.

The author would like to acknowledge helpful conversations and correspondence with B. C. Eaves, H. W. Kuhn, C. E. Lemke, H. E. Scarf, and A. W. Tucker, as well as the invitation from the latter that prompted the writing of this paper.

2. ALMOST-COMpletely-LABELED PATHS

A bimatrix game (A, B) consists of two m -by- n matrices $A = (a_{ij}: i \in I, j \in J)$ and $B = (b_{ij}: i \in I, j \in J)$, representing the payoffs to two players using pure strategies i and j respectively. It will be convenient to have the sets I and J disjoint, so we define $I = \{1, \dots, m\}$, $J = \{m + 1, \dots, m + n\}$, $K = I \cup J$. Mixed strategies are represented by vectors $s = (s_1, \dots, s_m) \in S$ and $t = (t_{m+1}, \dots, t_{m+n}) \in T$, where

$$\begin{cases} S = \{s \geq 0; \sum_I s_i = 1\} \\ T = \{t \geq 0; \sum_J t_j = 1\}. \end{cases}$$

The corresponding payoffs are $\sum_I \sum_J a_{ij} s_i t_j$ and $\sum_I \sum_J b_{ij} s_i t_j$. Geometrically, the sets S and T are simplices, of dimension $m - 1$ and $n - 1$, respectively. Define also

$$\begin{cases} \tilde{S} = S \cup \{s \geq 0: \sum_I s_i \leq 1 \text{ and } \prod_I s_i = 0\} \\ \tilde{T} = T \cup \{t \geq 0: \sum_J t_j \leq 1 \text{ and } \prod_J t_j = 0\}. \end{cases}$$

These extended domains are the boundaries of the simplices of the next higher dimension, "cut off" from the positive orthant by S and T respectively.

An equilibrium point of (A, B) is a pair $(s^*, t^*) \in S \times T$ such that

$$\begin{cases} \sum_I \sum_J a_{ij} s_i^* t_j^* = \max_{s \in S} \sum_I \sum_J a_{ij} s_i t_j^*, \\ \sum_I \sum_J b_{ij} s_i^* t_j^* = \max_{t \in T} \sum_I \sum_J b_{ij} s_i^* t_j. \end{cases}$$

An equivalent condition is

$$\begin{cases} \sum_J a_{i^*j} t_j^* = \max_{i \in I} \sum_J a_{ij} t_j^*, & \text{all } i^* \in I \text{ with } s_{i^*}^* > 0, \\ \sum_I b_{ij^*} s_i^* = \max_{j \in J} \sum_I b_{ij} s_i^*, & \text{all } j^* \in J \text{ with } t_{j^*}^* > 0. \end{cases}$$

We now define certain closed convex, polyhedral regions S^k in \tilde{S} , as follows:

$$\begin{cases} S^i = \{s \in \tilde{S} : s_i = 0\}, & \text{for } i \in I \\ S^j = \{s \in S : \sum_I b_{ij} s_i = \max_{\ell \in J} \sum_I b_{i\ell} s_i\}, & \text{for } j \in J. \end{cases}$$

The sets S^i , $i \in I$, cover all of $\tilde{S} - S$, as well as the relative boundary of S . The sets S^j , $j \in J$ (some of which may be empty) consist of those mixed strategies for the first player against which the pure strategy "j" is a best response by the second player. Since there always is at least one best response, they cover all of S . Hence the sets S^k , $k \in K$ cover all of \tilde{S} . Define the label of $s \in \tilde{S}$ to be the nonempty set

$$L'(s) = \{k: s \in S^k\}.$$

In exactly similar fashion, define regions T^k in \tilde{T} by

$$\begin{cases} T^i = \{t \in T: \sum_J a_{ij}t_j = \max_{\ell \in I} \sum_J a_{\ell j}t_j\}, & \text{for } i \in I, \\ T^j = \{t \in \tilde{T} : t_j = 0\}, & \text{for } j \in J, \end{cases}$$

and, for $t \in \tilde{T}$, define the label

$$L''(t) = \{k: t \in T^k\}.$$

Finally, let the label of the pair $(s, t) \in \tilde{S} \times \tilde{T}$ be

$$L(s, t) = L'(s) \cup L''(t),$$

We shall say that (s, t) is completely labeled if $L(s, t) = K$, and almost completely labeled if $L(s, t) = K - \{k\}$ for some $k \in K$.

THEOREM 1. If $(s, t) \in S \times T$, then (s, t) is an equilibrium point of (A, B) if and only if (s, t) is completely labeled.

This follows almost immediately from the definitions.

We now make a "nondegeneracy" assumption.

The game (A, B) is such that (1) each non-empty region S^k is $(m - 1)$ -dimensional; (2) the intersection of any two of the S^k is at most $(m - 2)$ -dimensional; (3) no point of \tilde{S} belongs to more than m of the S^k ; and (4) the analogous conditions to (1), (2), (3) hold for the regions T^k in \tilde{T} .

As the failure of any of these conditions would entail a special numerical relationship among the a_{ij} or b_{ij} ; we see that "almost all" bimatrix games are nondegenerate, in this sense.

Armed with this assumption, we can describe a graph \mathcal{S} in \tilde{S} , whose point set consists of all points in \tilde{S} that belong to at least $m - 1$ of the regions S^k . The points that belong to exactly m regions are the nodes of \mathcal{S} , while those that belong to exactly $m - 1$ regions make up the edges. In a given edge all points have the same label, which we shall consider the label of that edge. The nodes and edges of \mathcal{S} exhibit the following incidence relations: (1) each edge touches exactly two nodes (i.e., its endpoints), and (2) each node touches exactly m edges, each one omitting from its label a different member of the label of the node. We call two nodes adjacent if they are at opposite ends of the same edge, or, equivalently (since the regions are convex), if their labels differ in exactly one element.

An exactly analogous graph \tilde{S} can be described in \tilde{T} . The construction of these graphs is illustrated below for \mathfrak{S} , with $m = n = 3$, where we have taken B to be the identity matrix I_3 . In the equivalent planar diagram at the right, the exceptional node, 0 , is at the top and region ① is unbounded. It will be seen that the effect of our rather unusual extension of the mixed-strategy simplex from S to \tilde{S} has been to "close out" the graph, making it "regular of degree m " by adding one node and m edges.




Fig. 1

We now turn our attention to node pairs.* Let \mathcal{N} denote the set of all $(s, t) \in \tilde{\mathcal{S}} \times \tilde{\mathcal{T}}$ where s is a node of \mathcal{S} and t a node of \mathcal{T} . Define

$$\vartheta = \{(s, t) \in \mathcal{N} : L(s, t) = K\}$$

and, for each $k \in K$,

$$\vartheta^k = \{(s, t) \in \mathcal{N} : L(s, t) \supseteq K - \{k\}\}.$$

These are the node pairs that are completely or almost completely labeled; note that $\vartheta^k \supseteq \vartheta$, and that $k \neq \ell$ implies $\vartheta^k \cap \vartheta^\ell = \vartheta$. From Theorem 1 and our nondegeneracy assumption, it is easy to see that the members of ϑ are just the

*Implicitly, we are forming a kind of "product" graph out of \mathcal{S} and \mathcal{T} , in which the node-node pairs from \mathcal{S} , \mathcal{T} are the nodes and the node-edge and edge-node pairs are the edges.

The "graph" approach to path-following algorithms is more or less dual to the "pseudomanifold" (or regular simplicial complex) approach. Thus, the nodes and edges of \mathcal{S} and the $(m-1)$ -dimensional regions S^k correspond to the simplices of dimension $m-1$, $m-2$, and 0 , respectively, of the simplicial complex that is dual to the polyhedral complex in $\tilde{\mathcal{S}}$ of which \mathcal{S} is the regular, 1-dimensional skeleton. Kuhn [3] suggests viewing the L-H method as operating on a product of disjoint pseudomanifolds, which is itself trivially a pseudomanifold. However, the success of the method does not depend on \mathcal{S} actually being the skeleton of the dual of a pseudomanifold, since no conditions need be satisfied by the cells of intermediate dimension between 1 and $m-1$, nor need the regions S^k be convex or even connected. In an unpublished 1970 note, the present author proposed the use of combinatorially defined objects called "regular quasi-skeletons," as a generalized setting for the L-H method. But it is not clear that RQS's not arising from orientable pseudomanifolds would support any kind of orientation theory.

equilibrium points of (A, B) together with the node pair (0, 0).

Let us call two members of \mathcal{N} adjacent if their \mathcal{S} -components are the same and their \mathcal{J} -components adjacent, or their \mathcal{J} -components are the same and their \mathcal{S} -components adjacent. A subset \mathcal{R} of \mathcal{N} is said to be connected if a chain of members of \mathcal{R} can be found, each one adjacent to the next, joining any two given members of \mathcal{R} . A component of a subset \mathcal{R} of \mathcal{N} is a maximal connected subset of \mathcal{R} ; every subset of \mathcal{N} is the disjoint union of its components. A nonempty connected set \mathcal{R} is called a loop if every member of \mathcal{R} is adjacent to exactly two other members of \mathcal{R} , and a path if every member is adjacent to exactly two others except for two members (i.e., the endpoints) that are adjacent to just one each. A loop has at least three members, a path at least two.*

LEMMA 1. Let $k \in K$ be fixed. (a) Each member of ϑ is adjacent to exactly one member of ϑ^k . (b) Each member of $\vartheta^k - \vartheta$ is adjacent to exactly two members of ϑ^k .

To prove (a), take a completely labeled pair $(s, t) \in \vartheta$ and select the node r , $r = s$ or t , whose label includes k . Following the edge out of r whose label omits k yields a

*In our application, however, we shall find that paths of length two and loops of length three (or any odd length) cannot occur.

member of ϑ^k adjacent to (s, t) ; this is uniquely determined, since any other edge out of s or t would take us out of ϑ^k . To prove (b), take (s, t) in $\vartheta^k - \vartheta$ and note that there must be exactly one duplication in the label, say $\{h\} = L'(s) \cap L''(t)$. Following the edge out of s that omits h yields one adjacent member of ϑ^k , and following the edge out of t that omits h yields another; they are clearly distinct, and there are no others. Q.E.D.

From Lemma 1 it follows that each component of ϑ^k is either a path or a loop, and that the path endpoints are precisely the members of ϑ . Hence ϑ has an even number of members, and since $(0, 0)$ is not an equilibrium point, we have



THEOREM 2. (Lemke-Howson). Every non-degenerate bimatrix game has an odd number of equilibrium points (and hence at least one).

3. EXAMPLES

To exploit the "constructive" aspect of this existence proof, we need to be able to place ourselves on a path component of some ϑ^k . This is actually very easy, since the exceptional node pair $(0, 0)$ belongs to ϑ , and hence to every ϑ^k . Moreover, it is never on a loop. Thus, assuming we know how to read labels and follow edges in \mathfrak{S} and \mathfrak{J} , all we have to do to find an equilibrium point is select a value for k , start at $(0, 0)$, and follow the path to its end.

In working through the examples, the reader may enjoy placing two small checkers or coins on the page, one on \mathfrak{S} and one on \mathfrak{J} , and then making alternate moves. The "state of the system" at any given time is given by the positions of the two coins, plus an indication of which one moved last.

Let us try to explore ϑ^1 in Fig. 2. Starting at the exceptional node-pair (coins on "0" and "o"), we find that we must first make a move in \mathfrak{S} , sliding the coin along the edge that leads away from region ① (the unbounded, "outside" region as we have drawn it). We thereby arrive at "A", only to find that ④ is now doubly represented, in the label of the node pair (A, o) . Moving away from ④ in \mathfrak{J} takes us to "a", where we pick up contact with ③. We therefore move in \mathfrak{S} from "A" to "D", picking up ⑥; then in \mathfrak{J} from "a" to "x", picking up ④ again; and finally, in \mathfrak{S} , from "D" to "X", picking up ① for a complete label. So



Key:

p^1 OoAaDxX YyCzZ
 p^2 OoYy XxZz
 p^3 OoXx YyZz
 p^4 oOaXx yYcCdZz
 p^5 oOyY xXzZ ...bAcBdDb...
 p^6 oOxX yYzZ

Payoffs:

2	2	0	3	0	2
0	3	0	0	3	2
3	0	1	0	0	1

Fig. 2

the pair (X, x) is an equilibrium point, in which each player uses just his third pure strategy.

The φ^2 path leading out of $(0, o)$ is shorter, containing just the two other pairs (Y, o) and (Y, y) . The latter is another equilibrium point of the game, with each player using just his second pure strategy. This gives us the opportunity to start another φ^1 path, which will necessarily lead us to a third equilibrium point. In fact, moving away from ① at "Y" takes us to "C", duplicating ⑥ in the label; moving away from ⑥ at "y" takes us to "z", duplicating ③; and finally, moving away from ③ at "C" takes us to the equilibrium point (Z, z) . This equilibrium point involves a mixture of the second and third strategies for each player. As it happens, there are no other equilibrium points.*

The "Key" in Fig. 2 lists all of the φ^k 's, each "word" representing a different component, with the completely-labeled or almost completely-labeled node pairs indicated by adjacent letters. Especially interesting is the third component of φ^5 , which is a loop of length 6. Note that there are no paths from $(0, o)$ to (Z, z) or from (X, x) to (Y, y) ; the reason, as we shall see in the next

*The maximum number of equilibrium points for a non-degenerate 3×3 bimatrix game is seven; this may be attained by taking both A and B to be the identity matrix. We do not know if the maximum is similarly achieved in larger games; the question is related to the question of determining theoretical upper bounds for Lemke-Howson path lengths.


section, is that both the former have index +1 and both the latter -1, while according to Theorem 3 the endpoints of a path always have indices of opposite signs.*

The loop in θ^5 in Fig. 2 could not have been discovered by simple path-following. (Of course, in an example of this size, an exhaustive search is easy.) Figure 3 illustrates that inaccessible paths can also occur. Indeed, an inspection of the "Key" reveals six ways to get from (0, o) to (X, x) and six ways to get from (Y, y) to (Z, z), but no interconnection.** This shows that, in general, even the most thoroughgoing exploration of the path network radiating from the starting point (0, o) may not find all solutions, or even tell us whether there are any solutions un-
found. This is a serious practical drawback to the Lemke-Howson method, which is in other respects quite efficient computationally. The multiplicity of solutions is admittedly a weakness of the equilibrium-point solution concept, but there does not appear to be any good reason heuristically to prefer the solutions that are accessible to path-following, or to reject those that are not. An efficient method for finding all equilibrium points in bimatrix games would be an important contribution.

It is rather surprising that the 3×3 framework should have room for so much interesting behavior. But the apparent

*For an explanation of the arrows in Figs. 2 and 3 see the remarks accompanying Theorem 4.

**The example is due to Robert Wilson (private correspondence, 1970).



Key:

- | | | |
|-------|---------|---------|
| p^1 | OoAxX | yYbZz |
| p^2 | OoCaDxX | yYcBdZz |
| p^3 | OoXx | YyZz |
| p^4 | oOaXx | YyBzZ |
| p^5 | oOcAdXx | YyCbDzZ |
| p^6 | oOxX | yYzZ |

Payoffs:

0	3	0	0	2	3
2	2	0	3	2	0
3	0	1	0	0	1

Fig. 3

simplicity of the setting is deceptive; there are 64 node pairs in \mathcal{N} in any nondegenerate 3×3 game if none of the regions S^k, T^k are empty. In our two examples, about half of the 64 are actually used in tracing the almost-completely-labeled paths and loops; in larger games the proportion would presumably be much smaller.

4. INDEX THEORY

Given $(s, t) \in \tilde{S} \times \tilde{T}$, define the index matrix of (s, t) to be the $(m + n)$ -by- $(m + n)$ matrix $C = (c_{k\ell})$ with entries

$$c_{k\ell} = \begin{cases} 1 & \text{if } k = \ell \in I \cap L'(s) \\ b_{k\ell} & \text{if } k \in I \text{ and } \ell \in J \cap L'(s) \\ a_{\ell k} & \text{if } k \in J \text{ and } \ell \in I \cap L''(t) \\ 1 & \text{if } k = \ell \in J \cap L''(t) \\ 0 & \text{otherwise.} \end{cases}$$

This is illustrated below for $m = 3$, $n = 2$, $L'(s) = \{1, 3, 5\}$ (dots), and $L''(t) = \{3, 4\}$ (crosses). In this case, with only the second column all zero, we have $(s, t) \in \emptyset^2$.

•	•	•		
1	0	0	0	b_{15}
0	0	0	0	b_{25}
0	0	1	0	b_{35}
0	0	a_{34}	1	0
0	0	a_{35}	0	0
			x	x

Assume for the moment that all a_{ij} , b_{ij} are positive.* The index of (s, t) is then defined to be the sign of the determinant of the index matrix of (s, t) , that is, a number 1, -1, or 0. We note first that the index does not depend on the order in which the players or their pure strategies are numbered. It is obvious that any pair (s, t) not in θ will have index 0. Moreover, under our nondegeneracy assumption, no element of θ will have index 0. Our main object will be to prove that the indices of the endpoints of any path in any θ^k have opposite signs.

First we shall prove a lemma that "explains" the positivity assumption.

LEMMA 2. Let (A, B) be a nondegenerate bi-
matrix game with $A > 0, B > 0$. Form a new game
by adding nonnegative constants x and y to all
the elements of A and B respectively. Then $L',$
 $L'', \theta, \mathfrak{S}, \mathfrak{T}$, and the $\theta^k, k \in K$ are all un-
changed, as are the indices of all $(s, t) \in \tilde{\mathfrak{S}} \times \tilde{\mathfrak{T}}$

Proof. That the labels, graphs, and zero indices are unchanged is immediate from the definitions. With nondegeneracy, the determinant of the index matrix for $(s, t) \in \theta$ reduces, because of the blocks of zeroes, to

$$\pm(\det \bar{A})(\det \bar{B}),$$

*See the remark immediately following the proof of Lemma 2.

where \bar{A} and \bar{B} are the square submatrices associated with the positive parts of s and t and where the choice in the "+" symbol depends only on the labels of s and t . It will suffice to show that adding a positive constant does not change the sign of (say) $\det \bar{B}$. But \bar{B} has the property that left-multiplication by \bar{s} yields a constant, i.e.,

$$\bar{s}\bar{B} = (v, v, \dots, v)$$

for some $v > 0$, where we have written \bar{s} for the subvector made up of the positive components of s . Define $f(x)$ to be the determinant of the matrix \bar{B} with each entry increased by x . Then $f(-v) = 0$. Since f is linear in x ,* it follows that the sign of $f(x)$ is constant for $x \geq 0$. This completes the proof of Lemma 2.

In view of this lemma, it makes sense to define the indices in a nonpositive nondegenerate bimatrix game by adding constants to convert it to an equivalent positive game.

Now consider the m -by- $(m + n)$ matrix

$$[-I_m, B] = \begin{array}{|cccc|cccc} \hline -1 & 0 & \dots & 0 & b_{1,m+1} & \dots & & b_{1,m+n} \\ 0 & -1 & & 0 & b_{2,m+1} & & & \\ \dots & & \dots & & \dots & & & \\ 0 & & \dots & -1 & b_{m,m+1} & \dots & & b_{m,m+n} \\ \hline \end{array}$$

*To see this, subtract one column of the matrix from all the other columns.

Given an arbitrary node s of \mathfrak{S} , we can use the label $L'(s)$ to select columns of $[-I_m, B]$ to form a square matrix $B(s)$. Except for the minus signs, $B(s)$ will consist of the nonzero columns of the upper part of the index matrix of (s, t) , for any t .

LEMMA 3. Let s' be adjacent to s in \mathfrak{S} , and let the columns of $B(s)$ and $B(s')$ be so ordered that the two matrices are identical except in one column. Let $B > 0$. Then the determinants of $B(s)$ and $B(s')$ are opposite in sign.

This sort of lemma is familiar in linear programming,* but we do not have here quite the standard set-up, so we shall prove it afresh.

Let s be a node in \mathfrak{S} , and let $v = \max_J \sum_I b_{ij} s_i$; the latter is positive because $B > 0$. Define the vector p by the multiplication

$$(s_1, \dots, s_m)[-I_m, B] = (p_1, \dots, p_{m+n}),$$

and let $\delta_{Jk} = 0$ for $k \in I$, $\delta_{Jk} = 1$ for $k \in J$. Then we have

*When one pivots only on negative entries, the determinant of the "current basis matrix," which is the product of pivots to date, changes sign at each step.

$$p_k \left\{ \begin{array}{c} \bar{=} \\ < \end{array} \right\} \delta_{Jk} v, \quad \text{for } k \left\{ \begin{array}{c} e \\ \neq \end{array} \right\} L'(s).$$

Similarly, defining v' and p' analogously for the adjacent node s' , we have

$$p'_k \left\{ \begin{array}{c} \bar{=} \\ < \end{array} \right\} \delta_{Jk} v', \quad \text{for } k \left\{ \begin{array}{c} e \\ \neq \end{array} \right\} L'(s').$$

Let k^* and ℓ^* be the columns of $[-I_m, B]$ that make the difference between $B(s)$ and $B(s')$. That is, let $\{k^*\} = L'(s) - L'(s')$ and $\{\ell^*\} = L'(s') - L'(s)$. Then there are positive numbers α, α' such that

$$p_{\ell^*} = \delta_{J\ell^*} v - \alpha \quad \text{and} \quad p'_{k^*} = \delta_{Jk^*} v' - \alpha'.$$

For $0 \leq \lambda \leq 1$ define $B_\lambda = (1 - \lambda)B(s) + \lambda B(s')$. Multiplying $v's - vs'$ into B_λ , we obtain all zeroes except for the inner product of $v's - vs'$ with the replacement column; there we obtain

$$\begin{aligned} & v'[(1 - \lambda)p_{k^*} + \lambda p_{\ell^*}] - v[(1 - \lambda)p'_{k^*} + \lambda p'_{\ell^*}] \\ &= v'[(1 - \lambda)\delta_{Jk^*} v + \lambda(\delta_{J\ell^*} v - \alpha)] \\ & \quad - v[(1 - \lambda)(\delta_{Jk^*} v' - \alpha') + \lambda\delta_{J\ell^*} v'] \\ &= -v'\lambda\alpha + v(1 - \lambda)\alpha' \\ &= v\alpha' - \lambda(v\alpha' + v'\alpha). \end{aligned}$$

Thus, $v's - vs'$ annihilates the matrix B_λ when $\lambda = v\alpha' / (v\alpha' + v'\alpha)$. This number lies between 0 and 1 since v, v', α, α' all are positive. So the determinant of B_λ , which is linear in λ and not identically 0, has opposite signs at $\lambda = 0$ and $\lambda = 1$. This completes the proof of Lemma 3.

Now let (s, t) be an almost completely labeled node pair, say $(s, t) \in \vartheta^k - \vartheta$, and let h be the duplication in its label: $\{h\} = L'(s) \cap L''(t)$. Assuming $A > 0, B > 0$, define the S-index of (s, t) to be the sign of the determinant of the "S-index matrix" of (s, t) , obtained from the index matrix of (s, t) by a partial column transposition, namely, interchange the "upper" (or "I") portion of the k -th column (which is all zeroes) with the upper portion of the h -th column (which is the h -th unit vector if $h \in I$ or the h -th column of B if $h \in J$). The J-index of (s, t) is defined analogously, by making a transposition in the lower or "J" portion of the index matrix. By non-degeneracy and positivity, both of these indices are nonzero.

LEMMA 4. The S-index and the J-index for a given $(s, t) \in \vartheta^k - \vartheta$ have opposite signs.

Proof. The defining matrices differ only by a single column transposition. Q.E.D.

For completely labeled node pairs it will be convenient to define both the S-index and the J-index to be equal to the index, as previously defined.

LEMMA 5. Let (s, t) and (s', t) be adjacent node pairs in ϑ^k . Then their \mathfrak{S} -indices have opposite signs.

Proof. Regardless of whether one of (s, t) , (s', t) is in ϑ or not, their \mathfrak{S} -index-defining matrices differ only in the "upper" portions of their k -th columns, which reflect the replacement step between s and s' . Thus, the determinants of these matrices have the form $\pm(\det B(s))(\det \bar{A})$ and $\pm(\det B(s'))(\det \bar{A})$, with the same choice in the " \pm " symbol and with $B(s)$ and $B(s')$ related as in Lemma 3. Hence their signs are opposite. Q.E.D.

We can now define a "sense of direction" on the paths and loops of ϑ^k . Arbitrarily, let a "forward" step at (s, t) be a move to an adjacent member of ϑ^k that changes s if the \mathfrak{S} -index of (s, t) is positive, or that changes t if the \mathfrak{J} -index of (s, t) is positive. Similarly, a "backward" step changes s if the \mathfrak{S} -index is negative and t if the \mathfrak{J} -index is negative. Hence every move to an adjacent member of ϑ^k is either a forward or a backward step, and by Lemma 4 we have, at each almost-completely-labeled node pair, the choice between a forward and a backward step. Moreover, by Lemma 5, the reversal of any forward step is a backward step, and vice versa. Hence we have

THEOREM 3. Moving forward along any path in any ϑ^k leads to an equilibrium point with

index -1. Moving backward along any path in any ϑ^k leads to an equilibrium point (or the special point (0, 0)) with index +1. Hence there is exactly one more equilibrium point with index -1 than with index +1.

What we have done in the above has been to establish an orientation* on certain edges in the "product" graph of \mathfrak{S} and \mathfrak{T} , namely those edges that carry the paths and loops of the ϑ^k , $k \in K$. We might also attempt to orient the edges of the separate graph \mathfrak{S} (and similarly for \mathfrak{T}), by ascribing a "forward" sense to the directed edge $\langle s, s' \rangle$ whenever the directed edge $\langle (s, t), (s', t) \rangle$ represents a forward step in ϑ^k , for some t and k . Rather surprisingly, this works. The same edge $\langle s, s' \rangle$ may well participate in several different paths or loops, but, as Theorem 4 will show, all the orientations induced on it must agree. These induced orientations are illustrated by the multiple arrows in Figs. 2 and 3.**

THEOREM 4. Let s and s' be adjacent nodes in \mathfrak{S} , and let t, k, t', k' be such that (s, t)

*Compare Kuhn's observations [4] on oriented Sperner graphs.

**As it happens, almost every edge shown receives at least one arrowhead; this is because m and n are so small. If we interpret the arrowheads as flow units, then the solution nodes appear as sources and sinks (index +1 and -1 respectively), while the other nodes balance inflow and outflow.

and (s', t) are in ϑ^k and (s, t') and (s', t') are in $\vartheta^{k'}$. Then the \mathfrak{S} -indices of (s, t) and (s, t') are equal. Hence the move from (s, t) to (s', t) in ϑ^k and the move from (s, t') to (s', t') in $\vartheta^{k'}$ are either both forward steps or both backward steps.

Proof. If $t = t'$ the result is trivial. If $t \neq t'$ (and hence $k \neq k'$) we first show that t and t' are adjacent. Indeed, by definition of ϑ^k we have

$$L''(t) \cup [L'(s) \cap L'(s')] \cup \{k\} = K,$$

and since the three terms in this union have cardinality n , $m - 1$ and 1 respectively, the union is disjoint. Similarly, the union

$$L''(t') \cup [L'(s) \cap L'(s')] \cup \{k'\} = K$$

is disjoint, and so we see that the labels of t and t' differ in only one element:

$$L''(t) - L''(t') = \{k'\}; \quad L''(t') - L''(t) = \{k\}.$$

Hence t and t' are adjacent as claimed. It follows that the matrices $A(t)$ and $A(t')$, formed like $B(s)$ and $B(s')$ in Lemma 3, have oppositely-signed determinants.

If we now consider the defining matrices for the \mathfrak{S} -indices of (s, t) and (s, t') (whether or not one of them is in \emptyset), we see that their determinants have the form $\pm(\det B(s))(\det A(t))$ and $\pm(\det B(s))(\det A(t'))$, with opposite choices in the two " \pm " symbols this time, because columns k and k' must be transposed. Since $\det A(t)$ and $\det A(t')$ also have opposite signs, the two indices are equal. Q.E.D.

REFERENCES

- [1] S. S. Cairns, Introductory Topology (Ronald Press, New York 1961).
- [2] Ky Fan, "Simplicial maps from an orientable n-pseudo-manifold into S^m with the octahedral triangulation," J. Comb. Th. 2 (1967), 588-602.
- [3] H. W. Kuhn, "Systematic search on pseudomanifolds," SIAM Journal (to appear).
- [4] H. W. Kuhn, "A new proof of the fundamental theorem of algebra," in Mathematical Programming Study No. 1 (to appear).
- [5] C. E. Lemke and J. T. Howson, Jr., "Equilibrium points of bimatrix games," J. Soc. Indus. Appl. Math. 12 (1964), 413-423.

