

Very Preliminary Draft

A CANONICAL DECOMPOSITION ALGORITHM TO COMPUTE EQUILIBRIA OF N-PLAYER GAMES WITH ARBITRARY ACCURACY

SRIHARI GOVINDAN AND ROBERT WILSON

ABSTRACT. In this draft we extend the theoretical development in [1, Appendix C] to enable the Decomposition Algorithm to compute equilibria with arbitrary accuracy.

1. INTRODUCTION

In [1] we describe a Decomposition Algorithm that computes an ε -equilibrium of a game with any number N of players. For a game with the array G of payoffs from pure strategies, the algorithm traces a connected 1-dimensional path of ε -equilibria of games of the form $G \oplus \lambda g$ from $\lambda = \infty$ to $\lambda = 0$. Here g is a generic vector of bonuses for the players' pure strategies and an ε -equilibrium of G is obtained when the effect of these bonuses is nil because $\lambda = 0$. The key feature of the algorithm is that each player responds optimally to a mixture of profiles proposed by an auxiliary 'coordinator' whose objective is to minimize the difference between his proposals and the players' replies. The algorithm relies on a fixed finite set of feasible proposals for the coordinator that are the vertices of Kuhn's triangulation of the space of profiles of players' mixed strategies. The diameter δ of the simplices of this triangulation fixes the value of ε , and generally $\varepsilon = O(\delta)$, although the exact relation depends on the structure of the payoffs in G .

Figure 1 shows an example of the path of the algorithm, projected onto the simplex of one player's mixtures of three pure strategies. The figure shows the triangulation for $\delta = 1/5$. The coordinator's set of feasible proposals to this player is the set of vertices of this triangulation. The steps of the algorithm yield the sequence of circled points. At each stage the support of the coordinator's mixed strategy is the set of vertices of the smallest subsimplex that contains the circled point. The path starts in the lower right corner and ends in the boundary subsimplex where the path terminates with an arrow.

Date: September 24, 2007.

This work was funded in part by a grant from the National Science Foundation of the United States.

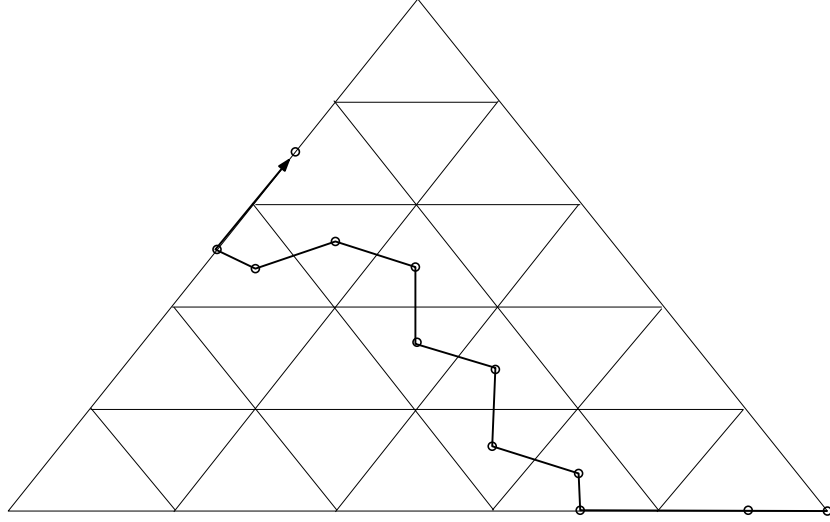


FIGURE 1. Illustrative Path of the Algorithm

In this note we extend the theoretical development in [1, Appendix C] to the case that triangulations with all diameters $d = 1/\alpha$, $\alpha = 1, 2, \dots$ are allowed. This is accomplished by extending the triangulation to the product $Z \times \Sigma$ of the integers Z and the strategy space Σ . This extension enables the Decomposition Algorithm to compute an equilibrium with arbitrary accuracy.

One starts by using the algorithm in [1] with a relatively small scale factor, say α_0 , to obtain an initial approximate equilibrium when $\lambda \downarrow 0$. The algorithm then proceeds through a sequence of values of α until the desired accuracy of the approximation is obtained. The sequence is not monotonic — the values of α can decrease, and indeed, whenever the path returns to α_0 one initially increases λ from zero until it returns to zero, whereupon one again increases α . Thus, the path generates a sequence of ε -equilibria of $G \oplus \lambda g$, where $\varepsilon = O(1/\alpha)$, along which either $\alpha = \alpha_0$ and $\lambda \geq 0$, or $\alpha \geq \alpha_0$ and $\lambda = 0$. One can terminate the computations at any level of α , provided $\lambda = 0$, and thereby obtain an ε -equilibrium where $\varepsilon = O(1/\alpha)$, or continue indefinitely to obtain increasingly accurate approximations of an equilibrium as $\alpha \uparrow \infty$.

Figure 2 illustrates a possible path of the algorithm for the case of a 1-dimensional strategy space, as well as other paths that are loops or inaccessible from $\lambda = \infty$. Figure 3 illustrates the form of the increasingly refined triangulations of the strategy space for $\alpha = 2, 3, 4, 5, 6$.

The algorithm has been implemented in an experimental code in the APL language that is available from the authors. Our initial computational experience indicates that if one wants to terminate with the ε -equilibrium obtained from α then one can start with a value of α_0




FIGURE 2. Illustrative Path of the Algorithm, and Other Paths

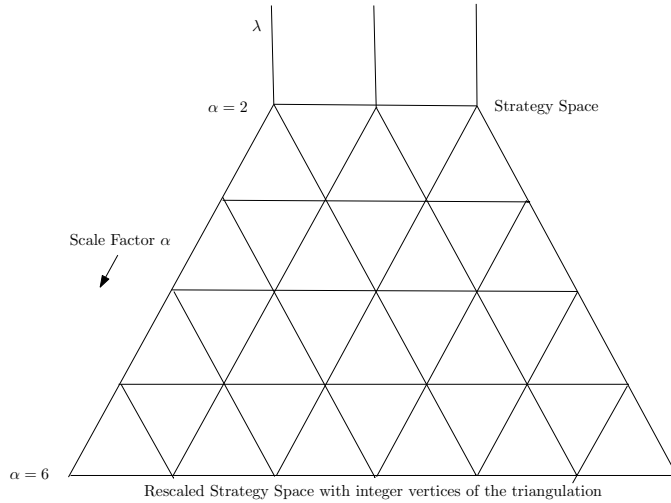


FIGURE 3. Triangulation of $Z \times \Sigma$

that is about 10% of α . Although one can start with $\alpha_0 = 1$, the algorithm can then be slow because there are so many approximate equilibria at this level that the algorithm returns many times to $\alpha_0 = 1$ before finally reaching an approximate equilibrium at this level that thereafter allows $\alpha \uparrow \infty$.

The following sections present the theoretical development of the increasingly refined triangulation in a form that closely parallels the development in [1, Appendix C]. For further details and additional references, consult the text of [1].

2. PRELIMINARIES

The player set is $\mathcal{N} = \{1, \dots, N\}$. Each player $n \in \mathcal{N}$ has a finite set of pure strategies denoted S_n . His set of mixed strategies is denoted Σ_n . $S \equiv \prod_n S_n$; $\Sigma \equiv \prod_n \Sigma_n$; for each n , $S_{-n} \equiv \prod_{m \neq n} S_m$ and $\Sigma_{-n} \equiv \prod_{m \neq n} \Sigma_m$. Let $d - 1$ be the dimension of Σ . The only subsets of S we consider here are product subsets $T = \prod_n T_n$. For each such subset T , denote by $\Sigma(T)$ the corresponding face of Σ , and let $d_T - 1$ be the dimension of $\Sigma(T)$.

With these player and strategy sets a game G is now specified as a collection of payoff functions, one per player, i.e. a game is specified as a vector G in $\mathcal{G} \equiv \mathbb{R}^{\mathcal{N} \times S}$. Given a game G , we write $G_n(\sigma)$ to denote player n 's payoff from the strategy profile $\sigma \in \Sigma$.

For ease in exposition, we label the pure strategies of all the players as consecutive positive integers. Specifically, each strategy set S_n is given as $\{i_{n-1}^* + 1, \dots, i_n^*\}$ where $i_0^* = 0$. (Thus player n has $i_n^* - i_{n-1}^*$ pure strategies, and $d - 1 = i_N^* - N$). Let $I = \{1, \dots, i_N^*\}$ be the set of all the pure strategies for all players. Σ is then a subset of \mathbb{R}_+^I ; and for each game G and each player n , the payoff function G_n has an obvious extension to the whole of \mathbb{R}_+^I , which is still denoted G_n . Let $I^* = \{i_1^*, \dots, i_N^*\}$ be the collection of the ‘‘last’’ strategies of the players. For a typical nonempty subset T_n of S_n , we denote the elements as $i^1(T_n) < \dots < i^*(T_n)$. For $T = \prod_n T_n$, $I(T) \subseteq I$ is the collection of all the strategies in T_n for each n , and $I^*(T)$ is the set of all $i^*(T_n)$.

Let $X = \{\alpha\sigma \in \mathbb{R}^I \mid \alpha \geq 1, \sigma \in \Sigma\}$. For each $x \in X$, $\sum_{i \in S_n} x_i$ is constant across n , and this common number is denoted \bar{x} . X is then d -dimensional. Let Λ be the set of points in X whose coordinates are all integers. For each subset T of S , $X(I(T))$ (or simply $X(T)$) is the subset of X consisting of those x for which $x_i = 0$ for all $i \notin I(T)$, and $\Lambda(T) \equiv \Lambda \cap X(T)$. $X(T)$ has dimension d_T .

For each $i \in I$, let e_i be the i -th unit vector in \mathbb{R}^I ; and, for any subset T of S , $e_T \equiv \sum_{i \in I^*(T)} e_i$. ($e_\emptyset = 0$.)

Given a game G we construct a new game \tilde{G} that has an extra player, called player 0 or the coordinator. The pure strategy set of player 0 is Λ . For each n , his set of pure strategies is $\{x_n \in \mathbb{R}^{S_n} \mid \sum_{i \in S_n} x_i \geq 1\}$. If the coordinator plays λ and the original players play the profile x , then the payoffs of the players in \tilde{G} are as follows. The payoff of player 0 is $\varphi(\lambda, x) = (\lambda + \bar{\lambda}e_I) \cdot (x - .5\lambda)$; for each $n \in \mathcal{N}$, his payoffs depend only on his strategy and that of the coordinator and are given by: $\tilde{G}_n(\lambda, x) = G_n(\lambda_{-n}, x)$. Given that player n 's

payoffs do not depend on the strategies of his opponents in G , we will write $\tilde{G}_n(\lambda, x_n)$ to denote his payoffs when he plays x_n and the coordinator plays λ . Also, we will write $\tilde{G}_n(\cdot, i)$ to denote his payoffs when plays the i -th unit vector in \mathbb{R}^{S_n} .

Let $\tilde{\Sigma}_0$ be the set of mixed strategies for player 0 whose support is finite. Then the payoff functions for all the players have an obvious extension when he employs a strategy in $\tilde{\Sigma}_0$ —we will still use φ and \tilde{G}_n to denote these extended functions.

3. PSEUDO-EQUILIBRIA

Fix a game $G \in \mathcal{G}$. For each player $n \in \mathcal{N}$, his payoff function \tilde{G}_n is linear in his own strategies, which form an unbounded set. Therefore, even his best response correspondence might not be well-defined (e.g. this happens when his payoff function G_n is positive). Consequently, the game \tilde{G} might not have an equilibrium. However, if we weaken the optimality requirement by asking that x_i be zero for each coordinate i that is an inferior reply, then we do get existence. The following definitions lay out this weaker requirement.

Definition 3.1. For each player n , $i \in S_n$ is an *inferior reply* against $\lambda \in \Lambda$ if there exists $j \in S_n$ such that $\tilde{G}_n(\lambda, i) < \tilde{G}_n(\lambda, j)$. A strategy x_n of player n is a *pseudo-best-reply* against λ if $x_i = 0$ for all inferior replies i .

Remark 3.2. An equivalent way to define a pseudo-best-reply is the following: x_n is a pseudo-best-reply if $(\sum_{i \in S_n} x_i)^{-1} x$ maximizes $\tilde{G}_n(\lambda, \sigma_n)$ over the set of $\sigma_n \in \Sigma_n$.

Definition 3.3. $(\sigma_0, x) \in \tilde{S}_0 \times \mathbb{R}^I$ is a *pseudo-equilibrium* of \tilde{G} if:

- (0) σ_0 is a best reply for player 0 against x , i.e. $\varphi(\sigma_0, x) \geq \varphi(\sigma'_0, x)$ for all $\sigma'_0 \in \tilde{\Sigma}_0$;
- (1) For each n , x_n is a pseudo-best-reply against λ .

We will now show the relationship between (certain) pseudo-equilibria of \tilde{G} and the equilibria of G . Let $K > 0$ be such that for each n , $\|G_n(\sigma) - G_n(\sigma')\| \leq K\|\sigma - \sigma'\|$ for all $\sigma, \sigma' \in \Sigma$.

Theorem 3.4. Fix $\varepsilon > 0$. Suppose (σ_0, x) is a pseudo-equilibrium of \tilde{G} such that $x \in X$ and $2K\|x - \lambda\| \leq \varepsilon(\bar{x})^N$ for each λ in the support of σ_0 . Then $\sigma \equiv (\bar{x})^{-1}x$ is an ε -Nash equilibrium, i.e. for each n , $G_n(\sigma_{-n}, \sigma'_n) \leq G_n(\sigma) + \varepsilon$ for all $\sigma'_n \in \Sigma_n$.

Proof. Let $\sigma'_n \in \Sigma_n$ be an arbitrary strategy for some player n . Let σ' be the profile (σ_{-n}, σ'_n) . Also, let $x' = (\bar{x})^{-1}\sigma'$. Observe first that, since (σ_0, x) is a pseudo-equilibrium, the bilinearity of \tilde{G}_n implies that $\tilde{G}_n(\sigma_0, x'_n) \leq \tilde{G}_n(\sigma_0, x_n)$. Next, our assumption on the support of σ_0 and

the Lipschitz nature of G_n imply that:

$$\begin{aligned} |G_n(x) - \tilde{G}_n(\sigma_0, x_n)| &= |G_n(x) - \sum_{\lambda \in \Lambda} \sigma_{0,\lambda} G_n(\lambda_n, x_n)| \leq \sum_{\lambda} \sigma_{0,\lambda} |G_n(x) - G_n(\lambda_n, x_n)| \\ &\leq \sum_{\lambda} \sigma_{0,\lambda} K \|x - (\lambda, x_n)\| \leq .5\varepsilon(\bar{x})^N. \end{aligned}$$

Likewise, $|G_n(x') - \tilde{G}_n(\sigma_0, x'_n)| \leq .5\varepsilon(\bar{x})^N$. Because $\tilde{G}_n(\sigma_0, x'_n) \leq \tilde{G}_n(\sigma_0, x_n)$,

$$G_n(x') - G_n(x) \leq G_n(x') - \tilde{G}_n(\sigma_0, x'_n) + \tilde{G}_n(\sigma_0, x_n) - G_n(x) \leq \varepsilon(\bar{x})^N.$$

Finally, by the multilinearity of G_n , $(\bar{x})^N (G_n(\sigma') - G_n(\sigma)) = G_n(x') - G_n(x)$ and the result is proved. \square

One implication of the above theorem is the following. If $(\sigma_0, x) \in \tilde{\Sigma}_0 \times X$ is a pseudo-equilibrium of \tilde{G} such that \bar{x} is large and the distance between σ_0 and x is small (relative to \bar{x})—i.e. the coordinator’s proposal and the strategies of the players in \mathcal{N} (when normalized) are close—then $(\bar{x})^{-1}x$ is an approximate equilibrium of G . In the next section we construct a subset of $\tilde{\Sigma}_0$ such that equilibria involving strategies in this subset have the property that the coordinator’s proposals are close to the responses of the original players.

4. THE COORDINATOR’S STRATEGIES

In this section, we describe a subset of the coordinator’s mixed strategy set that suffices for computing equilibria. This subset is obtained as follows. First we obtain a triangulation of X with Λ as its vertex set. Then we obtain the subset of mixed strategies for the coordinator consisting of those strategies whose support is included in the vertex set of some principal simplex of this triangulation. The procedure is derived from Kuhn’s procedure for triangulating a cube.

For each $x \in X$, define $z(x) \in \mathbb{R}^I$ by: for each n and $i \in S_n$, $z_i(x) = \sum_{i_{n-1}^* < j \leq i} x_j$. In other words, $z(s)$ is a representation of x as a profile of “cumulative distributions.” Let $Z = \{z(x) \mid x \in X\}$; let $V = \{z(\lambda) \mid \lambda \in \Lambda\}$. Z is easily seen to be the set of all $z \in \mathbb{R}^I$ such that $0 \leq z_{i_{n-1}^*+1} \leq z_{i_{n-1}^*+2} \leq \dots \leq z_{i_n^*} \leq z_{i_n^*}$ for each n , and $z_{i_m^*} = z_{i_n^*} \geq 1$ for all $n, m \in \mathcal{N}$; and V to be the set of points in Z whose coordinates are all integers. We first describe a triangulation of Z with V as the vertex set. The triangulation that it implies for X is then derived in a straightforward manner since the map sending x to $z(x)$ is a linear homeomorphism between X and Z . The following lemma is the basis for our triangulation.

Lemma 4.1. *Each $z \in Z$ has a unique decomposition of the form $z = \sum_{l=0}^k \alpha^l v^l$, where: $\alpha^l > 0$ for all l , and $\sum_l \alpha^l = 1$; $v^l \in V$ for all l , and $v^0 \preceq v^1 \preceq \dots \preceq v^k \leq v^0 + e_I$.*

Proof. Given $z \in X$, define $[z]$ to be the integer part of z , i.e., for all coordinates i , $[z]_i$ is the largest integer smaller than z_i . Clearly $[z]$ belongs to V . Let $r = z - [z]$. All the coordinates of r are nonnegative and strictly smaller than 1; also, the r_i 's are the same for all $i \in I^*$. Let $v^0 = [z]$. If $r = 0$, then set $k = 0$ and $\alpha^0 = 1$. Otherwise, we define a finite sequence a^1, \dots, a^k of positive numbers, and a corresponding sequence of pairwise disjoint subsets I^1, \dots, I^k of subsets of I , inductively, as follows. Let a^1 be the maximum over i of r_i and let I^1 be the coordinates that achieve this maximum. For $l > 1$, let a^l be the maximum over $\{i \notin \cup_{l' < l} I^{l'} \mid r_i \neq 0\}$ of r_i and let I^l be the set of coordinates that achieve this maximum. ($k = l - 1$ for the first integer l for which this maximum does not exist, that is, $r_i = 0$ for all $i \notin \cup_{l' < l} I^{l'}$.) For $0 < l \leq k$, let $v^l = v^{l-1} + e_{I^l}$. We claim now that each v^l belongs to V and that $v^0 \leq v^1 \leq \dots \leq v^k \leq v^0 + e_{I \setminus I^*}$. Indeed, to see this claim, observe that for each n and each strategy $i \neq i_n^*$, $r_i \leq r_{i+1}$ unless $[z]_i + 1 = [z]_{i+1}$; therefore, for each l , $v_i^l \leq v_{i+1}^l$. Also, for each l , I^l either contains I^* or is disjoint from it; therefore, for each l , $v_{i_n^*}^l$ is the same across all n . Thus, our claim is proved. Let $\alpha^0 = (1 - a^1)$ and, for $0 < l \leq k$, let $\alpha^l = (a^l - a^{l+1})$ where a^{k+1} is zero. $0 < \alpha^l < 1$ for all l . Moreover,

$$\begin{aligned} \sum_l \alpha^l v^l &= (1 - a^1)v^0 + (a^1 - a^2)(v^0 + e_{I^1}) + \dots + (a^{k-1} - a^k)(v^0 + e_{I^1} + \dots + e_{I^{k-1}}) \\ &\quad + a^k(v^0 + e_{I^1} + \dots + e_{I^k}) = v^0 + a^1 e_{I^1} + \dots + a^k e_{I^k} = [x] + r = x, \end{aligned}$$

which thus gives us a decomposition that is obviously unique. \square

Remark 4.2. It is easy to see that the $k + 1$ vertices in the lemma can equivalently be written in the form $v^0, v^0 + e_{I^1}, \dots, v^0 + e_{I^k}$ where $\emptyset \neq I^1 \subsetneq \dots \subsetneq I^k \subseteq I$, and for each l , I^l either contains I^* or is disjoint from it.

Define a simplicial complex \mathcal{D} as follows. The vertex set is V . A simplex in \mathcal{D} is the convex hull of a set of vertices in V of the form $v^0 \leq v^1 \leq \dots \leq v^k \leq v^0 + e_I$. It follows from Lemma 4.1 that \mathcal{D} is indeed a simplicial complex that triangulates X .

A simplex of \mathcal{L} is called a principal simplex if it is not a face of another simplex—or equivalently if it has a nonempty interior in X . There exists a simplex characterization of the principal simplices in \mathcal{L} , which we now provide.

Lemma 4.3. *A simplex D of \mathcal{D} is a principal simplex iff there is an ordering (i.e., a bijection) $\pi : \{1, \dots, d\} \rightarrow (I \setminus I^*) \cup \{I^*\}$ such that the vertex set of D is of the form $\{v^0, v^1, \dots, v^d\}$ where $v^l = v^{l-1} + e_{\pi(l)}$ for $1 \leq l \leq d$.*

Proof. As Z is d -dimensional, observe first that the principal simplices are those that have $d + 1$ vertices. Therefore, any simplex whose vertices are generated by an ordering as in the

lemma is clearly a principal simplex. To prove the other way around, suppose that L is a principal simplex with vertex set $v^0 \preceq v^1 \preceq \cdots \preceq v^d$. Since $v^d \leq v^0 + e_I$, and $|I \setminus I^*| = d$, the vertices of L satisfy the following properties: (i) for each $0 < l \leq d$, we have that $v^l - v^{l-1}$ is either e_{I^*} or e_i for some coordinate $i \in I \setminus I^*$; (ii) for each $i \in I \setminus I^*$ there exists a unique l such that $v^l - v^{l-1} = e_i$; and (iii) there exists a unique l such that $v^l - v^{l-1} = e_{I^*}$. There is now an implied ordering, and the vertices of L are obtained from v^0 and this ordering as given in the lemma. \square

Having triangulated Z we now show the properties of the triangulation, call it \mathcal{L} , that it implies for X . It turns out to be easier to characterize the principal simplices in \mathcal{L} than to characterize any arbitrary simplex of \mathcal{L} directly (which is the main reason for first studying Z). Suppose a principal simplex D of Z is obtained by an ordering π and a vertex v^0 as above. For $1 \leq l \leq d$, then v^l is obtained from v^{l-1} by adding 1 to each coordinate in $\pi(l)$. Therefore, for the corresponding points x^l and x^{l-1} in X , we have $x^l = x^{l-1} + \xi_{\pi(l)}$ where $\xi_{\pi(l)}$ equals e_{I^*} if $\pi(l) = I^*$ and it equals $e_{\pi(l)} - e_{\pi(l)+1}$ otherwise. We therefore have the following theorem, whose proof is obvious.

Theorem 4.4. *A simplex L is a principal simplex of \mathcal{L} iff there exists an ordering $\pi : \{1, \dots, d\} \rightarrow (I \setminus I^*) \cup \{I^*\}$ such its vertex set is of the form $\{\lambda^0, \dots, \lambda^d\}$ where for each $k > 0$, $\lambda^k = \lambda^{k-1} + \xi_{\pi(k)}$.*

Every principal simplex in \mathcal{L} then has a compact representation in the form of a vertex-ordering pair (λ, π) . A maximal proper face of a principal simplex L in \mathcal{L} is not a face of another principal simplex if it is contained in the boundary of X ; otherwise it is a face of exactly one other principal simplex.

For each $T \subseteq S$, the triangulation \mathcal{L} induces a triangulation, call it $\mathcal{L}(T)$, of $X(T)$, which we now study. To do so, again we first study the equivalent problem for Z . Let $Z(T)$ be the set of $z(x)$ such that $x \in X(T)$. More directly, $Z(T)$ is the set of $z \in Z$ such that for each n and $i \notin T_n$: $z_i = z_{i-1}$ if $i > i_{n-1}^* + 1$ and zero if $i = i_{n-1}^* + 1$. Let $\mathcal{D}(T)$ be the set of simplices of \mathcal{D} whose vertices belong to $V(T)$. Then $\mathcal{D}(T)$ is a triangulation of $X(T)$.

There exists a simple characterization of the principal simplices of $\mathcal{D}(T)$ analogous to that of the principal simplices of \mathcal{D} . To get at this, we need some more notation. For each n and $i^k(T_n)$ define $R(i^k(T_n)) \equiv \{i \in S_n \mid i^k(T_n) \leq i < i^{k+1}(T_n)\}$, where $i^{k+1}(T_n) = i_n^* + 1$ if $k = *$. Observe that for each $z \in Z(T)$, $z_i = z_{i^k(T_n)}$ if $i \in R(i^k(T_n))$. Therefore, if a vertex v in V is of the form $v^0 + e_{I'}$ for some $v^0 \in V(T)$, v also belongs to $V(T)$ only if for each $i \in I(T)$, I' either contains $R(i)$ or is disjoint from it. The next lemma now follows just like Lemma 4.3, its counterpart for \mathcal{D} .

Lemma 4.5. *A simplex of \mathcal{D} is a principal simplex of $\mathcal{D}(T)$ iff there is an ordering $\pi_T : \{1, \dots, d_T\} \rightarrow (I(T) \setminus I^*(T)) \cup \{I^*(T)\}$ such that the vertex set of D is of the form $v^0 \leq v^1 \leq \dots \leq v^{d_T}$ where: $v^0 \in V(T)$, and for $l \geq 1$, $v^l = v^{l-1} + e_{I^l}$, with $I^l = \cup_{i \in I^*(T)} R(i)$ if $\pi_T(l) = I^*(T)$ and $I^l = e_{R(\pi_T(l))}$ otherwise.*

In the above lemma if x^l and x^{l-1} are points in X that correspond to v^l and v^{l-1} above, then $x^l = x^{l-1} + \xi_{\pi_T(l)}$, where $\xi_{\pi_T(l)}$ is $e_{I^*(T)}$ if $\pi_T(l) = I^*(T)$ and it equals $e_{i^k(T_n)} - e_{i^{k+1}(T_n)}$ if $\pi_T(l) = i^k(T_n)$. The following theorem obtains readily from the previous lemma.

Theorem 4.6. *A simplex L of \mathcal{L} is a principal simplex of $\mathcal{L}(T)$ iff there is an ordering $\pi^T : \{1, \dots, d_T\} \rightarrow (I(T) \setminus I^*(T)) \cup \{I^*(T)\}$ such that the vertex set of L is of the form $\{\lambda^0, \dots, \lambda^{d_T}\}$ where $\lambda^0 \in \Lambda(T)$ and for each $l \geq 0$, $\lambda^{l+1} = \lambda^l + \xi_{\pi_T(l)}$.*

Like with X , a principal simplex in $\mathcal{L}(T)$ can be expressed as a pair (λ^0, π_T) , where $\lambda^0 \in \Lambda(T)$ and π_T is an ordering. For each T , let $\mathcal{L}^*(T)$ be the collection of the principal simplices of $\mathcal{L}(T)$ and let $\mathcal{K}^*(T)$ be the collection of maximal proper faces of the simplices in $\mathcal{L}^*(T)$. Let $\mathcal{L}^* = \cup_T \mathcal{L}^*(T)$ and $\mathcal{K}^* = \cup_T \mathcal{K}^*(T)$. Suppose $K \in \mathcal{K}^*(T)$ for some T , then K is a subset of at most two different elements of $\mathcal{L}^*(T)$; and it is a subset of exactly one element L of $\mathcal{L}^*(T)$ iff it belongs to the boundary of $X(T)$, i.e., $K \in \mathcal{L}^*(I(T) \setminus i)$ for some $i \in I(T)$, or $K \in \Sigma$. If K belongs to two different $L, L' \in \mathcal{L}^*(T)$, we use $L \wedge L'$ to denote K . If $K \in \mathcal{L}^*(I(T) \setminus i)$ and is a subset of $L \in \mathcal{L}^*(T)$, then we write L as $K \vee i$.

For each $L \in \mathcal{L}(T)$, let L_1 be the face of L that is contained in Σ . It is easy to check that the collection of simplices of the form L_1 triangulate Σ and that for each T , the principal simplices of Σ under this triangulation are the sets that are maximal faces of the principal simplices of $\mathcal{L}(T)$. Let $\mathcal{K}_1(T)$ be the collection of the principal simplices of $\Sigma(T)$: L_1 belongs to $\mathcal{K}_1(T)$ iff $L \in \mathcal{L}^*(T)$ and L_1 is a maximal face of L . Let $\mathcal{K}_1 = \cup_T \mathcal{K}_1(T)$. For each T , $K_1 \in \mathcal{K}_1(T)$ and $i \notin I(T)$, let $K_1 \vee_1 i$ be the unique simplex in $\mathcal{K}_1(I(T) \cup \{i\})$ that has K_1 as a maximal proper face. And for $K_1 \neq \tilde{K}_1$ in $\mathcal{K}_1(T)$ such that their intersection is a maximal proper face of each, let $K_1 \wedge_1 \tilde{K}_1$ denote this intersection.

For each $L \in \mathcal{L}^*$ (resp. $K \in \mathcal{K}^*$), let $\Sigma_0(L)$ (resp. $\Sigma_0(K)$) be the set of mixed strategies of the coordinator whose support is a subset of the vertex set of L (resp. K). Let Σ_0 be the union over $L \in \mathcal{L}^*$ of the sets $\Sigma_0(L)$. The properties of the simplices in \mathcal{L}^* imply the following theorem about Σ_0 .

Theorem 4.7. *For $L, L' \in \mathcal{L}^*$: $\Sigma_0(L)$ is a maximal proper face of $\Sigma_0(L')$ iff $L' = L \vee i$ for some i ; the intersection of $\Sigma_0(L)$ and $\Sigma_0(L')$ is a maximal proper face $\Sigma_0(K)$ of each iff $K = L \wedge L'$. These statements also hold for $K, K' \in \mathcal{K}_1$.*

5. THE COORDINATOR'S BEST REPLIES

This section is devoted to an analysis of the best reply correspondence of the coordinator. In view of Theorem 3.4, we will restrict the strategies of the original players to be in X .

For $i, j \in I$, let $\delta_{i,j} = e_i - e_j$. Given $\lambda \in \Lambda$, observe that $\lambda + \delta_{i,j}$ belongs to Λ only if i and j are strategies of the *same* player. Let \mathcal{J} be the collection of all subsets J of I with the property that J contains exactly one strategy for each player. For each $T \subseteq S$, let $\mathcal{J}(T)$ be the set of $J \in \mathcal{J}$ such that $J \subseteq I(T)$. Observe that for $\lambda \in \Lambda$, $\lambda + e_J$ belongs to Λ for $J \subseteq I$ only if J belongs to \mathcal{J} . In the following, when we talk of vectors of the form $\lambda + \delta_{i,j}$ or $\lambda + e_J$ it is to be implicitly understood that $i, j \in S_n$ for some n in the former case and $J \in \mathcal{J}$ in the latter.

We will say that two pure strategies $\lambda, \lambda' \in \Lambda$ are *immediate neighbors* if one is obtained from the other by adding a vector of the form $\delta_{i,j}$ or e_J ; we will say that they are *neighbors* if for each n , there exist subsets S_n^+ and S_n^- of S_n such that: (i) $\lambda' = \lambda + \sum_n (e_{S_n^+} - e_{S_n^-})$; (ii) $|S_n^+| - |S_n^-|$ is independent of n and belongs to $\{-1, 0, 1\}$. Observe that λ and λ' are neighbors if they are immediate neighbors. Furthermore, they are neighbors iff $|\lambda_i - \lambda'_i| \leq 1$ for all $i \in I$, and $|\bar{\lambda} - \bar{\lambda}'| \leq 1$.

When $x \in X$, the coordinator's payoffs are given by $\varphi(\lambda, x) = \lambda \cdot (x - .5\lambda) + N\bar{\lambda} \cdot (\bar{x} - .5\bar{\lambda})$. A simple computation yields the following useful formulae for all $(\lambda, x) \in \Lambda \times X$:

$$\varphi(\lambda, x) - \varphi(\lambda + \delta_{i,j}, x) = \lambda_i - \lambda_j - x_i + x_j + 1$$

and

$$\varphi(\lambda, x) - \varphi(\lambda + e_J, x) = \sum_{i \in J} (\lambda_i - x_i) + N(\bar{\lambda} - \bar{x}) + N.$$

Lemma 5.1. *If $\lambda \in \Lambda$ and $x \in X$ are such that either $|\bar{\lambda} - \bar{x}| \geq 1$, or $|\lambda_i - x_i| \geq 1$ for some $i \in I$, then λ has an immediate neighbor that is a better reply against x than λ . In particular, if λ is a best reply against x , and $x \in X(T)$ for some T , then $\lambda \in X(T)$ as well.*

Proof. The second statement follows trivially from the first, which we now prove. Suppose $\bar{\lambda} \leq \bar{x} - 1$. Then for each n there exists $i \in S_n$ such that $\lambda_i < x_i$. Letting J be the set of these i 's, one per player,

$$\varphi(\lambda, x) - \varphi(\lambda + e_J, x) = \sum_{i \in J} (\lambda_i - x_i) + N(\bar{\lambda} - \bar{x}) + N \leq \sum_{i \in J} (\lambda_i - x_i) < 0,$$

and $\lambda + e_J$ is a better reply against x than λ . If $\bar{\lambda} \geq \bar{x} - 1$, using an analogous argument, it can be shown that the strategy $\lambda - e_J$ is a better reply than λ against x , where J consists of one strategy i per player with the property that $\lambda_i > x_i$.

Suppose that $|\bar{\lambda} - \bar{x}| < 1$ and one of the following holds for some player n : (i) there exists $i \in S_n$ such that $x_i - \lambda_i \geq 1$; or (ii) there exists j such that $\lambda_j - x_j \geq 1$. We now have the following: under case (i) there exists $j \in S_n$ such that $x_j < \lambda_j$; and under case (ii) there exists $i \in S_n$ such that $x_i > \lambda_i$. In both cases, therefore, $\lambda + \delta_{ij}$ is a better reply than λ against σ_n , since

$$\varphi(\lambda, x) - \varphi(\lambda + \delta_{ij}, x) = 1 + \lambda_i - \lambda_j - x_i + x_j < 0.$$

□

The next lemma gives a simple test to determine if a pure strategy is a best reply for the coordinator.

Lemma 5.2. *Suppose λ is at least as good a reply against x as its immediate neighbors. Then λ is a best reply against x . More precisely:*

- (1) *A strategy that is not a neighbor of λ is not a best reply against x .*
- (2) *A neighbor $\lambda' = \lambda + \sum_n (e_{S_n^+} - e_{S_n^-})$ is at least as good a reply as λ iff*
 - (a) *$\lambda + \delta_{ij}$ is a best reply against x for all $n, i \in S_n^+, j \in S_n^-$.*
 - (b) *if $\bar{\lambda}' > \bar{\lambda}$, $\lambda + e_J$ is a best reply for each $J \in \mathcal{J}$ that is contained in $\cup_n S_n^+$.*
 - (c) *if $\bar{\lambda}' < \bar{\lambda}$, $\lambda - e_J$ is a best reply for each $J \in \mathcal{J}$ that is contained in $\cup_n S_n^-$.*

Proof. Since λ is at least as good a reply as each of its neighbors against x , by Lemma 5.1, $|\bar{\lambda} - \bar{x}| < 1$ and $|\lambda_i - x_i| < 1$ for all i . If λ' is not a neighbor of x , then: either $|\bar{\lambda} - \bar{\lambda}'| \geq 2$, in which case $|\bar{\lambda}' - \bar{x}| > 1$; or $|\lambda_i - \lambda'_i| \geq 2$ for some i , in which case $|\lambda'_i - x_i| > 1$. Therefore, in both these cases Lemma 5.1 implies that λ' is not a best reply.

Consider now a neighbor λ' as in Statement (2) of the lemma. Then $|\bar{\lambda} - \bar{\lambda}'| \leq 1$. Suppose $\bar{\lambda} \leq \bar{\lambda}'$, i.e., $\bar{\lambda}' - \bar{\lambda}$ is either 0 or 1. (The other case is analogous.) Then for each n , $|S_n^+| - |S_n^-| = \bar{\lambda}' - \bar{\lambda}$. There is no loss of generality in assuming that for each n , the sets S_n^+ and S_n^- are disjoint. If $\bar{\lambda}' - \bar{\lambda} = 1$, then fix $J \in \mathcal{J}$ such that J contains one strategy in each S_n^+ ; otherwise let J be the empty set. Since the sets $S_n^+ \setminus J$ and S_n^- have the same cardinality for each n , we can now choose any bijection sending for each n , every strategy $i \in S_n^+ \setminus J$ to $j(i) \in S_n^-$. Then,

$$\varphi(\lambda, x) - \varphi(\lambda', x) = \sum_{n, i \in S_n^+ \setminus J} (\varphi(\lambda, x) - \varphi(\lambda + \delta_{i, j(i)})) + \varphi(\lambda, x) - \varphi(\lambda + e_J, x).$$

By assumption the strategies $\lambda + \delta_{i, j(i)}$ and $\lambda + e_J$ are no better than λ . Therefore, λ is at least as good as λ' ; using Statement (1) it is now a best reply to x . Observe now that λ' is an equally good reply iff each of the strategies of the form $\lambda + \delta_{i, j(i)}$ or $\lambda + e_J$ is. Since the

set J and the bijection $i \rightarrow j(i)$ were arbitrary, Statements 2.a and 2.b follow. The proof of 2.c is analogous and therefore omitted. \square

Suppose $L \in \mathcal{L}(T)$ for some T , and it is generated by the pair (λ_0, π_T) . For each player n , let $T_n^+(L)$ be the subset of T_n consisting of those strategies $i^k(T_n)$ such that $\pi^{-1}(i^k(T_n)) < \pi^{-1}(i^{k-1}(T_n))$, i.e. these are coordinates that are first incremented by 1 before being decremented. Let $T_n^-(L) = T_n \setminus T_n^+(L)$ —these coordinates are decremented first by 1 before being incremented. Define $x(L) \in X(T)$ as follows: for each coordinate $i \notin I(T)$, $x_i(L) = 0$; for each n and $i \in T_n$, $x_i(L) = \lambda_i^0 + \alpha_n$ if $i \in T_n^+(L)$ and $x_i(L) = \lambda_i^0 - (1 - \alpha_n)$ if $i \in T_n^-(L)$, where $\alpha_n = (|T_n|(N + \sum_m 1/|T_m|))^{-1}(N - \sum_m |T_m^-(L)|/|T_m| + |T_n^-(L)|(N + \sum_m 1/|T_m|))$. For each n , $0 < \alpha_n < 1$ and thus $|x_i - \lambda_i^0| < 1$ for all i . $\bar{x}(L) = \bar{\lambda}^0 + (|N + \sum_m 1/|T_m|)^{-1}(N - \sum_m |T_m^-(L)|/|T_m|)$ and, therefore, $\bar{x} - \bar{\lambda}^0 < 1$.

Theorem 5.3. *For each T and $L \in \mathcal{L}(T)$:*

- (1) $x(L)$ is the unique point in $X(T)$ against which all the strategies in L are best replies for the coordinator.
- (2) For each n and $i \notin T$, the set of points in $X(\prod_{m \neq n} T_m \times (T_n \cup \{i\}))$ against which the strategies of L_n are best replies is the interval $[x(L), x(L \vee i)]$.

Proof. It is easily checked that λ_0 is at least as good a reply as its immediate neighbors. In fact, we have the following results: (i) an immediate neighbor of the form $\lambda_0 + \delta_{ij}$ is a best reply iff $i \in T_n^+(L)$ and $j \in T_n^-(L)$ for some player n ; (ii) an immediate neighbor of the form $\lambda_0 + e_J$ is a best reply iff $J \subseteq \cup T_n^+(L)$; and (iii) a neighbor of the form $\lambda_0 - e_J$ is never a best reply. Using Lemma 5.2 λ_0 is a best reply against $x(L)$. Also, every strategy in L is a neighbor of λ_0 and is of the form $\lambda_0 + \sum_n (e_{T_n^+} - e_{T_n^-})$ where $T_n^+ \subseteq T_n^+(L)$ and $T_n^- \subseteq T_n^-(L)$ for each n . Therefore, using points (i) and (ii) above, Lemma 5.2 shows that all the strategies in L are best replies. To complete the proof of point (1) we will now show that there exists at most point in $X(T)$ against which the strategies in L are equally good replies.

The set of points in $X(T)$ against which strategies in L are equally good replies satisfy the following system of linear equations in the variable $x \in \mathbb{R}^I$:

$$\begin{aligned} \varphi(\lambda^l, x) - \varphi(\lambda^{l+1}, x) &= 0 & 0 \leq l \leq d^{T_n} - 1 \\ \sum_{i \in T_1} x_i - \sum_{i \in T_n} x_i &= 0 & n \neq 1 \\ x_i &= 0 & i \notin I(T). \end{aligned}$$

It follows from simple linear algebra that there is at most one solution to this system in $X(T_n)$ —viewing the variable x as having $I(T) = d^T + N - 1$ coordinates, by dropping the coordinates not in $I(T)$, the above system has $d^T + N - 1$ linearly independent equations in $d^T + N - 1$ variables. Hence, there exists at most one point in $X(T)$ against which the points in L are best replies.

We turn now to point (2). Fix $i \notin I(T)$. Let C be the set of points in $X(I(T) \cup \{i\})$ against which the points of L are best replies. As before, elementary linear algebra involving the above system shows that C is a convex set whose dimension is at most 1. We know that C includes $x(L)$; and, applying point (1) of this theorem to the vertices of $L \vee i$, we have that it also contains $x(L \vee i)$. Therefore it contains the interval between the two points and in fact lies in the affine space A generated by this interval. Take a point x in A that is outside this interval. We will show that it does not belong to C , which proves point (2). Since x belongs to A , it is of the form $\beta x(L) + (1 - \beta)x(L_n \vee i)$. Furthermore, because $x(L_n)$ belongs to the boundary of $X(I(T) \cup \{i\})$, $\beta < 0$. Let λ be the vertex of $L_n \vee i$ that does not belong to L . Since i belongs to the support of λ but not to $x(L)$, by Lemma 5.1, λ is a worse reply against $x(L)$ than the vertices of L . Because it is a best reply against $x(L \vee i)$, it is, therefore, a strictly better reply against x than each vertex of L . Thus $x \notin C$. \square

We turn now to the case when the vertices of a set $K_1 \in \mathcal{K}_1$ are best replies. Fix a subset T of S . Let $x_1(T) \in \Sigma(T)$ be the strategy profile where each player n mixes uniformly over his strategies in T_n . The proof of the following theorem is similar to that of Theorem 5.3 and is, therefore, omitted.

Theorem 5.4. *For $K_1 \in \mathcal{K}_1(T)$ and $i \notin I(T)$:*

- (1) $x(T)$ is the unique point in Σ against which the strategies in K_1 are best replies.
- (2) The set of points in $\Sigma(I(T) \cup \{i\})$ against which the vertices of K_1 are best replies is $[x_1(T), x_1(T \cup \{i\})]$.

6. THE STRUCTURE OF PSEUDO-BEST-REPLIES

Our objective here is to choose a generic subset of the space \mathcal{G} of games such that for each game G in this set, the pseudo-best-replies of all the original players are “non-degenerate.”

In this section and the next, $g \in \mathbb{R}^I$ is a fixed vector. For each game $G \in \mathcal{G}$ and each $\alpha \in \mathbb{R}$, define the game $\tilde{G} \oplus \alpha g$ as follows: the strategy sets and the payoff function of the coordinator are as in \tilde{G} . The payoff to player n when he plays pure x_n and player 0 plays λ is $\tilde{G}(\lambda, x_n) + \sum_{i \in S_n} \alpha g_i x_i$.

Let $L \in \mathcal{L}(T)$ be a set generated by a pair (λ^0, π_T) for some T . Recall that L_1 is the face of L that is contained in Σ . In case L_1 is nonempty, there exists $0 \leq l_1 < d_T$ such that the vertex set of L_1 is $\{\lambda^0, \dots, \lambda_{l_1}\}$; in case L_1 is empty, take l_1 to be -1 .

Viewing strategies in $\Sigma_0(L)$ as points in \mathbb{R}^L (and not \mathbb{R}^Λ) consider now the following system of equations and inequalities in the variables (G, α, σ_0, v) in $\mathcal{G} \times \mathbb{R} \times \mathbb{R}^L \times \mathbb{R}^\mathcal{N}$:

$$\begin{aligned} \psi_{n,i}(G, \sigma_0, v) &\equiv \tilde{G}_n(\sigma_0, i) + \alpha g_i - v_n = 0 & \forall n, i \in T_n \\ \psi_{n,i}(G, \sigma_0, v) &\equiv \tilde{G}_n(\sigma_0, i) + \alpha g_i - v_n \leq 0 & \forall n, i \notin T_n \\ \psi_{0,*}(G, \sigma_0, v) &\equiv \sum_{k=0}^{d_T} \sigma_{0,k} - 1 = 0 \\ \psi_{0,k}(G, \sigma_0, v) &\equiv \sigma_{0,k} \geq 0 & \forall 0 \leq k \leq d_T \\ \psi_\alpha &\equiv \alpha \geq 0 \\ \psi_\perp &\equiv \alpha \left(\sum_{l>l_1} \sigma_{0,l} \right) = 0. \end{aligned}$$

Let $\Theta(L)$ be the set of solutions to this system and denote by $proj$ the natural projection from $\Theta(L)$ to \mathcal{G} . If $(G, \alpha, \sigma_0, v) \in \Theta(L)$, then for each n , each of the strategies in T_n yields a payoff of v_n in the game $\tilde{G} \oplus \alpha g$, and none of them is inferior; besides, the support of σ_0 is in L_1 unless $\alpha = 0$. This last property, given by the complementary slackness condition ψ_\perp , implies that every point in $\Theta(L)$ satisfies at least one of the inequalities weakly: either $\lambda_0 = 0$, or $\sigma_{0,l} = 0$ for all $l > l_1$. The following lemma shows that for generic games, almost all points satisfy exactly one inequality weakly.

Lemma 6.1. *There exists a closed, lower-dimensional subset \mathcal{G}_L of \mathcal{G} such that for each $G \notin \mathcal{G}_L$, and each $l \geq 1$, the set of points in $proj^{-1}(G)$ that satisfy exactly l of the inequalities weakly is 1-dimensional if $l = 1$; finite if $l = 2$; and empty if $l > 2$.*

Proof. For each positive integer l , let $\Theta^l(L)$ be the subset of $\Theta(L)$ consisting of points (G, α, σ_0, v) for which exactly l of the inequalities are satisfied with equality. To obtain the properties of the lemma, our first task is to prove that $\Theta^l(L)$ is a manifold of dimension $N|S| + 2 - l$. To this end, fix $(G, \alpha, \sigma_0, v) \in \Theta^l(L)$. Let J be the set of indices $j \neq \perp$ such that $\psi_j(G, \alpha, \sigma_0, v) = 0$. By the complementary slackness equation ψ_\perp , either J contains α or it contains $\psi_{0,k}$ for all $k > l_1$. Choose a neighborhood U of (G, α, σ_0, v) such that ψ_j is strictly negative on U for all j other than \perp that is not in J . Define $\psi^J : \mathcal{G} \times \mathbb{R} \times \mathbb{R}^L \times \mathbb{R}^\mathcal{N} \rightarrow \mathbb{R}^J$ to be the function given by the collection $(\psi_j)_{j \in J}$ of coordinate functions. The domain of ψ^J has dimension $N|S| + d_T + N + 2$; while $|J| = \sum_n |T_n| + 1 + l = d_T + N + l$. Therefore, if we

can show that the Jacobian of ψ^J evaluated at (G, α, σ_0, v) has full rank, then the implicit function theorem implies that (G, α, σ_0, v) has a neighborhood V , which can be taken to be a subset of U , such that $(\psi^J)^{-1}(0) \cap V$ is homeomorphic to an open subset of $\mathbb{R}^{N|S|+2-l}$, proving thereby that $\Theta^l(L)$ is a manifold of dimension $N|S|+2-l$. To prove that the Jacobian of ψ^J has full rank, observe that for each $(n, i) \in J$, and coordinate ψ_j of ψ^J , $\partial\psi_j/\partial G_n(i, s_{-n}) = 0$ for all $s_{-n} \in S_{-n}$ if $j \neq (n, i)$, while $\sum_{s_{-n}} \psi_i/\partial G_n(i, s_{-n}) > 0$. Thus, the derivative of $\psi_{n,i}$ is independent of the derivatives of ψ_j for all $j \neq (n, i)$. Since $\sum_k \sigma_{0,k} = 1$, there exists some k such that $(0, k) \notin J$. Therefore, the derivatives of the coordinate functions that correspond to restrictions on player 0's strategies are independent as well. Finally, the derivative of ψ_α if $\alpha \in J$ is obviously independent of the other derivatives as well. Thus the derivative of ψ^J has full rank and $\Theta^l(L)$ is a manifold of dimension $N|S|+2-l$.

Let $proj : \Theta(L) \rightarrow \mathcal{G}$ be the natural projection. For each $l > 2$, let \mathcal{G}_L^l be the closure of $proj(\Theta^l(L))$. Then, \mathcal{G}_L^l is a semialgebraic set of dimension at most $N|S|+2-l < N|S|$. For each $l > 2$ and $G \notin \mathcal{G}_L^l$, $proj^{-1}(G) \cap \Theta^l(L)$ is then empty.

For $l = 1, 2$, applying the generic local triviality theorem to the projection from $\Theta^l(L) \rightarrow \mathcal{G}$, we get that there exists a closed semialgebraic subset \mathcal{G}_L^l of \mathcal{G} such that for each component C of $\mathcal{G} \setminus \mathcal{G}_L^l$, there exists a semialgebraic set F_C such that $proj^{-1}(C) \cap \Theta^l(L)$ is homeomorphic to $C \times F_C$ under a map that sends $\{G\} \times F_C$ to $proj^{-1}(G) \cap \Theta^l(L)$. Since C is open in \mathcal{G} , $proj^{-1}(C) \cap \Theta^l(L)$ is an open subset of the semialgebraic manifold $\Theta^l(L)$ and hence has dimension $N|S|+2-l$. Therefore, F_C , and hence also $proj^{-1}(G) \cap \Theta^l(L)$, has dimension $1-l$. If $G \notin \mathcal{G}_L^l$, then $proj^{-1}(G) \cap \Theta^l(L)$ is $(2-l)$ -dimensional. (In particular, since all sets are semi-algebraic, $proj^{-1}(G) \cap \Theta^2(L)$ is finite.) Let \mathcal{G}_L be the union of the sets \mathcal{G}_L^l for $l \geq 0$. All the properties of the lemma hold for $G \notin \mathcal{G}_L$. \square

For each game G , let $\theta(G, L)$ be the set of points (α, σ_0) such that there exists v for which $(G, \alpha, \sigma_0, v) \in \Theta(L)$. Let $\theta_0(G, L)$ be the subset of $\theta(G, L)$ consisting of points of the form $(0, \sigma_0)$, and let $\theta_1(G, L)$ be the subset consisting of points (α, σ_0) where the support of σ_0 is in L_1 . For any fixed G , the set of solutions to the system of equations and inequalities described in this section are linear in the variables (α, σ_0, v) . Therefore, we now have the following theorems about the structure of $\theta_0(G, L)$ and $\theta_1(G, L)$, whose proofs are immediate from the above lemma.

Theorem 6.2. *For each $G \notin \mathcal{G}_L$, $\theta_0(G, L)$ is a closed interval such that for each $(0, \sigma_0) \in \theta_0(G, L)$:*

- (1) *If $(0, \sigma_0)$ belongs to the relative interior of $\theta_0(G, L)$, then σ_0 belongs to the relative interior of $\Sigma_0(L)$; and each $i \in I \setminus I(T)$ is an inferior reply against σ_0 in the game \tilde{G} .*

- (2) If $(0, \sigma_0)$ belongs to the boundary of $\theta_0(G, L)$ but σ_0 belongs to the relative interior of $\Sigma_0(L)$, then there is exactly one $i \in I \setminus I(T)$ that is not an inferior reply against σ_0 .
- (3) If σ_0 belongs to the relative boundary of $\Sigma_0(L)$ then it belongs to the relative interior of $\Sigma_0(K)$ for some maximal proper subset K of L ; and each $i \in I \setminus I(T)$ is an inferior reply against σ_0 .

Theorem 6.3. For $G \notin \mathcal{G}_L$, $\theta_1(G, L)$ is a closed interval that is nonempty only if L_1 is a maximal face of L , i.e, $L_1 \in \mathcal{K}_1(T)$. Moreover, for $(\alpha, \sigma_0) \in \theta_1(G, L)$:

- (1) If (α, σ_0) belongs to the relative interior of $\theta_1(G, L)$, then $\alpha > 0$; σ_0 belongs to the relative interior of $\Sigma_0(L_1)$; and each $i \in I \setminus I(T)$ is an inferior reply against σ_0 in the game $\tilde{G} \oplus \alpha g$.
- (2) If (α, σ_0) belongs to the boundary of $\theta_1(G, L)$ but σ_0 belongs to the relative interior of $\Sigma_0(L_1)$: then either $\alpha = 0$, and all $i \notin I(T)$ are inferior against σ_0 in \tilde{G} ; or $\alpha > 0$ and there is exactly one $i \in I \setminus I(T)$ that is not an inferior reply against σ_0 .
- (3) If σ_0 belongs to the relative boundary of $\Sigma_0(L_1)$ then it belongs to the relative interior of $\Sigma_0(K)$ for some maximal proper subset K_1 of L_1 ; $\alpha > 0$; and each $i \in I \setminus I(T)$ is an inferior reply against σ_0 in the game $\tilde{G} \oplus \alpha g$.

Let \mathcal{G}^* be the union of the sets $\mathcal{G} \setminus \mathcal{G}_L$ where L ranges over the principal simplices in \mathcal{L}^* . \mathcal{G}^* is then set a set of full measure; more precisely, its complement is a countable union of lower-dimensional semi-algebraic sets. Our final two theorems of the section show the relationship among the various sets $\theta(G, L)$, as L ranges over the principal simplices, for a game $G \in \mathcal{G}^*$. The two theorems have similar proofs and, so, the proof of the second is omitted.

Theorem 6.4. Suppose $G \in \mathcal{G}^*$ and $(0, \sigma_0) \in \theta_0(G, L)$ for some $L \in \mathcal{L}^*(T)$. Let $L' \neq L$ be another principal simplex in \mathcal{L}^* .

- (1) If $(0, \sigma_0)$ belongs to the interior of $\theta_0(G, L)$ then it does not belong to $\theta(G, L')$.
- (2) If $(0, \sigma_0)$ belongs to the boundary of $\theta_0(G, L)$ and σ_0 to the interior of $\Sigma_0(L)$ then $(0, \sigma_0)$ belongs to $\theta(G, L')$ iff $L' = L \vee i$ where i is as in Statement 2 of Theorem 6.2.
- (3) If σ_0 belongs to the interior of a maximal proper face $\Sigma_0(K)$, then $(0, \sigma_0)$ belongs to the $\theta(G, L')$ iff either $K = L'$ or $K = L \wedge L'$.

Proof. If $(0, \sigma_0)$ belongs to $\theta(G, L')$ then it belongs to $\theta_0(G, L')$. By applying Theorem 6.2 to $\theta(G, L')$ observe first that $(0, \sigma_0)$ belongs to $\theta(G, L')$ only if σ_0 belongs to the interior of either $\Sigma_0(L')$ or one of its maximal proper faces. Suppose first that σ_0 belongs to the

interior of $\Sigma_0(L)$. By Theorem 4.7, our observation implies that $(0, \sigma_0)$ belongs to $\Sigma_0(L')$ only if $L' = L \vee i$ for some $i \notin I(T)$. If $(0, \sigma_0)$ belongs to the interior of $\theta_0(G, L)$, then by statement 1 of Theorem 6.2 this coordinate i is an inferior reply and hence $(0, \sigma_0) \notin \theta(G, L)$. On the other hand if $(0, \sigma_0)$ is boundary point, then $\theta(G, L')$ contains $(0, \sigma_0)$ iff $L' = L \vee i$. Thus we have proved Statements 1 and 2.

Suppose now that σ_0 belongs to the relative interior of a maximal proper face $\Sigma_0(K)$. Then, using our observation above and also Theorem 4.7, $(0, \sigma_0)$ belongs to $\theta_0(G, L')$ only if $L' = K$ or $K = L \wedge L'$. Obviously if L' is either of these sets, then $(0, \sigma_0)$ belongs to $\theta_0(G, L')$. Thus, Statement 3 follows. \square

Theorem 6.5. *Suppose $G \in \mathcal{G}^*$ and $(\alpha, \sigma_0) \in \theta_1(G, L)$ for some $L \in \mathcal{L}^*(T)$. Let $L' \neq L$ be another principal simplex in \mathcal{L}^* .*

- (1) *If (α, σ_0) belongs to the interior of $\theta_1(G, L)$ then it does not belong to $\theta(G, L')$.*
- (2) *If (α, σ_0) belongs to the boundary of $\theta_0(G, L)$ and σ_0 to the interior of $\Sigma_0(L)$ then (α, σ_0) belongs to $\theta(G, L')$ iff $\alpha > 0$ and $L' = L \vee i$ where i is as in Statement 2 of Theorem 6.3.*
- (3) *If σ_0 belongs to the interior of a maximal proper face $\Sigma_0(K)$, then (α, σ_0) belongs to the $\theta(G, L')$ iff either $K = L'_1$ or $K = L_1 \wedge L'_1$.*

7. NEWTON PATHS

In this section, we show how to construct Newton paths for our algorithm when the game belongs to \mathcal{G}^* .

Take a game G that belongs to \mathcal{G}^* . Let $E^*(G)$ be the set of $(\alpha, \sigma_0, x) \in \mathbb{R}_+ \times \Sigma_0 \times X$ such that (σ_0, x) is an equilibrium of $\tilde{G} \oplus \alpha g$; and $\alpha = 0$ if the support of σ_0 is not contained in Σ . We analyze the structure of E^* through a series of claims.

Claim 7.1. *(α, σ_0, x) belongs to $E^*(G)$ only if (α, σ_0) belongs to $\theta(G)$.*

Proof. Let (a, σ_0, x) be a point in E^* . Let L^* be the simplex spanned by the support of σ_0 . There exists a unique face $X(T)$ of X whose relative interior contains the relative interior of the simplex L^* . Let L be an element of $\mathcal{L}(T)$ that has L^* as a subset. We will now show that (α, σ_0) belongs to $\theta(G, L)$.

Since the relative interior of L^* is contained in that of $X(T)$, for each $i \in I(T)$ there exists a strategy in L^* whose i -th coordinate is nonzero; and since this strategy is a best reply against x for player 0, it now follows from Lemma 5.1 that $x_i > 0$. Therefore, for each player n , each strategy in T_n is a pseudo-best-reply against σ_0 in the game $\tilde{G} \oplus \alpha g$. It now follows from its definition that $\theta(G, L)$ contains (α, σ_0) . \square

For each $L \in \mathcal{L}^*$, let $E^*(G, L)$ be the set of $(\alpha, \sigma_0, x) \in E^*(G)$ such that $(\alpha, \sigma_0) \in \theta(G, L)$. The above claim says that $E^*(G)$ is the union over all $L \in \mathcal{L}^*$ of $E^*(G, L)$.

Claim 7.2. *Suppose $(0, \sigma_0, x)$ is in $E^*(G, L)$ for some $L \in \mathcal{L}^*(T)$.*

- (1) *If $(0, \sigma_0)$ belongs to the interior of $\theta_0(G, L)$, then $x = x(L)$.*
- (2) *If $(0, \sigma_0)$ belongs to the boundary of $\theta_0(G, L)$ but σ_0 belongs to the interior of $\theta_0(G, L)$ then $x \in [x(L), x(L \vee i)]$ where i is as in Statement 2 of Theorem 6.2.*
- (3) *If σ_0 belongs to the interior of $\Sigma_0(K)$ for a maximal proper face K such that $K \in \mathcal{L}^*$, then $x \in [x(K), x(L)]$.*

Proof. If $(0, \sigma_0)$ belongs to the interior of $\theta_0(G, L)$ then by Statement 1 of Theorem 6.2, $x \in X(T)$. By Theorem 5.3 $x = x(L)$.

We turn now to Statement 2. In this case, by Statement 2 of Theorem 6.2, $x \in X(I(T) \cup \{i\})$. By Statement 2 of Theorem 5.3, $x \in [x(L), x(L \vee i)]$.

Finally, in the case of Statement 3, observe that $L = K \vee i$ for some i . Therefore, it follows from applying Statement 2 to $\theta(G, K)$. \square

The next claim provides the analog of the above result for $\theta_1(G, L)$. Its proof follows like above by combining Theorems 6.3 and 5.4.,

Claim 7.3. *Suppose (α, σ_0, x) is in $E^*(G, L)$ for some $L \in \mathcal{L}^*(T)$ and $(\alpha, \sigma_0) \in \theta_0(G, L)$.*

- (1) *If (α, σ_0) belongs to the interior of $\theta_1(G, L)$, then $x = x_1(T)$.*
- (2) *If $(0, \sigma_0)$ belongs to the boundary of $\theta_0(G, L)$ but σ_0 belongs to the interior of $\theta_0(G, L)$ then either (i) $\alpha = 0$ and $x \in [x_1(L_1), x(L)]$ or (ii) $\alpha > 0$ and $x \in [x(L), x(L \vee i)]$ where i is as in Statement 2 of Theorem 6.3.*
- (3) *If σ_0 belongs to a maximal proper face K of $\Sigma_0(L_1)$ then (i) $x = x_1(T)$ if $K = L_1 \wedge_1 K_1$ for some $K_1 \in \mathcal{K}_1$; (ii) $x \in [x(K), x(K_1)]$ if $K \in \mathcal{K}_1$.*

The only type of equilibrium that is not described by the above claims is the following: suppose $(0, \sigma_0, x) \in E^*(G)$ is such that σ_0 belongs to a maximal proper face of two principal simplices in \mathcal{L}^* , i.e., if its support is of the form $L \wedge L'$. In this case, the vertices of $L \wedge L'$ are obviously best replies against both $x(L)$ and $x(L')$. But, x need not lie in this interval. However, if we look at the subset of $E^*(G)$ where this is true, then we get a 1-dimensional manifold. Accordingly, let $E(G)$ be the set of $(\alpha, \sigma_0, x) \in E^*(G)$ such that if the support of σ_0 is of the form $L \wedge L'$, then $x \in [x(L), x(L')]$.

Theorem 7.4. *$E(G)$ is a 1-dimensional manifold without boundary.*

Proof. Fix $(\alpha, \sigma_0, x) \in E^*(G, L)$ for some $L \in \mathcal{L}^*$. We will prove that it has a neighborhood that is homeomorphic to an interval for the case where (α, σ_0) belongs to $\theta_0(G, L)$. The proof when (α, σ_0) belongs to $\theta_1(G, L)$ is similar. Suppose it belongs to the interior of $\theta_0(G, L)$. By Claim 2, $x = x(L)$, and $\theta_0(G, L) \times \{x\}$ is contained in $E^*(G, L)$. Since σ_0 belongs to the interior of $\theta_0(G, L)$, this set is actually a neighborhood of (α, σ_0, x) in $E^*(G)$ by Statement 1 of Theorem 6.4. Thus, (α, σ_0, x) has a neighborhood that is homeomorphic to an interval in this case.

If (α, σ_0) is in the boundary of $\theta_0(G, L)$ but σ_0 belongs to the interior of $\Sigma_0(L)$, then $x \in [x(L), x(L \vee i)]$. In this case, $(\theta_0(G, L) \times \{x(L)\}) \cup (\{\alpha, \sigma_0\} \times [x(L), x(L \vee i)]) \times (\theta_0(G, L \vee i) \times \{x(L \vee i)\})$ is a neighborhood in $E^*(G)$.

If σ_0 belongs to $\Sigma_0(K)$ for some $K \in \mathcal{L}^*$, then $L = K \vee i$ for some i , and we can apply the previous case to K to get the requisite neighborhood.

If σ_0 belongs to $\theta(G, L')$ for some $L' \neq L$ in $\mathcal{L}^*(T)$, then by assumption, $x \in [x(L), x(L')]$. Therefore, $(\theta_0(G, L) \times \{x(L)\}) \cup (\{\alpha, \sigma_0\} \times [x(L), x(L')]) \times (\theta_0(G, L') \times \{x(L')\})$ is a neighborhood in $E^*(G)$.

Finally, if (α, σ_0) belongs to the $\Sigma_0(K)$ for $K \in \mathcal{K}_1$, then $x \in [x_1(L), x(L)]$ and then $(\theta_1(G, L) \times \{x_1(L)\}) \times (\{\alpha, \sigma_0\} \times [x_1(L), x(L)]) \times (\theta_0(G, L) \times \{x(L)\})$ is a neighborhood of in $E^*(G)$. \square

Suppose the vector g is such that for each n , there exists a unique i such that $g_i > g_j$ for all $j \neq i$ in S_n . Then for all large α , for each player every strategy other than this strategy i is inferior in the game $G \oplus \alpha g$. Thus, the game has a unique pseudo-equilibrium where the players play this profile, call it $x(g)$, and the coordinator also chooses this profile. There is now a unique component of $E^*(G)$ that contains $(\alpha, x(g), x(g))$ for all large α . This component is obviously not compact. For each integer $M \geq 1$, this component must contain an equilibrium $(0, \sigma_0, x)$ where $\bar{x} = M$. This component then gives an approximate equilibrium with arbitrary accuracy.

REFERENCES

- [1] Govindan, S. and R. Wilson (2007), "A Decomposition Algorithm for N-Player Games," *Economic Theory*, special issue on computation of equilibria, to appear 2008.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF IOWA, IOWA CITY IA 52242, USA.

E-mail address: srihari-govindan@uiowa.edu

STANFORD BUSINESS SCHOOL, STANFORD, CA 94305-5015, USA.

E-mail address: rwilson@stanford.edu