

Exponentially Many Steps for Finding a Nash Equilibrium in a Bimatrix Game*

Rahul Savani and Bernhard von Stengel

Department of Mathematics, London School of Economics
Houghton St, London WC2A 2AE, United Kingdom
email: rahul@maths.lse.ac.uk, stengel@maths.lse.ac.uk

March 18, 2004

(minor corrections added May 13, 2004)

CDAM Research Report LSE-CDAM-2004-03[†]

Abstract. The Lemke–Howson algorithm is the classical algorithm for the problem NASH of finding one Nash equilibrium of a bimatrix game. It provides a constructive, elementary proof of existence of an equilibrium, by a typical “directed parity argument”, which puts NASH into the complexity class PPAD. This paper presents a class of bimatrix games for which the Lemke–Howson algorithm takes, even in the *best* case, *exponential* time in the dimension d of the game, requiring $\Omega((\theta^{3/4})^d)$ many steps, where θ is the Golden Ratio. The “parity argument” for NASH is thus explicitly shown to be inefficient. The games are constructed using pairs of dual cyclic polytopes with $2d$ suitably labeled facets in d -space.

1 Introduction

Note: This paper is directed at theoretical computer scientists, so this introduction emphasizes aspects of computational complexity. Our geometric construction of certain bimatrix games in the subsequent sections should be accessible to any game theorist.

Game theory is the formal study of conflict and cooperation. In computer science, game theory has attracted recent interest for economic aspects of the internet, such as electronic commerce, selfish routing in networks [28], and algorithmic mechanism design [23]. In complexity theory, game-theoretic ideas are basic for proving lower bounds for randomized algorithms [36] and in the competitive analysis of online algorithms [2].

*Rahul Savani is supported by an EPSRC doctoral grant. We thank the participants of the Bellairs Workshop on “Polytopes, Games, and Matroids” in Barbados in March 2003, in particular the organizer, Komei Fukuda, and Walter Morris and Jörg Rambau, for stimulating discussions and comments.

[†]This paper supersedes the earlier report LSE-CDAM-2003-14, “Long Lemke–Howson Paths”, which is obsolete.

A *bimatrix game* is a two-player game in strategic form, a basic model in non-cooperative game theory. The strategic form is specified by a finite set of “pure” strategies for each player, and a (for simplicity of input, integer) payoff for each player for each *strategy profile* (which is a tuple of strategies, one for each player). The game is played by each player independently and simultaneously choosing one strategy, whereupon the players receive their respective payoffs. A player is allowed to randomize according to a probability distribution on his pure strategy set, which defines a *mixed strategy* for that player. Players are then interested in maximizing their expected payoffs. A *Nash equilibrium* is a profile of (possibly mixed) strategies such that no player can gain by unilaterally choosing a different strategy, where the other strategies in the profile are kept fixed. Every game has at least one equilibrium in mixed strategies [22]. For two players, the game is specified by two $m \times n$ integer matrices A and B , where the m rows are the pure strategies i of player 1 and the n columns the pure strategies j of player 2, with resulting matrix entries a_{ij} and b_{ij} as payoffs to player 1 and 2, respectively. This is called a bimatrix game (A, B) .

A classic open problem is the complexity of the problem NASH of finding one Nash equilibrium of a bimatrix game. For the special case of zero-sum games, which are bimatrix games $(A, -A)$, this problem generalizes linear programming, but in general is not known to be polynomial. Together with factoring, NASH has been called “the most important concrete open question on the boundary of P today” [26]. A standard method for finding one Nash equilibrium of a bimatrix game is the *Lemke–Howson* (LH) algorithm [14]. In this paper, we present a class of games where this algorithm is exponential. The LH algorithm is a pivoting method related to the simplex algorithm for linear programming [4], which also has worst-case exponential behavior [13]. Our games show that even the *best-case* behavior of the LH algorithm can be exponential, for *any* choice of its free parameter (the first variable “to enter the basis”). To our knowledge, these are the first examples of this kind. Finding a Nash equilibrium in sub-exponential time must therefore go beyond this classic pivoting approach.

Exponentially long paths for finding solutions to a linear complementarity problem (LCP) are described in [21][9]. Although Nash equilibria are solutions to certain LCPs, bimatrix games do not define the LCPs studied in these papers.

NASH belongs to the complexity class TFNP of total function problems in NP [25, p. 229]: With the bimatrix game as input, the required output is a pair of mixed strategies (as the decision problem whether an equilibrium exists is trivial); due to Lemma 1 (of [22]) below, the mixed strategy probabilities are rational numbers of polynomial length as they solve certain linear equations, and the equilibrium property is verified in polynomial time. The class TFNP does not have complete problems unless $NP = co-NP$ [19] (almost by definition, the TFNP-version of SAT, which for a SAT-formula produces a satisfying truth assignment if one exists or else the answer “no”, when reduced to a TFNP-complete problem, would give a problem that is both NP- and co-NP-complete).

More specifically, NASH belongs to the subclass of TFNP call PPAD, of problems with a parity argument for directed graphs [24, p. 516]. The parity argument states that a directed graph (defined implicitly), where the indegree and outdegree of every node is at most one, consists of cycles and directed paths, so that there are as many starting points as endpoints of these paths. An instance of a problem in PPAD is specified by a polynomial-time algorithm for finding at least one starting point, and for finding the neighbor of a point in the graph or else declaring it as an endpoint. The possible endpoints (of which at least one exists) are the allowed function values. The LH algorithm is a special case. It uses a trivial *artificial equilibrium* as starting point, and a freely chosen starting edge, and then a unique “complementary” pivoting rule for determining a next “basic solution”. It thereby traces the vertices of a certain polytope and ends at an equilibrium. The edges of the graph are directed (so the direction of the path can be determined even without knowing the past history) by a geometric orientation [29]; see also Figure 1 below. The parity argument may be inefficient if the paths are not of polynomial length. Our paper shows explicitly that this inefficiency may occur for NASH, by giving games that produce exponentially long Lemke–Howson paths.

A related question is if NASH may be complete for the class PPAD. This seems to require encoding an arbitrary polynomial-time Turing machine computation into complementary pivoting steps for polynomial-sized payoff matrices, which looks difficult.

Unlike the set of solutions to a linear program, or equivalently [4, p. 290] of equilibria of a zero-sum game, the set of Nash equilibria of a general bimatrix game is not convex. Thus, NASH cannot be approached in an obvious way by interior point methods. Moreover, the set of all Nash equilibria is computationally “hard” in the sense that various associated decision problems are NP-complete, e.g. if the game has only one Nash equilibrium, or one with a certain support size, or with a player’s payoff above a given bound [7] (the *support* of a mixed strategy is the set of pure strategies that have positive probability). Therefore, any method for finding one Nash equilibrium, e.g. by divide-and-conquer or incrementally (which is not obvious at any rate) must be weaker than characterizing the set of all equilibria in the end, if such a method is to be polynomial (unless $P = NP$).

The games that we construct are hard to solve for the LH algorithm. Unfortunately, it is easy to guess the support of the equilibrium, and thereby find it, since the only Nash equilibrium of our games is fully mixed. A next step in the construction, which we have not yet done, would thus be to *hide* the equilibrium support so that it is no longer the set of all pure strategies.

Equilibrium enumeration methods (see [33], [35], [10], [1] and the survey in [31]) can be modified to terminate once the first equilibrium is found, and would have to be tested on our games as well. These methods are designed to produce all rather than just one equilibrium, which cannot even be polynomial in the *output* size (unless $P = NP$), since deciding if a game has only one Nash equilibrium is NP-complete ([7], see above). There is no a priori reason to assume that these methods are good for finding just one equilibrium. Similarly,

general algorithms for finding equilibria in games with any number of players are not likely to be fast. These include path-following algorithms [6], which typically specialize to pivoting in the two-player case (for a generalization of LH to more than two players see [27], [34]), and algorithms for approximating fixed points [17][24].

Finding an ε -*approximate* equilibrium (see e.g. [16]) is a different problem than NASH. In our games, all points on the LH path fulfill the equilibrium condition except for one pure strategy, and it is possible that the payoff “error” for that strategy goes below ε after a small number of steps; we have not investigated this. Similarly, a pivoting method implemented in floating-point arithmetic may quickly and erroneously produce an “equilibrium” due to rounding errors. Numerical problems arise since our payoffs are derived from the moment curve, which leads to ill-conditioned matrices. Working implementations of the LH algorithm use exact integer arithmetic [32].

Our work is most closely related to that of Morris [20], who used dual cyclic polytopes to produce exponentially long paths, called “Lemke” paths in [20], for a related method. As we will explain in Section 2, this can be interpreted as the LH method for finding a *symmetric* equilibrium of a symmetric bimatrix game. However, these games have additional non-symmetric equilibria that are found very quickly by the general LH algorithm, and are therefore not useful for our purpose. Morris showed that Lemke paths cannot be used to address the Hirsch conjecture [12]. This famous conjecture states a tight linear bound on the shortest path between any two vertices of a polytope, for which the best known bounds are not even polynomial [11]. A *polynomial* pivoting algorithm for NASH (or even for finding a symmetric Nash equilibrium of a symmetric game, using the symmetrization in (2) below), applied to zero-sum games, would answer that question as well.

Section 3 describes our construction. The LH paths are defined purely combinatorially in terms of the supports of, and best responses to, the mixed strategies that it traces. These correspond to known bit patterns that encode the vertices of dual cyclic polytopes, which are one of the few classes of polytopes whose face structure is known in arbitrary dimension. Linear recurrences for the various path lengths give rise to their exponential growth. The longest path lengths are given by every third Fibonacci number, growing with $\theta^{3d/2}$ for a $d \times d$ game, where θ is the Golden Ratio. Shorter path lengths are obtained by certain sums of these, the shortest length being $\Omega(\theta^{3d/4})$.

Section 4 concludes with open problems.

2 Games, polytopes, and the Lemke–Howson algorithm

Given a bimatrix game (A, B) with $m \times n$ payoff matrices A and B , a mixed strategy for player 1 is a vector x in \mathbb{R}^m with nonnegative components that sum to one, and a mixed strategy for player 2 a similar vector y in \mathbb{R}^n . All vectors are column vectors; the row vector corresponding to x is written as the transpose x^\top . The support of a mixed strategy is the set

of pure strategies that have positive probability. A *best response* to \mathbf{y} is a mixed strategy \mathbf{x} of player 1 that maximizes his expected payoff $\mathbf{x}^\top \mathbf{A}\mathbf{y}$, and a best response to \mathbf{x} is a mixed strategy \mathbf{y} of player 2 that maximizes her expected payoff $\mathbf{x}^\top \mathbf{B}\mathbf{y}$. A Nash equilibrium is a pair of mutual best responses, that is, a mixed strategy pair $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ so that $\bar{\mathbf{x}}^\top \mathbf{A}\bar{\mathbf{y}} \geq \mathbf{x}^\top \mathbf{A}\bar{\mathbf{y}}$ and $\bar{\mathbf{x}}^\top \mathbf{B}\bar{\mathbf{y}} \geq \bar{\mathbf{x}}^\top \mathbf{B}\mathbf{y}$ for all other mixed strategies \mathbf{x} and \mathbf{y} . Best responses are characterized by the following combinatorial condition, which we state only for a mixed strategy \mathbf{x} of player 1.

Lemma 1 [22] *Let \mathbf{x} and \mathbf{y} be mixed strategies of player 1 and 2, respectively. Then \mathbf{x} is a best response to \mathbf{y} if and only if all strategies in the support of \mathbf{x} are pure best responses to \mathbf{y} .*

Proof. Let $(\mathbf{A}\mathbf{y})_i$ be the i th component of $\mathbf{A}\mathbf{y}$, which is the expected payoff to player 1 when playing row i . Let $u = \max_i (\mathbf{A}\mathbf{y})_i$. Then

$$\mathbf{x}^\top \mathbf{A}\mathbf{y} = \sum_i x_i (\mathbf{A}\mathbf{y})_i = u - \sum_i x_i (u - (\mathbf{A}\mathbf{y})_i).$$

Since the sum $\sum_i x_i (u - (\mathbf{A}\mathbf{y})_i)$ is nonnegative, $\mathbf{x}^\top \mathbf{A}\mathbf{y} \leq u$. The expected payoff $\mathbf{x}^\top \mathbf{A}\mathbf{y}$ achieves the maximum u if and only if that sum is zero, that is, if $x_i > 0$ implies $(\mathbf{A}\mathbf{y})_i = u$, as claimed. \square

A game (\mathbf{A}, \mathbf{B}) is *symmetric* if $\mathbf{A} = \mathbf{B}^\top$, so it does not change when the players change roles. The game of “chicken” with $\mathbf{A} = \mathbf{B}^\top = \begin{pmatrix} 2 & 2 \\ 4 & 1 \end{pmatrix}$ is an example. Its equilibria, in terms of probability vectors, are the bottom left pure strategy pair $((0, 1)^\top, (1, 0)^\top)$ with payoffs 4, 2 to players 1, 2, the top right pure strategy pair $((1, 0)^\top, (0, 1)^\top)$ with payoffs 2, 4, and the mixed strategy pair $((1/3, 2/3)^\top, (1/3, 2/3)^\top)$ with payoffs 2, 2. The mixed strategy equilibrium is the only symmetric equilibrium. Its probabilities are uniquely determined by the condition that the pure strategies in the support of the opponent’s strategy must both be best responses (by Lemma 1) and hence have equal expected payoff.

In a mixed equilibrium, the probabilities are uniquely given by the pair of supports if the corresponding sub-matrices have full rank; the support sizes are then equal. This holds if the game is *nondegenerate*, defined by the property that the number of pure best responses to any mixed strategy never exceeds the size of its support (see [31] for a detailed discussion). Even in a degenerate bimatrix game, any Nash equilibrium is a convex combination of extreme equilibria [35][10], which are also determined by linear equalities. The LH algorithm can be extended to degenerate games by standard lexicographic perturbation techniques [14][31]. All games considered here are nondegenerate.

By Lemma 1, an equilibrium is given if any pure strategy of a player is either a best response (to the opponent) or is played with probability zero (by the player himself). This

can be captured by polytopes [37][8] whose facets represent pure strategies, either as best responses or having probability zero. We explain first the simpler case of symmetric equilibria of a symmetric game with $d \times d$ payoff matrix C to player 1, say. Let

$$S = \{z \in \mathbb{R}^d \mid z \geq \mathbf{0}, Cz \leq \mathbf{1}\} \quad (1)$$

where $\mathbf{0}$ and $\mathbf{1}$ denote vectors with all entries 0 and 1, respectively, and inequalities holding for all components. We can assume that C is nonnegative and has no zero column by adding a constant to all payoffs, which does not change the best response structure, so that the polyhedron S is bounded and thus a polytope. We assume there are no redundant inequalities in $Cz \leq \mathbf{1}$, which would correspond to dominated strategies [31]. Then the game is nondegenerate if and only if the polytope S is *simple*, that is, every vertex lies on exactly d facets of the polytope [37][8]. A facet is obtained by making one of the inequalities defining the polytope *binding*, that is, by converting it into an equality.

Lemma 2 [33] *A mixed strategy pair (x, y) is a symmetric Nash equilibrium of the game (C, C^\top) if and only if $x = y = u \cdot z$ and $z \in S$ in (1), $z \neq \mathbf{0}$, $u = 1/\sum_i z_i$, and $z^\top(\mathbf{1} - Cz) = 0$, where z must be a vertex of S by nondegeneracy.*

Proof. Let $z \in S$, $z \neq \mathbf{0}$ and $u = 1/\sum_i z_i$. Then $u > 0$, and $u \cdot z$ is a mixed strategy x . The condition $Cz \leq \mathbf{1}$ is equivalent to $Cx \leq \mathbf{1}u$. The orthogonality condition $z^\top(\mathbf{1} - Cz) = 0$ is equivalent to $x^\top(\mathbf{1}u - Cx) = 0$, so that for each positive component x_i of x (of which there is at least one), $(Cx)_i = u = \max_k (Cx)_k$. Thus, by Lemma 1, x is a best response to itself, that is, (x, x) is a symmetric equilibrium. Conversely, any such equilibrium (x, x) with $u = \max_k (Cx)_k$, which is positive, and $z = 1/u \cdot x$, gives a vector z with the stated properties.

The vector z is on d facets of S since for each i , either $z_i = 0$ or $(Cz)_i = 1$. If z was not a vertex but on a higher-dimensional face of S , any vertex of that face would be on additional facets, contradicting nondegeneracy of the game as S would then not be a simple polytope. \square

In the game of chicken above, $z = (1/6, 1/3)^\top$ gives the symmetric equilibrium. The vector z has to be re-scaled to become a mixed strategy x . The equilibrium payoff u , normalized to 1 in $Cz \leq \mathbf{1}$, is the scaling factor. The converse mapping from x to z defines a projective transformation of a polyhedron representing the upper envelope of expected payoffs to the polytope S [31].

The conditions in Lemma 2 define an LCP [3], usually stated as: find z so that


$$z \geq \mathbf{0}, \quad q + Mz \geq \mathbf{0}, \quad z^\top(q + Mz) = 0,$$

here with data $M = -C$, $q = \mathbf{1}$. This LCP has a trivial solution $z = \mathbf{0}$, which is not a Nash equilibrium. However, $z = \mathbf{0}$ is an *artificial equilibrium* which is the starting point of the LH algorithm (in our first version for symmetric games, giving what [20] calls ‘‘Lemke paths’’).

It is useful to *label* the facets of S [29]. For each pure strategy i , the facets defined by $z_i = 0$ and by $(Cz)_i = 1$ both get label i . Every vertex has the label of the facets it lies on. The complementarity condition $z^\top(\mathbf{1} - Cz) = 0$ then means that z is *completely labeled* (has all labels i), since then each i is either not played or a best response, as required in equilibrium. Since S is simple, such a completely labeled vertex has each label exactly once.

The LH algorithm is started from the completely labeled vertex $z = \mathbf{0}$ by choosing one label k that is *dropped*, meaning that label k is no longer required. This is the only free choice of the algorithm, which from then on proceeds in a unique manner. By leaving the facet with label k , a unique edge is traversed whose endpoint is another vertex, which lies on a new facet. The label, say j , of that facet is said to be *picked up*. If this is the missing label k , the algorithm terminates at a completely labeled vertex. Otherwise, j is clearly *duplicate*, and the next edge is (uniquely) chosen by leaving the facet that so far had label j , and the process repeated. The LH method generates a sequence of *k-almost complementary* edges and vertices (having all labels except possibly k , where k occurs only at the starting point and endpoint). The resulting path cannot repeat a vertex as this would offer a second way to proceed when that vertex is first encountered, which is not the case (since S is simple). Hence, it terminates at a Nash equilibrium.

FIGURE 1



This is illustrated in Figure 1 for dimension 3. Point a is completely labeled, being adjacent to facets with labels 1, 2, 3. Dropping label 1, it proceeds to point b picking up label 2, now duplicate. The next point is c with duplicate label 3, and finally d where the missing label 1 is picked up, which terminates the path.

As in the simplex algorithm [4], edge traversal is implemented algebraically by *pivoting* with variables entering and leaving a basis, the nonbasic variables representing the facets. The only difference is the rule for choosing the next entering variable, which in linear programming is done such as to improve the objective function. Here, it is the *complementary pivoting* rule where the nonbasic variable with duplicate label enters the basis.

Furthermore, the path is directed, giving a “directed parity argument” [24] which puts the problem in the class PPAD, rather than just in PPA. In Figure 1, the starting point a has an *orientation*, with the labels 1, 2, 3 in clockwise order. When label 1 is dropped, the remaining labels keep their orientation (in one dimension less) relative to the edges of the path. In Figure 1, label 2 is always to the left and label 3 always to the right of the edge. At the endpoint of the path, the missing label is picked up at the other end of the edge, so that

the orientation of that vertex is opposite to that of the starting vertex of the path; in Figure 1, point d has labels 1, 2, 3 in anticlockwise order. This generalizes to higher dimensions, where orientation is defined as the sign of a certain determinant. The endpoints of any LH path have opposite orientation, which leads to an *index theory* of equilibria [29][6]. Knowing the orientation of the artificial equilibrium, the orientation of any k -almost complementary edge can be determined directly, which gives the PPAD property.

For nonsymmetric bimatrix games (A, B) , or even for finding nonsymmetric equilibria of symmetric games as in the game of “chicken” above, the LH algorithm is applied as follows, which is its standard form. Let

$$z = \begin{pmatrix} x \\ y \end{pmatrix}, \quad C = \begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix}. \quad (2)$$

The polytope S of dimension $d = m + n$ in (1) is then the *product* $P \times Q$ of the polytopes

$$P = \{x \in \mathbb{R}^m \mid x \geq \mathbf{0}, B^\top x \leq \mathbf{1}\}, \quad Q = \{y \in \mathbb{R}^n \mid Ay \leq \mathbf{1}, y \geq \mathbf{0}\}. \quad (3)$$

Any Nash equilibrium (x, y) of (A, B) is again given by $z^\top(\mathbf{1} - Cz) = 0$, which is equivalent to $x^\top(\mathbf{1} - Ay) = 0$ and $y^\top(\mathbf{1} - B^\top x) = 0$. These conditions state that x is a best response to y and vice versa, where x and y have to be normalized to represent mixed strategies. The only difference to Lemma 2 is that this normalization has to be done separately for x and y , rather than for the entire vector z . It is easy to see that in equilibrium $x = \mathbf{0}$ if and only if $y = \mathbf{0}$, where $(\mathbf{0}, \mathbf{0})$ is the artificial equilibrium.

The LH algorithm is then applied as before, where a label corresponds either to a strategy i of player 1 or a strategy j of player 2. These have to be distinct, so it is convenient to number the n strategies of player 2 as $m + 1, \dots, m + n$, as suggested in [29]. A duplicate label then represents a pure strategy that has both probability zero and is a best response. If this is a strategy i of player 1, for example, it determines the facet $x_i = 0$ in P or $(Ay)_i = 1$ in Q , corresponding to the respective i th inequality in (3).

The LH path on the edges of $S = P \times Q$ is a subgraph of the *product graph* of the edge graphs of P and Q . This means that edges are alternately traversed in P and Q , keeping the vertex in the other polytope fixed. A duplicate label picked up in Q is dropped in P and vice versa. This is the standard view of the LH algorithm; for further details see [31].

3 Lemke–Howson on labeled dual cyclic polytopes

We construct square games with $m = n = d$ strategies for each player. Similar to [30], these are derived from *dual cyclic polytopes* [37][8] in dimension d with $2d$ facets. A standard way of obtaining a cyclic polytope P' in dimension d with $2d$ vertices is to take the convex hull of $2d$ points $\mu(t_i)$ on the *moment curve* $\mu: t \mapsto (t, t^2, \dots, t^d)^\top$ for $1 \leq i \leq 2d$.

2d. However, the polytopes in (3) are defined by inequalities and not as the convex hull of points. In the dual of a polytope, its vertices are re-interpreted as normal vectors of facets. The polytope P' is first translated so that it has the origin $\mathbf{0}$ in its interior, for example by subtracting the arithmetic mean $\bar{\mu}$ of the points $\mu(t_i)$ from each such point. The resulting vectors $c_i = \mu(t_i) - \bar{\mu}$ then define the dual cyclic polytope

$$P'' = \{z \in \mathbb{R}^d \mid c_i^\top z \leq 1, 1 \leq i \leq 2d\}.$$

A vertex u of such a polytope is characterized by a bitstring $u_1 u_2 \cdots u_{2d}$ of length $2d$, with the k th bit u_k indicating whether u is on the k th facet ($u_k = 1$) or not ($u_k = 0$). The polytope is simple, so exactly d bits are 1, the other d bits are 0. Assume that $t_1 < t_2 < \cdots < t_{2d}$ when defining the k th facet of P'' by the binding inequality $(\mu(t_k) - \bar{\mu})z \leq 1$. Then the vertices of P'' are characterized by the 0-1 strings fulfilling the *Gale evenness* condition [5]: A bitstring represents a vertex if and only if any substring of the form $01 \cdots 10$ has even length, so 0110 , 011110 , etc., is allowed, but not 010 , 01110 , and so on. A maximal substring of 1's is called a *run*. We only consider *even* dimension d , where the allowed odd runs of 1's at both ends of the string can be glued together to form an even run, which shows the cyclic symmetry of the Gale evenness condition. Let $G(d)$ be the set of these Gale evenness bitstrings of length $2d$ with d ones.

Both P and Q in (3) will be dual cyclic polytopes with a special order of their inequalities corresponding to the facet labels. A suitable affine transformation [30, p. 560] gives P from P' , and Q in a similar manner, so that the first d inequalities (for the pure strategies of player 1) in P have the form $x \geq \mathbf{0}$, and the second d inequalities (for the pure strategies of player 2) in Q are $y \geq \mathbf{0}$. The remaining d inequalities $B^\top x \leq \mathbf{1}$ in P and $Ay \leq \mathbf{1}$ in Q then determine the game (A, B) . For further details of the construction see Appendix A. The game data is of polynomial size in d (so the running time of an algorithm is polynomial in the size of the game if and only if it is polynomial in d).

The equilibrium condition and the LH algorithm depend on which facets a vertex belongs to, as encoded in the Gale evenness bitstrings in $G(d)$, and on the facet labels. These are defined by permutations l and l' of $1, \dots, 2d$ for P and Q , respectively. For a vertex u of P , which we identify with its bitstring in $G(d)$, its labels are given by $l(k)$ where $u_k = 1$, and the labels of a vertex v of Q are $l'(k)$ where $v_k = 1$, for $1 \leq k \leq 2d$. The k th facet of P (corresponding to the k th position in a bitstring) has label $l(k) = k$, so l is simply the identity permutation. The k th facet of Q has label $l'(k)$. The permutation l' has the fixed points $l'(1) = 1$ and $l'(d) = d$, and otherwise exchanges adjacent numbers, as follows:

$$l'(k) = \begin{cases} k, & k = 1, d, \\ k + (-1)^k, & 2 \leq k \leq d - 1, \\ k - (-1)^k, & d + 1 \leq k \leq 2d. \end{cases} \quad (4)$$

The artificial equilibrium e_0 is a vertex pair (u, v) so that u is labeled with $1, \dots, d$ and v with $d+1, \dots, 2d$. In terms of bitstrings, $u = 1^d 0^d$ (which are d ones followed by d zeros) and $v = 0^d 1^d$, which both fulfill Gale evenness, and have the indicated labels under l and l' , respectively, so that

$$e_0 = (1^d 0^d, 0^d 1^d) \in G(d) \times G(d). \quad (5)$$

A similar Nash equilibrium e_1 is readily identified, which has full support.

Lemma 3 *Let $e_1 = (0^d 1^d, 1^d 0^d)$. This is the only Nash equilibrium of the game.*

Proof. Let (u, v) be a completely labeled vertex pair, and suppose that $u_d = 1$. If $u_{d+1} = 1$, then $v_d = 0$ and $v_{d+2} = 0$ (via complementarity, since $l'(d+1) = d+2$) so $v_{d+1} = 0$ by Gale evenness, and thus $u_{d+2} = 1$. Continuing in that way, all 1's to the right of the d th bit u_d of u (which is 1) have to come in pairs. Similarly, if $u_{d-1} = 1$, then $v_{d-2} = 0$ by complementarity, which with $v_d = 0$ implies $v_{d-1} = 0$. This means that the 1's to the left of u_d come in pairs if there is a zero to the left of them. In the latter case, the run of 1's containing u_d has odd length, so it must include u_{2d} , but then is too long. Hence, the only possibility where $u_d = 1$ is when $(u, v) = e_0$. Similarly, $u_d = 0$ implies $(u, v) = e_1$. \square

Hence, all LH paths, for any dropped label, lead from e_0 to e_1 . Denote by $\pi(d, k)$ the path when label k is dropped in dimension d , regarded as a sequence $(u^0, v^0) (u^1, v^1) \dots (u^L, v^L)$ of vertex pairs in $P \times Q$, that is, in $G(d) \times G(d)$, and let $L(d, k) = L$ be the length of that path.




FIGURE 2




FIGURE 3

As an example, Figure 2 shows $\pi(4, 8)$, with P on the left and Q on the right. The numbers at the top are the labels $l(k)$ and $l'(k)$ for $k = 1, \dots, 8$. The starting point e_0 is the vertex pair $e_0 = (u, v) = (11110000, 00001111)$. We first drop label 8 in Q , so the bit v_7 (since $l'(7) = 8$) changes from 1 to 0, which by Gale evenness gives the bit string 10001101 as the new vertex v' in Q . In Figure 2, u is connected to both v and v' by a sloped line. These sloped lines (forming the middle zigzag path) indicate the vertex pairs on the LH path, which we use since in each step only one vertex changes but the other stays fixed. In v' , label 1 has been picked up, which is now duplicate and dropped in P , giving the next vertex 01111000. The new duplicate label is 5 and in the next step dropped in Q , giving vertex 10011001. In that manner, the path proceeds until it ends at e_1 .

All paths can be expressed in terms of the two special paths $\pi(d, 1)$ and $\pi(d, 2d)$. These have certain symmetries.

Lemma 4 *Let $L = L(d, 1)$ and let (u^i, v^i) be the i th vertex pair of the path $\pi(d, 1)$. Then for $0 \leq i \leq L$, $(u^i, v^i) = (v^{L-i}, u^{L-i})$.*

Proof. The particular names of the labels do not matter, so we can re-name them for both P and Q with the permutation l' in (4), the k th facet in P getting label $l'(l(k))$, which is $l'(k)$, and in Q label $l'(l'(k))$, which is $l(k)$. But then P and Q switch roles, e_0 is exchanged with e_1 , label 1 stays the same, and the path backwards corresponds to $\pi(d, 1)$ itself as claimed. Examples are Figures 5, 7, 10 for $d = 2, 4, 6$ in Appendix B. \square

For the path $\pi(d, 2d)$, we disregard the first vertex pair and the last two vertex pairs. The remaining path, which will call $B(d)$, is point-symmetric in each polytope, by reversing the bitstrings while ignoring the zero bit for the missing label. Figure 3 shows this for $d = 4$ where the disregarded rows and columns are struck out. Examples are Figures 6, 9, 11, 12 for $d = 2, 4, 6, 8$ in Appendix B.

Lemma 5 *Let $L = L(d, 2d)$ and let (u^i, v^i) be the i th vertex pair of the path $\pi(d, 2d)$ for $0 \leq i \leq L = L(d, 2d)$. Then for $1 \leq i \leq L - 2$,*

$$u_k^i = u_{2d-k}^{L-1-i} \quad (1 \leq k \leq 2d - 1), \quad (6)$$

$$v_1^i = v_{2d}^{L-1-i} = 1, \quad (7)$$

$$v_k^i = v_{2d-k}^{L-1-i} \quad (2 \leq k \leq 2d - 2), \quad (8)$$

$$u_{2d}^i = v_{2d-1}^i = 0. \quad (9)$$

In (u^1, v^1) , the duplicate label is 1, which is then dropped in P, and never picked up again.

Proof. An example for the following arguments is provided by Figure 2. Equation (9) holds because label $2d$ is missing for all $i = 1, \dots, L - 2$. After one step on $\pi(d, 2d)$, the vertex pair

$$(u^1, v^1) = (1^d 0^d, 10^{d-1} 1^{d-2} 01) \quad (10)$$

is reached. The duplicate label is 1, which has been picked up in Q, and will next be dropped in P from u^1 . The last vertex pair of $\pi(d, 2d)$ is $(u^L, v^L) = e_1^d$. This is reached by picking up label $2d$ in P. The previous vertex pair is therefore $(u^{L-1}, v^{L-1}) = (0^{d-1} 1^d 0, 1^d 0^d)$, where label d is duplicate. The vertex pair (u^{L-2}, v^{L-2}) is therefore

$$(u^{L-2}, v^{L-2}) = (0^{d-1} 1^d 0, 1^{d-1} 0^d 1), \quad (11)$$

with duplicate label $2d - 1$. This vertex pair is reached from

$$(u^{L-3}, v^{L-3}) = (0^{d-2} 1^d 00, 1^{d-1} 0^d 1) \quad (12)$$

by picking up this label $2d - 1$ in P . Equations (10) and (11) describe the starting vertex pair (u^1, v^1) and ending vertex pair (u^{L-2}, v^{L-2}) of the path $B(d)$.

The mapping r defined by $r(k) = 2d - k$ for $1 \leq k \leq 2d - 1$, and $r(2d) = 2d$, serves as a relabeling to prove the claimed symmetry of $B(d)$. Equation (6) is equivalent to

$$u_k^i = u_{r(k)}^{L-1-i} \quad (1 \leq k \leq 2d - 1). \quad (13)$$

We will show shortly that (7) and (8) together are essentially equivalent to the equations, similar to (13),

$$v_{l'(k)}^i = v_{l'(r(k))}^{L-1-i} \quad (1 \leq k \leq 2d - 1). \quad (14)$$

When label k is dropped in Q from v^i , then the bit v_j^i in position j so that $l'(j) = k$ changes from one to zero. Since l' is its own inverse, $j = l'(k)$. So, (14) has to be read as: the bit of vertex v^i in *position* $l'(k)$, which has label k , is equal to the bit of vertex v^{L-1-i} in position $l'(r(k))$, which has label $r(k)$.

In addition to the two equations (13) and (14), we will show by induction on i : If label k is duplicate in the vertex pair (u^i, v^i) , then label $r(k)$ is duplicate in the vertex pair (u^{L-1-i}, v^{L-1-i}) , and when label k is dropped in one polytope (P or Q) going from (u^i, v^i) to (u^{i+1}, v^{i+1}) , then label $r(k)$ is dropped in the same polytope going backwards from (u^{L-1-i}, v^{L-1-i}) to (u^{L-2-i}, v^{L-2-i}) .

Equations (10) and (11) show that (13) holds for $i = 1$, and the duplicate label is 1 in (u^1, v^1) and $r(1)$ in (u^{L-2}, v^{L-2}) . When i is odd, then in the step from (u^i, v^i) to (u^{i+1}, v^{i+1}) the duplicate label is dropped in P , and $v^i = v^{i+1}$. Similarly, the backwards step from (u^{L-1-i}, v^{L-1-i}) to (u^{L-2-i}, v^{L-2-i}) is also done by dropping the label in P . If in these two steps, the duplicate labels are k and $r(k)$, respectively, then (13) also holds for $i + 1$ instead of i , because r preserves the Gale evenness condition (r is a reversal of the bitstrings and cyclic shift by one position), so the vertices u^i and u^{i+1} are joined by an edge in P if and only if u^{L-1-i} and u^{L-2-i} are joined by an edge (defined by the same labels, using the relabeling r) in P . Moreover, a new label, say k' , is picked up in P which is the duplicate label of the vertex pair (u^{i+1}, v^{i+1}) , and the duplicate label in the vertex pair (u^{L-2-i}, v^{L-2-i}) is then $r(k')$. This shows the inductive step from i to $i + 1$ when i is odd.

We now consider Q . The set $\{2, 3, \dots, 2d - 2\}$ is mapped to itself under both r and l' defined in (4); note that both bijections map d to itself. It is then easy to see that (8) is equivalent to (14) for $2 \leq k \leq 2d - 2$. For $k = 1$, we have $l'(k) = 1$ and $l'(r(k)) = 2d$, so (14) for $k = 1$ (or $k = 2d - 1$, which gives the same equation) follows from (7); we will prove the stronger assertion (7).

As inductive hypothesis, assume that for some i equations (7), (13), and (14) hold; moreover, that the duplicate label is k and dropped from v^i in Q (that is, i is even), and that the duplicate label to be dropped when going backwards from v^{L-1-i} to v^{L-2-i} is $r(k)$. This is true for $i = 2$ by (12), where $k = d + 1$. Suppose that label k' is picked up in Q in vertex v^{i+1} . We want to show that v^{L-2-i} has the new label $r(k')$.

Call two labels (of any vertex) in Q *adjacent* if they are the labels of two cyclically adjacent positions in the bitstring, as determined by l' . Unlike in P , if some labels a and b are adjacent in Q , then there are cases where $r(a)$ and $r(b)$ are not adjacent. This occurs when (a, b) is any one of the pairs $(1, 3)$, $(2d - 3, 2d)$, $(2d, 2d - 1)$. Similarly, the labels $r(a)$ and $r(b)$ are adjacent when (a, b) is $(2d - 1, 2d - 3)$, $(3, 2d)$, or $(2d, 1)$ but then a and b are not. However, we show that this does not matter, considering each of these three cases (and their symmetric counterparts) in turn.

Because of the inductive assumption (7), and (9), the adjacent labels $2d, 2d - 1, 1$ in Q (in positions $2d - 1, 2d$, and 1 , respectively) correspond to the bits $0, 1, 1$ for both current vertices v^i and v^{l-1-i} . First, suppose that when dropping label k from v^i , it matters that labels 1 and 3 are adjacent. That is, the run of ones that is changed in going from v^i to v^{i+1} includes positions 1 and 2 (which have labels 1 and 3 , respectively). This would cause a problem since labels $r(1) = 2d - 1$ and $r(3) = 2d - 3$ in v^{l-1-i} are not adjacent (because they correspond to positions $2d$ and $2d - 2$, respectively). This could occur in two cases: First, when $k = 2d - 1$, because then the above bits $0, 1, 1$ would change to $0, 0, 1$ and the even run of ones starting in position $2d$ (with label $2d - 1$) would “shift to the right” across positions 1 and 2 , which have labels 1 and 3 in Q . But then, by the inductive assumption, the dropped label $r(k)$ when going backwards from v^{l-1-i} to v^{l-2-i} is 1 , which is the label of the rightmost bit in $0, 1, 1$, which shifts left to become $1, 1, 0$, so that the label picked up in v^{l-2-i} would be $2d$, contradicting (9). The second case occurs when the said string $0, 1, 1$ would shift left to become $1, 1, 1$ (where the third one, in position 1 , is shifted in from position 2), again contradicting (9).

Second, suppose it matters that labels $2d - 3$ and $2d$ are adjacent in v^i but $r(2d - 3)$ and $r(2d)$ in v^{l-1-i} are not. This can only happen when label $2d$ is picked up, which it is not, by (9). Third, the fact that labels $2d$ and $2d - 1$ are adjacent in v^i but $r(2d)$ and $r(2d - 1)$ in v^{l-1-i} are not does not matter either, because label $2d$ is not picked up.

These three cases have their counterparts where labels $r(2d - 1)$ and $r(2d - 3)$ in v^{l-1-i} are adjacent, but $2d - 1$ and $2d - 3$ in v^i are not; the reasoning is identical to the first case above. The second and third case are that $r(3)$ and $r(2d)$, and $r(2d)$ and $r(1)$, respectively, are adjacent in v^{l-1-i} , which is again unproblematic.

This completes the inductive step from i to $i + 1$ for all i , so that equations (6) to (9) hold throughout. \square

Two vertices of $G(n)$ are connected by an edge if and only if the corresponding bitstrings differ only by two substrings which are $1^k 0$ for one bitstring and $0 1^k$ for the other (where k is even), using the cyclic symmetry of the Gale evenness bitstrings if necessary. For example, the vertices v^0 and v^1 in Figure 2 are 00001111 and 10001101 , where the substrings are those in positions $7, 8, 1$. These use the cyclic symmetry since the substrings in question involve both position $2d$ and position 1 . We say that such an edge *wraps around*

(the end of the 0-1 string). If the mentioned substrings 1^k0 and 01^k are contiguous substrings of positions 1 through $2d$, the edge does *not* wrap around, like, for example, the edge connecting vertices $u^0 = 11110000$ and $u^1 = 01111000$ in Figure 2.

Lemma 6 *No edge of $\pi(d, 1)$ wraps around (in either polytope): If the edge connects (u, v) to (u', v) , then the edge connecting u and u' in P does not wrap around, and if the edge of $\pi(d, 1)$ connects (u, v) to (u, v') , then the edge connecting v and v' in Q does not wrap around.*

Proof. The first edge of $\pi(d, 1)$ joining e_0 to $(01^d0^{d-1}, 0^d1^d)$ does not wrap around, and neither does the last edge joining $(0^d1^d, 01^d0^{d-1})$ to e_1 . In all other edges, position 1 is zero in both polytopes, so none of these edges wraps around either. \square

For any two paths π and π' on $G(n) \times G(n)$, we denote by $\pi + \pi'$ the path obtained by joining the last vertex pair of π to the first pair of π' , assuming this is possible. The length of the new path is the sum of the lengths of π and π' plus one; the number of *vertex pairs* is simply the respective sum.

The following central theorem describes how paths $\pi(d, 1)$ and $\pi(d, 2d)$ are composed of other such paths, possibly from lower dimension. Appendix B shows these paths as patterns of bitstrings that illustrate this, as indicated in detail in the proof.

Theorem 7 *Let $A(d) = \pi(d, 1)$ and $B(d) = (u^1, v^1) \dots (u^{L-2}, v^{L-2})$ where (u^i, v^i) is the i th vertex pair of $\pi(d, 2d)$, $0 \leq i \leq L = L(d, 2d)$. Then there are paths $C(d)$ and mappings $\alpha, \beta, \beta', \gamma, \gamma'$ defined on vertex pairs, and extended to sequences of vertex pairs, so that*

$$A(d) = \beta(B(d)) + C(d), \quad (15)$$

$$C(d) = \alpha(A(d-2)) + \beta'(B(d)), \quad (16)$$

$$B(d) = \gamma(A(d-2)) + \gamma'(C(d-2)). \quad (17)$$

Proof. Overview: The path $C(d)$ is simply a tail segment of $A(d)$. The mappings are given as follows: β and β' are defined on $G(d) \times G(d)$,

$$\beta(u, v) = (u, 0v_2v_3 \dots v_{2d-2}1v_{2d}),$$

and β' is determined by β due to Lemma 4. Furthermore, $\alpha, \gamma, \gamma': G(d-2) \times G(d-2) \rightarrow G(d) \times G(d)$. With \overleftarrow{u} defined as the bitstring u reversed,

$$\alpha(u, v) = (0\overleftarrow{u}110, 0\overleftarrow{v}110). \quad (18)$$

With $c = 2d - 4$,

$$\gamma(u_1 \dots u_c, v) = (u_111u_2 \dots u_c00, 10v01). \quad (19)$$

We obtain γ' from γ by Lemma 5.

First we show, equivalent to (15) and (16), that

$$A(d) = \beta(B(d)) + \alpha(A(d-2)) + \beta'(B(d)). \quad (20)$$

Note that only positions 1 and $2d-1$ in Q (corresponding to the missing label in $A(d)$ and $B(d)$, respectively) are changed by the mapping β , and these positions are constant throughout $B(d)$ by (7) and (9). The starting point (u^1, v^1) of $B(d)$ is given by (10), and in the first step of $B(d)$ label 1 is dropped in P . The path $A(d)$ is also started by dropping label 1 in P from e_0^d . Now $\beta(u^1, v^1) = e_0^d$, as required, and in the first step of $A(d)$ and $B(d)$ the label to be dropped is 1 in P . As (u^1, v^1) and e_0^d differ only in positions that are constant throughout $B(d)$, the path $B(d)$ maps to $\beta(B(d))$ and thereby represents the initial part of $A(d)$. An example is $B(6)$ in Figure 11 (omitting from $\pi(6, 12)$ the first and the last two vertex pairs), and $A(6)$ in Figure 12. By (11), the endpoint of $\beta(B(d))$ is

$$\beta(0^{d-1}1^d0, 1^{d-1}0^d1) = (0^{d-1}1^d0, 01^{d-2}0^{d-1}11). \quad (21)$$

The duplicate label is $2d-1$, which has been picked up in P . So in the next step of $A(d)$, label $2d-1$ is dropped in Q and label $2d-3$ is picked up, giving the vertex pair

$$(u^*, v^*) = (0^{d-1}1^d0, 01^{d-2}0^{d-2}110). \quad (22)$$

(For the path $\pi(d, 2d)$, label d would be picked up instead at this stage, as stated in the proof of Lemma 5.) This is the edge of $A(d)$ which joins $\beta(B(d))$ to $\alpha(A(d-2))$ in (20).

We are now at the start of $C(d)$ and want to show that this path segment starts with $\alpha(A(d-2))$ with α in (18). Indeed, the starting vertex pair of $C(d)$ is $(u^*, v^*) = \alpha(e_0^{d-2})$. The duplicate label is $2d-3$, which is to be dropped in P in the next step. The subsequent steps are represented by $\alpha(A(d-2))$ since in the lower-dimensional polytope, label 1 is dropped, which is mapped by α to label $2d-3$ of the higher-dimensional polytope, considering α also as an injective map of labels, obtained in the obvious way from (18), namely $\alpha(k) = 2d-2-k$ for $1 \leq k \leq 2d-4$. Essentially, the subsequent steps in $A(d-2)$ map into higher dimension by (18) and by Lemma 6; we only need to check complementarity of the constant positions in higher dimension. In the higher dimension, position 1 with the missing label 1 is zero in both polytopes, consistent with (18). Positions $2d-1$ and $2d$ are also complementary by (18). For positions $2d-3$ and $2d-2$, we have complementarity because $2d-3$ is zero as it is obtained from the position with the missing label 1 in lower dimension. This shows that the initial segment of $C(d)$ is indeed $\alpha(A(d-2))$.

In the last step of $A(d-2)$, label 1 is picked up in Q . So in the last step of $\alpha(A(d-2))$, label $2d-2$ is picked up in Q . Then we are at the vertex pair $(v^*, u^*) = \alpha(e_1^{d-2})$, which is $(01^{d-2}0^{d-2}110, 0^{d-1}1^d0)$ by (22). We have shown that the initial part of $A(d)$ in (20) is $\beta(B(d)) + \alpha(A(d-2))$ and that the starting point and endpoint of $\alpha(A(d-2))$ are (u^*, v^*)

and (v^*, u^*) , respectively. Then the rest of the path $A(d)$ in (20) is obtained by Lemma 4: The next vertex pair, obtained from (v^*, u^*) by dropping label $2d - 2$ in P , is

$$(u', v') = (01^{d-2}0^{d-1}11, 0^{d-1}1^d0), \quad (23)$$

which agrees with Lemma 4 and (21). Thus, the remainder is the path $\beta(B(d))$ backwards but with the bitstrings for P and Q exchanged. However, using the symmetry of $B(d)$ in Lemma 5, this part of the path can be expressed as $\beta'(B(d))$ with a suitably defined mapping β' , similar to β , which exchanges the bitstrings for P and Q . This shows (20). (Figures 7, 10, 11 illustrate the case $d = 6$.)

We now show (17). (For $d = 8$, Figures 12 and 10 show $B(8)$ and $A(6)$, and Figures 7 and 11 together give $C(6)$ as in (16).) The first part of $B(d)$ is indeed $\gamma(A(d - 2))$: Both $B(d)$ and $A(d - 2)$ start by dropping label 1 in P , and the starting point of $B(d)$ is $\gamma(e_0^{d-2})$. Then $B(d)$ proceeds like $\gamma(A(d - 2))$ because of Lemma 6 and since complementarity holds for the constant positions in higher dimension, which is easily checked using (19). Next, by (20),

$$\gamma(A(d - 2)) = \gamma[\beta(B(d - 2)) + \alpha(A(d - 4)) + \beta'(B(d - 2))]. \quad (24)$$

Now consider the starting point (u'', v'') of $\beta'(B(d - 2))$, which is (u', v') given by (23) but with $d - 2$ instead of d . Furthermore, consider the endpoint of $\beta'(B(d - 2))$, that is, the endpoint e_1^{d-2} of $A(d - 2)$. The images of these points under γ are

$$\begin{aligned} \gamma(u'', v'') &= \gamma(01^{d-4}0^{d-3}11, 0^{d-3}1^{d-2}0) = (01^{d-2}0^{d-3}1100, 10^{d-2}1^{d-2}001) \\ \gamma(e_1^{d-2}) &= \gamma(0^{d-2}1^{d-2}, 1^{d-2}0^{d-2}) = (0110^{d-3}1^{d-2}00, 101^{d-2}0^{d-1}1). \end{aligned}$$

This shows that these two vertex pairs $\gamma(u'', v'')$ and $\gamma(e_1^{d-2})$ are mirror images of each other under the symmetry of $B(d)$ described in Lemma 5. This means that the endpoint $\gamma(e_1^{d-2})$ of $\gamma(A(d - 2))$ is already in the second half of $B(d)$. The central part of $B(d)$, given by the last part of $\gamma(A(d - 2))$ in (24), is $\gamma[\beta'(B(d - 2))]$. Therefore, there is a mapping γ' so that

$$B(d) = \gamma[\beta(B(d - 2)) + \alpha(A(d - 4)) + \beta'(B(d - 2))] + \gamma'[\alpha(A(d - 4)) + \beta'(B(d - 2))],$$

because the paths $A(d - 4)$ and $B(d - 2)$ are symmetric and therefore do not have to be written backwards. This representation of $B(d)$ is equivalent to (17) as claimed. \square

Let a_n be the number of vertex pairs of $A(2n)$, which is one more than the length $L(2n, 1)$ of that path. Let b_n and c_n be that number for $B(2n)$ and $C(2n)$, respectively. That is,

$$a_n = L(2n, 1) + 1, \quad b_n = L(2n, 4n) - 2 \quad (n \geq 1). \quad (25)$$

Then the concatenation of paths in (15) implies $a_n = b_n + c_n$, in (16) $c_n = a_{n-1} + b_n$, and in (17) $b_n = a_{n-1} + c_{n-1}$. Moreover, the paths $\pi(2, 1)$ and $\pi(2, 4)$ have length $4 =$

$a_1 - 1 = b_1 + 2$. This shows that the numbers $b_1, c_1, a_1, b_2, c_2, a_2, \dots$ are the Fibonacci numbers $2, 3, 5, 8, 13, 21, \dots$ given by

$$f_0 = 1, \quad f_1 = 2, \quad f_{n+1} = f_n + f_{n-1} \quad (n \geq 1), \quad (26)$$

that is,

$$a_n = f_{3n}, \quad b_n = f_{3n-2} \quad (n \geq 1). \quad (27)$$

So both the lengths of $\pi(d, 1)$ and of $\pi(d, 2d)$ for even $d = 2n = 2, 4, 6, \dots$ are given by every third Fibonacci number (minus one and plus two, respectively). These are the longest paths. They occur several times, since $L(d, 1) = L(d, d)$ and $L(d, d+1) = L(d, d+2) = L(d, 2d-1) = L(d, 2d)$. As shown next, this is due to the symmetry of the Gale evenness condition and of the labelings. Other paths $\pi(d, k)$ are given as concatenations of these paths in lower dimension. They are characterized, for all possible dropped labels k , in the following theorem. The lengths of these paths for $d \leq 14$ are shown in Table 1 in Appendix B.

Theorem 8 *The LH path lengths for any dropped label are characterized by (25), (26), (27), and*

- (a) $L(d, k) = L(d, d+1-k)$ and $L(d, d+k) = L(d, 2d+1-k)$, for $1 \leq k \leq d$;
- (b) $L(d, k) = L(d, k+1)$ for even k when $2 \leq k \leq d-2$, and odd k when $d+1 \leq k \leq 2d-1$;
- (c) $L(d, k) = L(k, 1) + L(d-k, 1)$ for even k and $2 \leq k \leq d-2$;
- (d) $L(d, d+k) = L(k, 2k) + L(d-k+2, 2(d-k+2)) - 4 = b_{k/2} + b_{d/2-k/2+1}$ when k is even and $4 \leq k \leq d-2$.

Proof. Overview. Claim (a) is proved using a cyclic shift by d of each string in $G(d)$ followed by a reversal, which leaves $G(d)$ invariant and is compatible with the labelings l and l' . Claim (b) is proved like Lemma 4. For (c), the paths $A(k)$ and $A(d-k)$ are concatenated with extension mappings similar to (15), (16), (17). A similar argument applies to (d) using the paths $B(k)$ and $B(d-k+2)$. Using (b), cases (c) and (d) cover all possible dropped labels. The range of k in (d) can be restricted because of (a) and (b); by (25), we could have allowed $k = 2$ in (d), but there would be nothing to prove.

For (a), let ψ be defined by $\psi(k) = d - k + 1$ and $\psi(d+k) = 2d - k + 1$ for $k = 1, \dots, d$. This is a cyclic shift by d followed by a reversal of positions, which leaves the set $G(d)$ invariant. Furthermore, ψ commutes with the labelings l and l' of P and Q , so the Lemke–Howson algorithm proceeds in the same manner. That is to say, the vertex pairs on the path $\pi(d, k)$, seen as a pairs of 0-1 strings, are step by step the same when the positions in each string are permuted according to ψ . Under ψ , the first vertex pair e_0 (and similarly the last pair e_1) is mapped to itself, but the positions are changed as above. This means that

under ψ , the path $\pi(d, k)$ is mapped to $\pi(d, d - k + 1)$, so these two paths have the same length. Figures 7 and 8 in Appendix B illustrate the case $d = 4, k = 1$.

To show (b), let $2 \leq k \leq d - 2$. As in Lemma 4, the relabeling l' in (4) applied to both P and Q shows that $\pi(d, k)$ corresponds to the path $\pi(d, k + 1)$ backwards, so these paths have the same length.

Claim (c) requires more work. For any paths A and B on $G(n) \times G(n)$, considered as sequences of vertex pairs, let AB denote the path A joined to the path B , where the endpoint of A is equal to the starting point of B . The length (number of edges) of AB is the sum of the lengths of A and B .

We prove the the following statement, which clearly implies (c): Let k be even and $2 \leq k \leq d - 2$. Then

$$\pi(d, k) = \alpha(A(k)) \beta(A(d - k))$$

with $\alpha : G(k) \times G(k) \rightarrow G(d) \times G(d) = P \times Q$,

$$\alpha(u, v) = (u_k \cdots u_1 1^{d-k} 0^{d-k} u_{2k} \cdots u_{k+1}, v_k \cdots v_1 0^{d-k} 1^{d-k} v_{2k} \cdots v_{k+1}), \quad (28)$$

and $\beta : G(d - k) \times G(d - k) \rightarrow G(d) \times G(d)$,

$$\beta(u, v) = (0^k u 1^k, 1^k v 0^k).$$

The starting point of $\pi(d, k)$ is e_0^d . As required, $\alpha(e_0^k) = e_0^d$. In the first step of $\pi(d, k)$, label k is dropped in P . Position 1 of the lower dimensional polytopes, given by the bit u_1 and v_1 in (28), is mapped to position k in both P and Q in the higher dimension. In P , position k has label k , which is missing in $\pi(d, k)$. This missing label in the higher dimensional polytope P corresponds to the missing label 1 in the lower dimension. Because α preserves the adjacency of labels cyclically, and since by Lemma 6 the path $A(k)$ does not wrap around, the first $L(k, 1)$ steps of $\pi(d, k)$ proceed according to $\alpha(A(k))$; all we need to check is the complementarity of the positions of the higher dimensional polytope which are constant according to α . Complementarity of positions $k + 2, \dots, 2d - k$ is immediate. Position k and $k + 1$ in P and Q , respectively, correspond to the missing label k of $\pi(d, k)$ and are thus both zero throughout. Finally, position $k + 1$ in P , with label $k + 1$, is 1 according to (28), and is complementary since position k in Q with label $k + 1$ is 0 throughout, as it corresponds to the missing label in the lower dimensional polytope. At the end of the first $L(k, 1)$ steps, the vertex pair $\alpha(e_1^k)$ is reached, where

$$\alpha(e_1^k) = (0^k 1^{d-k} 0^{d-k} 1^k, 1^k 0^{d-k} 1^{d-k} 0^k) = \beta(e_0^{d-k}),$$

and so this is also the starting point of $\beta(A(d - k))$, as required. In a similar way as before, one can see that this is the second part of $\pi(d, k)$, which ends in $\beta(e_1^{d-k}) = e_1^d$. This shows (c). Figures 13, 7 and 10 in Appendix B illustrate the case $d = 10, k = 4$.

Let k be even and $4 \leq k \leq d - 2$. To show (d), we construct suitable mappings γ and δ so that

$$\pi(d, d + k) = e_0^d + \gamma(B(k)) \delta(B(d - k + 2)) + e_1^d. \quad (29)$$

These mappings are $\gamma: G(k) \times G(k) \rightarrow G(d) \times G(d) = P \times Q$,

$$\gamma(u, v) = (u_1 1^{d-k} u_2 \cdots u_{2k} 0^{d-k}, v_1 0^{d-k} v_2 \cdots v_{2k} 1^{d-k}), \quad (30)$$

and $\delta: G(d - k + 2) \times G(d - k + 2) \rightarrow G(d) \times G(d) = P \times Q$,

$$\begin{aligned} \delta(u, v) = & (v_{d-k+2} \cdots v_{2d-2k+2} 0^{k-2} 1^k 0 v_2 \cdots v_{d-k+1}, \\ & u_{d-k+2} \cdots u_{2d-2k+2} 1^{k-2} u_{2d-2k+3} 0^{k-1} u_1 \cdots u_{d-k+1}). \end{aligned}$$

It can be verified that these mappings preserve the adjacency of relevant labels and complementarity. The path $\pi(d, d + k)$ starts as follows: after dropping from e_0^d label $d + k$ in Q , which is in position $d + k - 1$, the vertex pair is $(1^d 0^d, 10^{d-1} 1^{k-2} 0 1^{d-k+1})$, which is equal to $\gamma(u^1, v^1)$ for the first vertex pair (u^1, v^1) of $B(k)$ as in (10) (with k instead of n), also with duplicate label 1. The path continues as described in (30) since the bits v_1 and v_{2k} in (30) stay constant according to (7). The last vertex pair of $\gamma(B(k))$ is, by (11) and (30), equal to

$$\gamma(0^{k-1} 1^k 0, 1^{k-1} 0^k 1) = (0 1^{d-k} 0^{k-2} 1^k 0^{d-k+1}, 1 0^{d-k} 1^{k-2} 0^k 1^{d-k+1}).$$

This is equal to $\delta(1^{d-k+2} 0^{d-k+2}, 10^{d-k+1} 1^{d-k} 0 1)$, which is δ applied to the first vertex pair of $B(d - k + 2)$, using (10) with $d - k + 2$ instead of d . The duplicate label $d + k - 1$ to be dropped, in position $d - k$ in Q , is the image of the bit u_1 under δ , where this bit u_1 is dropped in $B(d - k - 2)$ by Lemma 5. Note that $\delta(u, v)$ ignores the bits $u_{2d-2k+4}$, and v_1 , $v_{2d-2k+3}$, and $v_{2d-2k+4}$, which are constant throughout $B(d - k + 2)$ by Lemma 5. The last vertex pair of $\delta(B(d - k + 2))$ is

$$\delta(0^{d-k+1} 1^{d-k+2} 0, 1^{d-k+1} 0^{d-k+2} 1) = (0^{d-1} 1^k 0 1^{d-k}, 1^d 0^d)$$

with duplicate label d , which has just been picked up in Q as the image of bit $u_{2d-2k+3}$ under δ . When this label d is dropped in P , the endpoint e_1^d is reached, which terminates the path $\pi(d, d + k)$. This completes the proof of (29). According to (29), the length of $\pi(d, d + k)$ is the sum of the lengths of $B(k)$ and of $B(d - k + 2)$ plus two (for the first edge from e_0^d and last edge to e_1^d), which shows (d). The case $d = 10$, $k = 4$ is illustrated by Figures 14, 9 and 12 in Appendix B. \square

It is easy to see that the shortest path lengths are obtained as follows: If d is divisible by four, that is, $d/2$ is even, then the shortest path length occurs when dropping label $d/2$, and is given by $L(d, d/2) = 2a_{d/2} - 2$ according to Theorem 8(c). If $d/2$ is odd, then the shortest path length occurs for dropped label $3d/2$, where $L(d, 3d/2) = L(d, 3d/2 + 1) = 2b_{d/2+1}$ by Theorem 8(b) and (d). When $d/2$ is even, the path when dropping label $3d/2$ is only

two steps longer than when dropping label $d/2$ since then $L(d, 3d/2) = b_{d/2} + b_{d/2+1} = b_{d/2} + a_{d/2} + c_{d/2} = 2a_{d/2}$. Therefore, the shortest path results essentially when dropping label $3d/2$.

The Fibonacci numbers are given by the well-known explicit expression

$$f_n = K\theta^n + \bar{K}\bar{\theta}^n, \quad \theta, \bar{\theta} = 0.5 \pm 0.5\sqrt{5}, \quad K, \bar{K} = 0.5 \pm 0.3\sqrt{5},$$

where $\theta = 1.618\dots$ is the Golden Ratio and $K = 1.170\dots$. Then f_n is $K\theta^n$ rounded to the nearest integer since $\bar{K}\bar{\theta}^n$ is less than 0.5 and at any rate exponentially small. By Theorem 8(d), the sequence of shortest LH path lengths $L(2n, 3n)$ for $n = d/2 = 1, 2, 3, \dots$ is 4, 10, 16, 42, 68, 178, \dots , which is the sequence of Fibonacci numbers (multiplied by two) with every third number omitted. These shortest lengths grow with the square root of the longest lengths, which is still exponential.

Corollary 9 *There are $d \times d$ games, for even d , where each LH path has length $\Omega(\theta^{3d/4})$.*

A similar construction using a similar labeling to (4) is possible for odd d , but there the path lengths are less symmetric than those in Theorem 8 for even d . We do not need this since it is trivial to obtain an odd-dimensional game from the next lower even dimension by adding a strictly dominated strategy for each player.

4 Conclusions and open questions

In this paper, we have presented a construction of $d \times d$ games with a unique equilibrium which is found by the LH algorithm using an exponential number of steps, for any dropped label. Unfortunately, the equilibrium is easily *guessed* since it has full support, so our games are not “hard to solve” by other methods. This holds because the permutation U' in (4) gives the artificial equilibrium e_0 in (5), but also its complement e_1 in Lemma 3, which is the completely mixed equilibrium (see also condition (35) in Appendix A below).

It is an open problem to find a suitable construction, perhaps extending ours, where the game does not have a completely mixed equilibrium, but where all equilibria are still found using an exponential number of LH steps. If the support of an equilibrium is $d/2$, say, then it could be *hidden* by randomly permuting the players’ strategies, so that it would take an exponential number, like some fraction of $\binom{d}{d/2}$, of guesses to find that support.

Lemke’s algorithm [15] is closely related to the LH algorithm. It solves an LCP [3] by introducing an auxiliary vector and variable into the system, and pivots according to the same complementary pivoting rule as LH, until the extra variable becomes zero, thereby computing an equilibrium. This method can be given an interpretation in game-theoretic terms [32]. Its extra flexibility given by the choice of numerical values in the auxiliary vector, rather than just of finitely many starting edges as in the LH algorithm, deserves further study.

It may happen that the equilibrium in our construction is found quickly from a suitable, easily found starting point. A possibly related study is [18].

Using the games presented here as test cases raises the additional problem that the moment curve gives rise to notoriously ill-conditioned matrices. As a consequence, numerical problems arise when the pivoting steps are implemented using floating-point arithmetic. These numerical problems may possibly be avoided by using points on the trigonometric moment curve, as mentioned in Appendix A. As mentioned there, an open problem is the required numerical accuracy of these points.

Appendix A: Generating game matrices

In this appendix, we describe how to obtain games from representations of cyclic polytopes. In principle, this has already been described in Proposition 2.1 of [30] for general polytopes. In our case, as well as in the construction of games with a large number of equilibria in [30], the polytopes are dual cyclic polytopes in dimension d with $2d$ facets, with a labeling of the facets of each polytope that has a certain structure. This structure allows a further simplification: Only one of dual cyclic polytopes, say P , has to be brought into the form (3), where d of the inequalities are simply nonnegativities and the other d inequalities define the payoff matrix of one player, here B . The other polytope, and thus the payoff matrix A of the other player, is then simply obtained by a suitable permutation of the rows and columns of B^\top . We first explain this construction, which is summarized in Proposition 10 below.

Secondly, we apply this to a representation of cyclic polytopes in dimension $d = 6$ derived from the so-called trigonometric moment curve. In this low dimension, the coordinates on that curve can be approximated by small integers, which gives rise to small game matrix entries.

As indicated at the beginning of Section 3, a standard way of obtaining a cyclic polytope in dimension d with $2d$ vertices is, first, to consider $2d$ points $\mu(t_i)$ on the moment curve $\mu: t \mapsto (t, t^2, \dots, t^d)^\top$ for $1 \leq i \leq 2d$. Suppose that $t_1 < t_2 < \dots < t_{2d}$. Then the vertices of that polytope are characterized by the 0-1 strings fulfilling the Gale evenness condition. The *polar* (or dual) polytope [37][8] is obtained by translating the polytope so that it has the origin $\mathbf{0}$ in its interior, for example by subtracting the arithmetic mean $\bar{\mu}$ of the points $\mu(t_i)$ from each such point. The resulting vectors $c_i = \mu(t_i) - \bar{\mu}$ then define the polar cyclic polytope

$$P' = \{z \in \mathbb{R}^d \mid c_i^\top z \leq 1, 1 \leq i \leq 2d\}. \quad (31)$$

As described in [30, p. 560], if $P' = \{z \in \mathbb{R}^d \mid Cz \leq \mathbf{1}, Dz \leq \mathbf{1}\}$ with $d \times d$ matrices C and D , then an affine transformation of P' is given by

$$P = \{x \in \mathbb{R}^d \mid x \geq \mathbf{0}, -DC^{-1}x \leq r\}, \quad r = \mathbf{1} - DC^{-1}\mathbf{1}. \quad (32)$$

Since $\mathbf{0}$ is a vertex of the simple polytope P , the vector r is positive, and the second d inequalities in (32) can be re-normalized so that the right hand side is one. With the diagonal matrix S with entries $s_{ii} = 1/r_i$ with r as in (32), and $s_{ij} = 0$ for $i \neq j$, we can rewrite (32) as

$$P = \{x \in \mathbb{R}^d \mid x \geq \mathbf{0}, -SDC^{-1}x \leq \mathbf{1}\}. \quad (33)$$

Affine transformations leave the combinatorial structure (that is, the face incidences) of a polytope unchanged, so P is a cyclic polytope with facets characterized by Gale evenness strings. These Gale evenness strings refer to the $2d$ inequalities defining P according to the ordering in (33), that is, $x_1 \geq 0$ being the first inequality obtained from the first point $\mu(t_1)$ on the moment curve, $x_2 \geq 0$ corresponding to $\mu(t_2)$, and so on.

Consider the polytope \overline{Q} defined by

$$\overline{Q} = \{\overline{y} \in \mathbb{R}^d \mid -SDC^{-1}\overline{y} \leq \mathbf{1}, \overline{y} \geq \mathbf{0}\}, \quad (34)$$

which is identical to P in (33) except that the first and last d inequalities are interchanged.

In a dual cyclic polytope like P in (33), each inequality defines a facet (obtained by converting the inequality to an equality). We say that a facet of P has *label* k (for $k = 1, \dots, 2d$) if it corresponds to the k th inequality in the description of the polytope in (33). Similarly, a facet of \overline{Q} has label k if it corresponds to the k th inequality in (34). If P and \overline{Q} in (33) and (34) are the polytopes P and Q in (3), they define a symmetric bimatrix game with payoff matrices (A, B) where $B^\top = A = -SDC^{-1}$.

The vertices of a dual cyclic polytope are given by the sets of d facets each vertex lies on. Encoded as bitstrings, these sets are characterized by the Gale evenness condition explained at the beginning of Section 3, with $G(d)$ as the set of these bitstrings. We assume d is even. Then the Gale evenness condition is preserved by a cyclic rotation of the bitstrings, in particular by d positions, as used in the definition (34) of \overline{Q} . Thus, the vertices of both P and \overline{Q} correspond to the Gale evenness strings in the set $G(d)$. A bitstring u in $G(d)$ defines the vertex x of P obtained by converting the k th inequality in (33) to an equality whenever $u_k = 1$, for $k = 1, \dots, 2d$. In the same manner, v in $G(d)$ defines the vertex y of \overline{Q} where the k th inequality in (34) is binding whenever $v_k = 1$.

In our construction, as well as in [30], the polytopes P and Q in (3) are dual cyclic polytopes but the games are not symmetric, because the facets of Q are not labeled in their original order. Instead, a certain permutation λ is used to obtain Q from \overline{Q} , by letting the k th facet of \overline{Q} in the description (34) have label $\lambda(k)$ in Q , for $k = 1, \dots, 2d$. In our construction, we used the permutation $\lambda = l'$ defined in (4). With $\lambda(S) = \{\lambda(k) \mid k \in S\}$ for $S \subseteq 1, \dots, 2d$, this permutation has the property

$$\lambda(\{1, \dots, d\}) = \{1, \dots, d\} \quad (35)$$

(and thus $\lambda(\{d+1, \dots, 2d\}) = \{d+1, \dots, 2d\}$). This condition implies that the pair $(u, v) = e_0$ in (5) is complementary. The corresponding vertex pair of $P \times Q$ is the artificial equilibrium $(\mathbf{0}, \mathbf{0})$. (Property (35) also implies that e_1 in Lemma 3 is complementary, which defines the completely mixed equilibrium.)

The following proposition describes the construction of a bimatrix game (A, B) using P in (33), and Q defined by \overline{Q} in (34) with labels given by a permutation λ fulfilling (35). The proposition shows how to obtain A from B^\top by permuting rows and columns suitably.

Proposition 10 *Consider a pair of dual cyclic polytopes in dimension d with $2d$ facets, with each vertex set represented by the set of Gale evenness strings $G(d)$. Let λ be a permutation of $\{1, \dots, 2d\}$ that fulfills (35). For $k = 1, \dots, 2d$, a vertex u in $G(d)$ of the first polytope*

has the labels k where $u_k = 1$, a vertex v in $G(d)$ of the second polytope has the labels $\lambda(k)$ where $v_k = 1$. A vertex pair (u, v) is complementary if it has all labels. Then a $d \times d$ bimatrix game (A, B) with Nash equilibria corresponding to complementary vertex pairs, where the artificial equilibrium corresponds to e_0 in (5), is given by $B^\top = -SDC^{-1}$ as in (33), using a representation (31) of a dual cyclic polytope consistent with the Gale evenness ordering. The matrix entries $a(i, j)$ of A are obtained from the matrix entries $b(i, j)$ of B by

$$a(\lambda(i), \lambda(j + d) - d) = b(j, i) \quad (1 \leq i, j \leq d). \quad (36)$$

Proof. For the characterization of equilibria, the combinatorial structure of the dual cyclic polytopes suffices, as given by the Gale evenness strings $G(d)$. Both P in (33) and \bar{Q} in (34) are representations of such polytopes. By assumption, for $k = 1, \dots, 2d$, the k th inequality of (33) has label k , and the k th inequality of (34) has label $\lambda(k)$.

Let $B^\top = -SDC^{-1}$, define the matrix A with entries $a(i, j)$ by (36), and let

$$Q = \{y \in \mathbb{R}^d \mid Ay \leq \mathbf{1}, y \geq \mathbf{0}\}. \quad (37)$$

The polytopes P and Q in (33), (37) correspond to the bimatrix game (A, B) , as in (3). The facets of Q have labels in the order of the inequalities in (37). It suffices to show that these labels k correspond to the labels $\lambda(k)$ of \bar{Q} stated above.

In detail, the inequalities in (34) are

$$\bar{Q} = \left\{ \bar{y} \in \mathbb{R}^d \mid \sum_{j=1}^d b(j, i) \bar{y}_j \leq 1 \quad (1 \leq i \leq d), \right. \\ \left. \bar{y}_j \geq 0 \quad (1 \leq j \leq d) \right\}. \quad (38)$$

For $1 \leq j \leq d$, the $(d + j)$ th inequality in (38) has label $\lambda(d + j)$. Hence, it should appear as the $\lambda(d + j)$ th inequality in (37), which by (35) is the inequality $y_{\lambda(d+j)-d} \geq 0$. This is achieved by the correspondence between y in (37) and \bar{y} in (38) given by $y_{\lambda(d+j)-d} = \bar{y}_j$.

The i th of the first d inequalities in (38), for $1 \leq i \leq d$, has label $\lambda(i)$. It should appear as the $\lambda(i)$ th inequality in (37). That inequality has the form

$$\sum_{l=1}^d a(\lambda(i), l) y_l \leq 1,$$

which by (35) can be rewritten as

$$\sum_{j=1}^d a(\lambda(i), \lambda(d + j) - d) y_{\lambda(d+j)-d} \leq 1,$$

which by (36) is the i th inequality of (38) as claimed. \square

Consider the following 6×6 bimatrix game (A, B) with

$$A = \begin{bmatrix} -180 & 72 & -333 & 297 & -153 & 270 \\ -30 & 17 & -33 & 42 & -3 & 20 \\ -81 & 36 & -126 & 126 & -36 & 90 \\ 90 & -36 & 126 & -126 & 36 & -81 \\ 20 & -3 & 42 & -33 & 17 & -30 \\ 270 & -153 & 297 & -333 & 72 & -180 \end{bmatrix}, \quad B = \begin{bmatrix} 72 & 36 & 17 & -3 & -36 & -153 \\ -180 & -81 & -30 & 20 & 90 & 270 \\ 297 & 126 & 42 & -33 & -126 & -333 \\ -333 & -126 & -33 & 42 & 126 & 297 \\ 270 & 90 & 20 & -30 & -81 & -180 \\ -153 & -36 & -3 & 17 & 36 & 72 \end{bmatrix}.$$

The matrix A is obtained from B via (36) with $\lambda = \iota'$ in (4). The matrix B is obtained as in Proposition 10. The underlying representation (31), however, is not based on points $(t, t^2, t^3, t^4, t^5, t^6)$ of the moment curve, but on points $\nu(t)$ of the *trigonometric* moment curve, $\nu(t) = (\cos t, \sin t, \cos 2t, \sin 2t, \cos 3t, \sin 3t)$. These points also give rise to cyclic polytopes [37, p. 75f] [8, p. 67]. For $t = i\pi/6$ for $i = 1, \dots, 12$, the first pair of coordinates of $\nu(t)$ denote the vertices of a regular 12-gon, the second pair those of a regular hexagon, used twice, and the third pair those of a square, used three times. The origin is in the interior of the convex hull of these vertices, so the polytope does not have to be translated to obtain its polar. The combinatorial structure is preserved by choosing suitable integer coordinates near the points on the circle, which are shown in Figure 4; payoffs have been multiplied by 18 to obtain integers. (The square is represented perfectly; choosing as its vertices instead the points $(1, 0), (0, 1), (-1, 0), (0, -1)$, say, would not change B as the affine transformation that produces (32) always gives the unit vectors as the normal vectors of the first d facets of P .) It is an open problem to find suitable approximations with small integers in higher dimensions that preserve the combinatorial structure.

The bimatrix game (A', B) with

$$A' = \begin{bmatrix} -81 & 36 & -126 & 126 & -36 & 90 \\ -180 & 72 & -333 & 297 & -153 & 270 \\ 20 & -3 & 42 & -33 & 17 & -30 \\ -30 & 17 & -33 & 42 & -3 & 20 \\ 270 & -153 & 297 & -333 & 72 & -180 \\ 90 & -36 & 126 & -126 & 36 & -81 \end{bmatrix}$$

is obtained from the permutation $\lambda(k) = k - (-1)^k$ for $1 \leq k \leq 12$ in (36). This permutation is used in [30], and the game (A', B) has 75 equilibria.




FIGURE 4: Approximation of points on the trigonometric moment curve by small integers. The circled numbers refer to the labels $i = 1, \dots, 12$ of the vertices, which become facets in the dual cyclic polytope P .

Appendix B: Examples of path lengths and paths

The following table and figures show the empirical evidence leading to our main Theorems 7 and 8. Table 1 shows the path lengths and their exponential growth, and that the lengths of the short paths $\pi(d, 3d/2)$ for $d = 2, 4, 6, \dots$ are given by the Fibonacci numbers times two, with every third Fibonacci number omitted. Figures 7, 10, 11, and 12 illustrate Theorem 7, and Figures 8, 13, and 14 show cases of Theorem 8.

label	dimension						
	2	4	6	8	10	12	14
1	4	20	88	376	1596	6764	28656
2	4	8	24	92	380	1600	6768
3	4	8	24	92	380	1600	6768
4	4	20	24	40	108	396	1616
5		10	24	40	108	396	1616
6		10	88	92	108	176	464
7		10	36	92	108	176	464
8		10	36	376	380	396	464
9			16	146	380	396	464
10			16	146	1596	1600	1616
11			36	42	612	1600	1616
12			36	42	612	6764	6768
13				42	152	2586	6768
14				42	152	2586	28656
15				146	68	618	10948
16				146	68	618	10948
17					152	178	2592
18					152	178	2592
19					612	178	644
20					612	178	644
21						618	288
22						618	288
23						2586	644
24						2586	644
25							2592
26							2592
27							10948
28							10948

TABLE 1: Path lengths for different dropped labels.

In the following figures, each row displays two pivoting steps, one in P and one in Q, so the number of the last row is to be multiplied by two to obtain the path length.




FIGURE 5: $\pi(2, 1)$.




FIGURE 6: $\pi(2, 4)$




FIGURE 7: $\pi(4, 1)$




FIGURE 8: $\pi(4, 4)$

	1	2	3	4	5	6	7	8		1	3	2	4	6	5	8	7
0	1	1	1	1	0	0	0	0	0	0	0	0	0	1	1	1	1
1	0	1	1	1	1	0	0	0	0	1	0	0	0	1	1	0	1
2	0	1	1	0	1	1	0	0	0	1	0	0	1	1	0	0	1
3	0	0	1	1	1	1	0	0	0	1	0	1	1	0	0	0	1
4	0	0	0	1	1	1	1	0	0	1	1	1	0	0	0	0	1
5	0	0	0	0	1	1	1	1	1	1	1	1	1	0	0	0	0

FIGURE 9: $\pi(4, 8)$

	1	2	3	4	5	6	7	8	9	10	11	12		1	3	2	5	4	6	8	7	10	9	12	11
0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
1	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	1	1	0	1	1	1	1	1
2	0	1	1	1	1	0	1	1	0	0	0	0	0	0	0	0	0	1	1	0	0	1	1	1	1
3	0	1	1	0	1	1	1	1	0	0	0	0	0	0	0	0	0	1	1	0	0	0	1	1	1
4	0	1	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	1	1	0	0	1	1	0	1
5	0	1	1	0	1	1	0	1	1	0	0	0	0	0	0	0	0	0	1	1	0	1	1	0	1
6	0	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	0	0	1	1	1	1	0	1	1
7	0	1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	0	1	1
8	0	1	1	0	1	1	0	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	0	1	1
9	0	1	1	0	0	1	1	0	1	1	0	0	0	0	0	0	0	1	1	1	1	0	0	1	1
10	0	1	1	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	1	1
11	0	0	1	1	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	0	1	0	0	0	1
12	0	0	1	1	0	1	1	0	1	1	0	0	0	0	0	0	0	1	1	0	0	1	0	0	1
13	0	0	1	1	1	1	0	0	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	1	1
14	0	0	0	1	1	1	1	0	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	1	1
15	0	0	0	1	1	0	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	1	1
16	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	0	1	1
17	0	0	0	0	0	1	1	1	1	1	1	1	0	0	0	0	0	1	1	0	0	0	1	1	0
18	0	0	0	0	1	1	1	1	0	1	1	0	0	0	0	0	0	1	1	0	0	0	1	1	0
19	0	0	0	1	1	0	1	1	0	1	1	0	0	0	0	0	0	1	1	0	0	1	1	0	0
20	0	0	0	1	1	1	1	0	0	1	1	0	0	0	0	0	0	1	1	0	0	1	1	0	0
21	0	0	1	1	1	1	0	0	0	1	1	0	0	0	0	0	0	1	1	0	1	0	1	1	0
22	0	0	1	1	0	1	1	0	0	1	1	0	0	0	0	0	0	1	1	0	0	1	1	0	0
23	0	0	1	1	0	0	1	1	0	1	1	0	0	0	0	0	0	1	1	0	0	0	1	1	0
24	0	1	1	0	0	0	1	1	0	1	1	0	0	0	0	0	0	1	1	1	1	0	0	1	1
25	0	1	1	0	0	1	1	0	0	1	1	0	0	0	0	0	0	1	1	0	1	0	1	1	0
26	0	1	1	0	1	1	0	0	0	1	1	0	0	0	0	0	0	1	1	1	1	0	1	1	0
27	0	1	1	1	1	0	0	0	0	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	0
28	0	1	1	1	1	0	0	0	0	0	1	1	0	0	0	0	0	1	1	1	1	1	1	0	0
29	0	1	1	0	1	1	0	0	0	0	1	1	0	0	0	0	0	1	1	0	1	1	1	0	0
30	0	1	1	0	0	1	1	0	0	0	1	1	0	0	0	0	0	1	1	1	1	0	1	1	0
31	0	1	1	0	0	0	1	1	0	0	1	1	0	0	0	0	0	1	1	0	0	1	1	0	0
32	0	0	1	1	0	0	1	1	0	0	1	1	0	0	0	0	0	1	1	0	1	1	0	0	0
33	0	0	1	1	0	1	1	0	0	0	1	1	0	0	0	0	0	1	1	1	1	1	0	0	0
34	0	0	1	1	1	1	0	0	0	0	1	1	0	0	0	0	0	1	1	1	1	1	0	0	0
35	0	0	0	1	1	1	1	0	0	0	0	1	1	0	0	0	0	1	1	0	1	1	0	0	0
36	0	0	0	1	1	0	1	1	0	0	1	1	0	0	0	0	0	1	1	0	0	1	1	0	0
37	0	0	0	0	1	1	1	1	0	0	1	1	0	0	0	0	0	1	1	0	0	1	1	0	0
38	0	0	0	0	0	1	1	1	1	0	1	1	0	0	0	0	0	1	1	0	1	1	0	0	0
39	0	0	0	0	1	1	0	1	1	0	1	1	0	0	0	0	0	1	1	0	1	1	0	0	0
40	0	0	0	1	1	0	0	1	1	0	1	1	0	0	0	0	0	1	1	1	1	1	0	0	0
41	0	0	0	1	1	0	0	0	1	1	1	1	0	0	0	0	0	1	1	1	1	1	0	0	0
42	0	0	0	0	1	1	0	0	1	1	1	1	0	0	0	0	0	1	1	1	1	0	0	0	0
43	0	0	0	0	0	1	1	0	1	1	1	1	0	0	0	0	0	1	1	1	1	0	0	0	0
44	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0

FIGURE 10: $\pi(6, 1)$





FIGURE 11: $\pi(6, 12)$



References

- [1] Audet, C., P. Hansen, B. Jaumard, and G. Savard (2001), Enumeration of all extreme equilibria of bimatrix games. *SIAM Journal on Scientific Computing* **23**, 323–338.
- [2] Borodin, A. and R. El-Yaniv (1998), *Online Computation and Competitive Analysis*. Cambridge Univ. Press, Cambridge.
- [3] Cottle, R. W., J.-S. Pang, and R. E. Stone (1992), *The Linear Complementarity Problem*. Academic Press, San Diego.
- [4] Dantzig, G. B. (1963), *Linear Programming and Extensions*. Princeton University Press, Princeton.
- [5] Gale, D. (1963), Neighborly and cyclic polytopes. In: *Convexity*, Proc. Symposia in Pure Math., Vol. 7, ed. V. Klee, American Math. Soc., Providence, Rhode Island, 225–232.
- [6] Garcia, C. B., and W. I. Zangwill (1981), *Pathways to Solutions, Fixed Points, and Equilibria*. Prentice-Hall, Englewood Cliffs.
- [7] Gilboa, I. and E. Zemel (1989), Nash and correlated equilibria: some complexity considerations. *Games and Economic Behavior* **1**, 80–93.
- [8] Grünbaum, B. (2003), *Convex Polytopes, 2nd ed.* Springer, New York.
- [9] Fathi, Y. (1979), Computational complexity of LCPs associated with positive definite symmetric matrices. *Mathematical Programming* **17**, 335–344.
- [10] Jansen, M. J. M. (1981), Maximal Nash subsets for bimatrix games. *Naval Research Logistics Quarterly* **28**, 147–152.
- [11] Kalai, G., and D. J. Kleitman (1992), A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bull. Amer. Math. Soc.* **26**, 315–316.
- [12] Klee, V., and P. Kleinschmidt (1987), The d-step conjecture and its relatives. *Math. of Oper. Res.* **12**, 718–755.
- [13] Klee, V., and G. J. Minty (1972), How good is the simplex algorithm? In: *Inequalities, III*, Proc. Third Sympos., UCLA, 1969, ed. O. Shisha, Academic Press, New York, 159–175.
- [14] Lemke, C. E. and J. T. Howson, Jr. (1964), Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* **12**, 413–423.
- [15] Lemke, C. E. (1965), Bimatrix equilibrium points and mathematical programming. *Management Science* **11**, 681–689.
- [16] Lipton, R. J., E. Markakis, and A. Mehta (2003), Playing large games using simple strategies. *Proc. 4th ACM conf. Electronic Commerce*, San Diego, 36–41.
- [17] McKelvey, R. D., and A. McLennan (1996), Computation of equilibria in finite games. In: *Handbook of Computational Economics, Vol. I*, eds. H. M. Amman, D. A. Kendrick, and J. Rust, Elsevier, Amsterdam, 87–142.

- [18] Megiddo, N. (1986), On the expected number of linear complementarity cones intersected by random and semi-random rays. *Mathematical Programming* **35**, 225–235.
- [19] Megiddo, N., and C. H. Papadimitriou (1991), On total functions, existence theorems and computational complexity (Note). *Theoretical Computer Science* **81**, 317–324.
- [20] Morris, W. D., Jr. (1994), Lemke paths on simple polytopes. *Math. of Oper. Res.* **19**, 780–789.
- [21] Murty, K. G. (1978), Computational complexity of complementary pivot methods. *Mathematical Programming Study 7: Complementary and Fixed Point Problems*, 61–73.
- [22] Nash, J. F. (1951), Non-cooperative games. *Annals of Mathematics* **54**, 286–295.
- [23] Nisan, N., and A. Ronen (2001), Algorithmic mechanism design. *Games and Economic Behavior* **35**, 166–196. [Extended Abstract in *Proc. 31st STOC* (1999), 129–140.]
- [24] Papadimitriou, C. H. (1994), On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences* **48**, 498–532.
- [25] Papadimitriou, C. H. (1994), *Computational Complexity*. Addison-Wesley, Reading, Mass.
- [26] Papadimitriou, C. H. (2001), Algorithms, games, and the internet. In: *Proc. 33rd STOC*, 749–753.
- [27] Rosenmüller, J. (1971), On a generalization of the Lemke–Howson algorithm to noncooperative N-person games. *SIAM J. Appl. Math.* **21**, 73–79.
- [28] Roughgarden, T. (2002), *Selfish Routing*. PhD thesis, Cornell University.
- [29] Shapley, L. S. (1974), A note on the Lemke–Howson algorithm. *Mathematical Programming Study 1: Pivoting and Extensions*, 175–189.
- [30] von Stengel, B. (1999), New maximal numbers of equilibria in bimatrix games. *Discrete and Computational Geometry* **21**, 557–568.
- [31] von Stengel, B. (2002), Computing equilibria for two-person games. Chapter 45, *Handbook of Game Theory, Vol. 3*, eds. R. J. Aumann and S. Hart, North-Holland, Amsterdam, 1723–1759.
- [32] von Stengel, B., A. H. van den Elzen, and A. J. J. Talman (2002), Computing normal form perfect equilibria for extensive two-person games. *Econometrica* **70**, 693–715.
- [33] Vorob’ev, N. N. (1958), Equilibrium points in bimatrix games. *Theory of Probability and its Applications* **3**, 297–309.
- [34] Wilson, R. (1971), Computing equilibria of N-person games. *SIAM J. Appl. Math.* **21**, 80–87.
- [35] Winkels, H.-M. (1979), An algorithm to determine all equilibrium points of a bimatrix game. In: *Game Theory and Mathematical Economics*, eds. O. Moeschlin and D. Pallaschke, North-Holland, Amsterdam, 137–148.
- [36] Yao, A. C. (1977), Probabilistic computation: towards a unified measure of complexity. In: *Proc. 18th FOCS*, 222–227, 1977.
- [37] Ziegler, G. M. (1995), *Lectures on Polytopes*. Springer, New York.