

Unit Vector Games

Rahul Savani*

Bernhard von Stengel[†]

February 15, 2016

Published in: *International Journal of Economic Theory* 12 (2016), 7–27. doi: 10.1111/ijet.12077

Abstract

McLennan and Tourky (2010) showed that “imitation games” provide a new view of the computation of Nash equilibria of bimatrix games with the Lemke–Howson algorithm. In an imitation game, the payoff matrix of one of the players is the identity matrix. We study the more general “unit vector games”, which are already known, where the payoff matrix of one player is composed of unit vectors. Our main application is a simplification of the construction by Savani and von Stengel (2006) of bimatrix games where two basic equilibrium-finding algorithms take exponentially many steps: the Lemke–Howson algorithm, and support enumeration.

Keywords: bimatrix game, Nash equilibrium computation, imitation game, Lemke–Howson algorithm, unit vector game

1 Introduction

A bimatrix game is a two-player game in strategic form. The Nash equilibria of a bimatrix game correspond to pairs of vertices of two polyhedra derived from the payoff matrices. These vertex pairs have to be “completely labeled”, which expresses the equilibrium condition that every pure strategy of a player (represented by a “label”) is either a best response to the other player’s mixed strategy or played with probability zero.

This polyhedral view gives rise to algorithms that compute a Nash equilibrium. A classical method is the algorithm by Lemke and Howson (1964) which follows a path of “almost completely labeled” polytope edges that terminates at Nash equilibrium. The Lemke–Howson (LH) algorithm has been one inspiration for the complexity class PPAD defined by Papadimitriou (1994) of computational problems defined by such path-following arguments, which includes more general equilibrium problems such as the computation of approximate Brouwer fixed

*Department of Computer Science, University of Liverpool, Liverpool L69 3BX, United Kingdom. Email: rahul.savani@liverpool.ac.uk

[†]Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom. Email: stengel@nash.lse.ac.uk

points. An important result proved by Chen and Deng (2006) and Daskalakis, Goldberg, and Papadimitriou (2009) states that every problem in the class PPAD can be reduced to finding a Nash equilibrium of a bimatrix game, which makes this problem “PPAD-complete”. (The problem of finding the Nash equilibrium at the end of a *specific* path is a much harder, namely PSPACE-complete, see Goldberg, Papadimitriou, and Savani 2013.)

If an algorithm takes exponentially many steps (measured in the size of its input) for certain problem instances, these are considered “hard” instances for the algorithm. Savani and von Stengel (2006) constructed bimatrix games that are hard instances for the LH algorithm. Their construction uses “dual cyclic polytopes” which have a well-known vertex structure for any dimension and number of linear inequalities. Morris (1994) used similarly labeled dual cyclic polytopes where all “Lemke paths” are exponentially long. A Lemke path is related to the path computed by the LH algorithm, but is defined on a *single* polytope that does not have a product structure corresponding to a bimatrix game. The completely labeled vertex found by a Lemke path can be interpreted as a symmetric equilibrium of a symmetric bimatrix game. However, as in the example in Figure 4 below, such a symmetric game may also have nonsymmetric equilibria which here are easy to compute, so that the result by Morris (1994) seemed unsuitable to describe games that are hard to solve with the LH algorithm.

The “imitation games” defined by McLennan and Tourky (2010) changed this picture. In an imitation game, the payoff matrix of one of the players is the identity matrix. The mixed strategy of that player in any Nash equilibrium of the imitation game corresponds exactly to a symmetric equilibrium of the symmetric game defined by the payoff matrix of the other player. In that way, an algorithm that finds a Nash equilibrium of a bimatrix game can be used to find a symmetric Nash equilibrium of a symmetric game. (The converse statement that a bimatrix game can be “symmetrized”, see Proposition 2 below, is an earlier folklore result stated for zero-sum games by Gale, Kuhn, and Tucker 1950.)

In one sense the two-polytope construction of Savani and von Stengel (2006) was overly complicated: the imitation games by McLennan and Tourky (2010) provide a simple and elegant way to turn the single-polytope construction of Morris (1994) into exponentially-long LH paths for bimatrix games. In another sense, the construction of Savani and von Stengel was not redundant. Namely, the square imitation games obtained from Morris (1994) have a single completely mixed equilibrium that is easily computed by equating all payoffs for all pure strategies. Savani and von Stengel (2006) extended their construction of square games with long LH paths (and a single completely mixed equilibrium) to non-square games that are *simultaneously* hard for the LH algorithm and “support enumeration”, which is another natural and simple algorithm for finding equilibria. The support of a mixed strategy is the set of pure strategies that are played with positive probability. Given a pair of supports of equal size, the mixed strategy probabilities are found by equating all payoffs for the other player’s support, which then have to be compared with payoffs outside the support to establish the equilibrium property (see Dickhaut and Kaplan 1991).

In this paper, we extend the idea of imitation games to games where one payoff matrix is arbitrary and the other is a set of unit vectors. We call these *unit vector games*. An imitation game is an example of a unit vector game, where the unit vectors form an identity matrix. The main result of this paper is an application of unit vector games: we use them to extend Morris’s

construction to obtain non-square bimatrix games that use only one dual cyclic polytope, rather than the two used by Savani and von Stengel, and which are simultaneously hard *both* for the LH algorithm and support enumeration. This result (Theorem 11) was first described by Savani (2006, Section 3.8).

Before presenting this construction in Section 3, we introduce in Section 2 the required background on labeled best response polytopes for bimatrix games, in an accessible presentation due to Shapley (1974) that we think every game theorist should know. We define unit vector games and the use of imitation games, and their relationships to the LH algorithm. We will make the case that unit vector games provide a general and simple way to construct bimatrix games using a single labeled polytope.

To our knowledge, unit vector games were first defined and used by Balthasar (2009, Lemma 4.10) in a different context, namely in order to prove that a symmetric equilibrium of a non-degenerate symmetric game that has positive “symmetric index” can be made the *unique* symmetric equilibrium of a larger symmetric game by adding suitable strategies (Balthasar 2009, Theorem 4.1).

2 Unit vector games

In this section, we first describe in Section 2.1 how labeled polyhedra capture the “best-response regions” of mixed strategies where a particular pure strategy of the other player is a best response, and how these are used to identify Nash equilibria. In Section 2.2 we introduce unit vector games, whose equilibria correspond to completely labeled vertices of a single labeled polytope. In Section 2.3 we discuss the role of imitation games for symmetric games and their symmetric equilibria. Finally, in Section 2.4, we show how Lemke paths defined for single labeled polytopes are “projections” of seemingly more general LH paths in the case of unit vector games (Theorem 5).

2.1 Nash equilibria of bimatrix games and polytopes

Consider an $m \times n$ bimatrix game (A, B) . We describe a geometric-combinatorial “labeling” method, due to Shapley (1974), that allows an easy identification of the Nash equilibria of the game. It has an equivalent description in terms of polytopes derived from the payoff matrices.

Let $\mathbf{0}$ be the all-zero vector and let $\mathbf{1}$ be the all-one vector of appropriate dimension. All vectors are column vectors and C^\top is the transpose of any matrix C , so $\mathbf{1}^\top$ is the all-one row vector. Let X and Y be the mixed-strategy simplices of the two players,

$$X = \{x \in \mathbb{R}^m \mid x \geq \mathbf{0}, \mathbf{1}^\top x = 1\}, \quad Y = \{y \in \mathbb{R}^n \mid y \geq \mathbf{0}, \mathbf{1}^\top y = 1\}. \quad (1)$$

It is convenient to identify the $m + n$ pure strategies of the two players by separate *labels* where the labels $1, \dots, m$ denote the m pure strategies of the row player 1 and the labels $m + 1, \dots, m + n$ denote the n pure strategies of the column player 2.

Consider mixed strategies $x \in X$ and $y \in Y$. We say that x has label $m + j$ for $1 \leq j \leq n$ if j is a pure best response of player 2 to x . Similarly, y has label i for $1 \leq i \leq m$ if i is a pure best response of player 1 to y . In addition, we say that x has label i for $1 \leq i \leq m$ if $x_i = 0$, and that

y has label $m + j$ for $1 \leq j \leq n$ if $y_j = 0$. That is, a mixed strategy such as x has label i (one of the player's own pure strategies) if i is not played.

In a Nash equilibrium, every pure strategy that is played with positive probability is a best response to the other player's mixed strategy. In other words, if a pure strategy is not a best response, it is played with probability zero. Hence, a mixed strategy pair (x, y) is a Nash equilibrium if and only if every label in $\{1, \dots, m + n\}$ appears as a label of x or of y . The Nash equilibria are therefore exactly those pairs (x, y) in $X \times Y$ that are *completely labeled* in this sense.

As an example, consider the 3×3 game (A, B) with

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 2 & 4 \\ 3 & 2 & 0 \\ 0 & 2 & 0 \end{pmatrix}. \quad (2)$$

The labels 1, 2, 3 represent the pure strategies of player 1 and 4, 5, 6 those of player 2. Figure 1 shows X and Y with these labels shown as circled numbers. The interiors of these triangles are covered by *best-response regions* labeled by the other player's pure strategies, which are closed polyhedral sets where the respective pure strategy is a best response. For example, the best-response region in Y with label 1 is the set of those (y_1, y_2, y_3) such that $y_1 \geq y_2$ and $y_1 \geq y_3$, due to the particularly simple form of A in (2). The outsides of X and Y are labeled with the players' own pure strategies where these are not played. These outside facets are opposite to the vertex where only that pure strategy is played; for example, label 1 is the label of the facet of X opposite to the vertex $(1, 0, 0)$. In Figure 1 there is only one pair (x, y) that is completely labeled, namely $x = (\frac{1}{3}, \frac{2}{3}, 0)$ with labels 3, 4, 5 and $y = (\frac{1}{2}, \frac{1}{2}, 0)$ with labels 1, 2, 6, so this is the only Nash equilibrium of the game.




Figure 1 Labeled mixed strategy sets X and Y for the game (2).

The subdivision of X and Y into best-response regions is most easily seen with the help of the “upper envelope” of the payoffs to the other player, which are defined by the following polyhedra. Let

$$\bar{P} = \{(x, v) \in X \times \mathbb{R} \mid B^\top x \leq \mathbf{1}v\}, \quad \bar{Q} = \{(y, u) \in Y \times \mathbb{R} \mid Ay \leq \mathbf{1}u\}. \quad (3)$$

For the example (2), the inequalities $B^\top x \leq \mathbf{1}v$ state that $3x_2 \leq v$, $2x_1 + 2x_2 + 2x_3 \leq v$, $4x_1 \leq v$, which say that v is at least the best-response payoff to player 2. If one of these inequalities

is tight (holds as equality), then v is exactly the best-response payoff to player 2. The left-hand diagram in Figure 2 shows these “best-response facets” of \bar{P} , and their projection to X by ignoring the payoff variable v , which defines the subdivision of X into best-response regions as in the left-hand diagram in Figure 1.




Figure 2 Best-response facets of the polyhedron \bar{P} in (3), and the polytope P in (4), for the game in (2).

Throughout this paper, assume (without loss of generality) that A and B^\top are non-negative and have no zero column. Then v and u in $B^\top x \leq \mathbf{1}v$ and $Ay \leq \mathbf{1}u$ are always positive. By dividing these inequalities by v and u , respectively, and writing x_i instead of x_i/v and y_j instead of y_j/u , the polyhedra \bar{P} and \bar{Q} are replaced by P and Q ,

$$P = \{x \in \mathbb{R}^m \mid x \geq \mathbf{0}, B^\top x \leq \mathbf{1}\}, \quad Q = \{y \in \mathbb{R}^n \mid Ay \leq \mathbf{1}, y \geq \mathbf{0}\}, \quad (4)$$

which are bounded and therefore polytopes. For B in (2), P is shown on the right in Figure 2.

Both polytopes P and Q in (4) are defined by $m+n$ inequalities that correspond to the pure strategies of the player, which we have denoted by the labels $1, \dots, m+n$. We can now identify the labels, as pure best responses of the other player, or unplayed own pure strategies, as tight inequalities in either polytope. That is, a point x in P has label k if the k th inequality in P is tight, that is, if $x_k = 0$ for $1 \leq k \leq m$ or $(B^\top x)_{k-m} = 1$ for $m+1 \leq k \leq m+n$. Similarly, y in Q has label k if $(Ay)_k = 1$ for $1 \leq k \leq m$ or $y_{k-m} = 0$ for $m+1 \leq k \leq m+n$. Then (x, y) in $P \times Q$ is *completely labeled* if x and y together have all labels in $\{1, \dots, m+n\}$. With the exception of $(\mathbf{0}, \mathbf{0})$, these completely labeled points of $P \times Q$ represent (after rescaling to become pairs of mixed strategies) exactly the Nash equilibria of the game (A, B) .

The pair (x, y) in $P \times Q$ is completely labeled if

$$x_i = 0 \text{ or } (Ay)_i = 1 \text{ for all } i = 1, \dots, m, \quad y_j = 0 \text{ or } (B^\top x)_j = 1 \text{ for all } j = 1, \dots, n. \quad (5)$$

Because x , $\mathbf{1} - Ay$, y , and $\mathbf{1} - B^\top x$ are all non-negative, the *complementarity* condition (5) can also be stated as the orthogonality condition

$$x^\top (\mathbf{1} - Ay) = 0, \quad y^\top (\mathbf{1} - B^\top x) = 0. \quad (6)$$

The characterization of Nash equilibria as completely labeled pairs (x, y) holds for arbitrary bimatrix games. For considering algorithms, it is useful to assume that the game is *nondegenerate* in the sense that no point in P has more than m labels, and no point in Q has more than

n labels. Clearly, for a nondegenerate game, in an equilibrium (x, y) each label appears exactly once either as a label of x or of y .

Nondegeneracy is equivalent to the condition that the number of pure best responses against a mixed strategy is never larger than the size of the support of that mixed strategy. It implies that P is a *simple* polytope in the sense that no point of P lies on more than m facets, and similarly that Q is a simple polytope. A facet is obtained by turning one of the inequalities that define the polytope into an equality, provided that the inequality is irredundant, that is, cannot be omitted without changing the polytope. A redundant inequality in the definition of P and Q may also give rise to a degeneracy if it corresponds to a pure strategy that is weakly (but not strictly) dominated by, or payoff equivalent to, a mixture of other strategies. For a detailed discussion of degeneracy see von Stengel (2002).

2.2 Unit vector games and a single labeled polytope

The components of the k th unit vector e_k are 0 except for the k th component, which is 1. In an $m \times n$ unit vector game (A, B) , every column of A is a unit vector in \mathbb{R}^m . The matrix B is arbitrary, and without loss of generality B^\top is non-negative and has no zero column.

In this subsection, we consider such a unit vector game (A, B) . Let the j th column of A be the unit vector $e_{\ell(j)}$, for $1 \leq j \leq n$. Then the sequence $\ell(1), \dots, \ell(n)$ together with the payoff matrix B completely specifies the game.

For this game, the polytope Q in (4) has a very special structure. For $1 \leq i \leq m$, let

$$N_i = \{j \mid \ell(j) = i, 1 \leq j \leq n\} \quad (7)$$

so that N_i is the set of those columns j whose best response is row i . These sets N_i are pairwise disjoint, and their union is $\{1, \dots, n\}$. Then clearly

$$Q = \{y \in \mathbb{R}^n \mid \sum_{j \in N_i} y_j \leq 1, 1 \leq i \leq m, y \geq \mathbf{0}\}. \quad (8)$$

That is, except for the order of inequalities, Q is the product of m simplices of the form $\{z \in \mathbb{R}^{N_i} \mid \sum_{j \in N_i} z_j \leq 1, z \geq \mathbf{0}\}$, for $1 \leq i \leq m$. If each N_i is a singleton, then, by (7), A is a permuted identity matrix, $n = m$, each simplex is the unit interval, and Q is the n -dimensional unit cube.

In any bimatrix game, the polytopes P and Q in (4) each have $m+n$ inequalities that correspond to the pure strategies of the two players. Turning the k th inequality into an equality typically defines a facet of the polytope, which defines the label k of that facet, $1 \leq k \leq m+n$.

In our unit vector game (A, B) where the j th column of A is the unit vector $e_{\ell(j)}$, for $1 \leq j \leq n$, we introduce the *labeled polytope* P^ℓ ,

$$P^\ell = \{x \in \mathbb{R}^m \mid x \geq \mathbf{0}, B^\top x \leq \mathbf{1}\}, \quad (9)$$

where the $m+n$ inequalities of P^ℓ have the labels i for the first m inequalities $x_i \geq 0$, $1 \leq i \leq m$, and the j th inequality of $B^\top x \leq \mathbf{1}$ has label $\ell(j)$, for $1 \leq j \leq n$. That is, P^ℓ is just the polytope P in (4) except that the last n inequalities are labeled with $\ell(1), \dots, \ell(n)$, each of which is a number in $\{1, \dots, m\}$. A point x of P^ℓ is *completely labeled* if every number in $\{1, \dots, m\}$

appears as the label of an inequality that is tight for x . In particular, if P^ℓ is a simple polytope with one label for each facet, then x is completely labeled if x is a vertex of P^ℓ so that the m facets that x lies on together have all labels $1, \dots, m$.

The following proposition shows that with these labels, P^ℓ carries all the information about the unit vector game, and the polytope Q is not needed. The proposition was first stated in a dual version by Balthasar (2009, Lemma 4.10), and in essentially this form by Véggh and von Stengel (2015, Proposition 1). Its proof also provides the first step of the proof of Theorem 5 below.

Proposition 1 *Consider a labeled polytope P^ℓ with labels as described following (9). Then x is a completely labeled point of $P^\ell - \{\mathbf{0}\}$ if and only if for some $y \in Q$ the pair (x, y) is (after scaling) a Nash equilibrium of the $m \times n$ unit vector game (A, B) where $A = [e_{\ell(1)} \cdots e_{\ell(n)}]$.*

Proof. Let $(x, y) \in P \times Q - \{(\mathbf{0}, \mathbf{0})\}$ be a Nash equilibrium, so it has all labels in $\{1, \dots, m+n\}$. Then x is a completely labeled point of P^ℓ for the following reason. If $x_i = 0$ then x has label i . If $x_i > 0$ then y has label i , that is, $(Ay)_i = 1$, which requires that for some j we have $y_j > 0$ and the j th column of A is equal to e_i , that is, $\ell(j) = i$. Because $y_j > 0$, and (x, y) is completely labeled, x has label $m+j$ in P , that is, $(B^\top x)_j = 1$, which means x has label $\ell(j) = i$ in P^ℓ , as required.

Conversely, let x be a completely labeled point of $P^\ell - \{\mathbf{0}\}$. Then for each i in $\{1, \dots, m\}$ with $x_i > 0$, label i for x comes from a binding inequality $(B^\top x)_j = 1$ with label $\ell(j) = i$, that is, for some $j \in N_i$ in (7). Let $y_j = 1$ and $y_h = 0$ for all $h \in N_i - \{j\}$, and do this for all i with $x_i > 0$. It is easy to see that the pair (x, y) is a completely labeled point of $P \times Q$. \square

The game in (2) is a unit vector game. For this game, the polytope P in (4) is shown on the right in Figure 2, where we have shown only the labels 4, 5, 6 for the “best response facets”. In addition, the facets with labels 1, 2, 3 where $x_1 = 0$, $x_2 = 0$, $x_3 = 0$ are the facets, hidden in this picture, at the back right, back left, and bottom of the polytope, respectively. In the polytope P^ℓ , the labels 4, 5, 6 are replaced by 1, 2, 3 because the corresponding columns of A are the unit vectors e_1, e_2, e_3 . Figure 3 shows this polytope in such a way that there is only one hidden facet, with label 1 where $x_1 = 0$. Apart from the origin $\mathbf{0}$, the only completely labeled point of P^ℓ is x as shown, which is part of a Nash equilibrium (x, y) as stated in Proposition 1.

A polytope like in (9) that has a label for each facet provides a particularly natural way to describe equilibrium-finding algorithms, as described in Section 2.4 below.

2.3 Reductions between equilibria of bimatrix and unit vector games

A method that “solves” a unit vector game in the sense of finding one equilibrium, or all equilibria, of the game, can be used to solve an arbitrary bimatrix game. The first step in seeing this is the fact that the equilibria of a bimatrix game correspond to the symmetric equilibria of a suitable symmetric game. This “symmetrization” has been observed for zero-sum games by Gale, Kuhn, and Tucker (1950) and seems to be a folklore result for bimatrix games.

Proposition 2 *Let (A, B) be a bimatrix game, and $(x, y) \in P \times Q - \{(\mathbf{0}, \mathbf{0})\}$ in (4). Then (x, y) (suitably scaled) is a Nash equilibrium of (A, B) if and only if (z, z) (suitably scaled) is a symmetric equilibrium of (C, C^\top) with $z = (x, y)$ and $C = \begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix}$.*




Figure 3 The polytope P^ℓ for the unit vector game (2). The hidden facet at the back has label 1, written on the left.

Proof. This holds by (6) because (z, z) is an equilibrium of (C, C^\top) if and only if $z \neq \mathbf{0}$, $z \geq \mathbf{0}$, $Cz \leq \mathbf{1}$, and $z^\top(\mathbf{1} - Cz) = 0$. \square

By Proposition 2, finding an equilibrium of a bimatrix game can be reduced to finding a symmetric equilibrium of a symmetric bimatrix game. The converse follows from the following proposition, due to McLennan and Tourky (2010, Proposition 2.1), with the help of imitation games. They define an imitation game as an $m \times m$ bimatrix game (A, B) where B is the identity matrix. Here, we define an imitation game as a special unit vector game (A, B) where A (rather than B) is the identity matrix I . The reason for this (clearly not very material) change is that this game is completely described by the polytope P^ℓ in (9), which corresponds to P in (4) and compared to Q has a more natural description because the m inequalities $x \geq \mathbf{0}$ with labels $1, \dots, m$ are listed first.

Proposition 3 *The pair (x, x) is a symmetric Nash equilibrium of the symmetric bimatrix game (C, C^\top) if and only if (x, y) is a Nash equilibrium of the imitation game (I, C^\top) for some y .*




Figure 4 Labeled mixed-strategy sets X and Y for the symmetric game (C, C^\top) in (10).

As an example, consider the symmetric game (C, C^\top) with

$$C = \begin{pmatrix} 0 & 3 & 0 \\ 2 & 2 & 2 \\ 4 & 0 & 0 \end{pmatrix}, \quad C^\top = \begin{pmatrix} 0 & 2 & 4 \\ 3 & 2 & 0 \\ 0 & 2 & 0 \end{pmatrix}, \quad (10)$$

so that $C^\top = B$ in (2). Figure 4 shows the labeled mixed-strategy simplices X and Y for this game. In addition to the symmetric equilibrium (x,x) where $x = (\frac{1}{3}, \frac{2}{3}, 0)$, the game has two non-symmetric equilibria (a,b) and (b,a) where $a = (\frac{1}{2}, \frac{1}{2}, 0)$ and $b = (0, \frac{2}{3}, \frac{1}{3})$. A method that just finds a Nash equilibrium of a bimatrix game may not find a symmetric equilibrium when applied to this game, which shows the use of Proposition 3. The corresponding imitation game (I, C^\top) is just (A, B) in (2), which has the unique equilibrium (x,y) where (x,x) is the symmetric equilibrium of (C, C^\top) .




Figure 5 (Left) Best-response regions for identifying symmetric equilibria. (Right) Degenerate symmetric game (11) with a unique symmetric equilibrium.

The left-hand diagram in Figure 5 shows the mixed strategy simplex X subdivided into regions of pure best responses against the mixed strategy itself, which corresponds to the polytope P^ℓ in Figure 3. The (in this case unique) symmetric equilibrium is the completely labeled point x .

The right-hand diagram in Figure 5 shows this subdivision of X for another game (C, C^\top) where

$$C = \begin{pmatrix} 0 & 4 & 0 \\ 2 & 2 & 2 \\ 4 & 0 & 0 \end{pmatrix}, \quad C^\top = \begin{pmatrix} 0 & 2 & 4 \\ 4 & 2 & 0 \\ 0 & 2 & 0 \end{pmatrix}. \quad (11)$$

This game is degenerate because the mixed strategy $x = (\frac{1}{2}, \frac{1}{2}, 0)$ has three pure best responses. This mixed strategy x also defines the unique symmetric equilibrium (x,x) of this game. However, the corresponding equilibria (x,y) of the imitation game (I, C^\top) are not unique, because due to the degeneracy any convex combination of $(\frac{1}{2}, \frac{1}{2}, 0)$ and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ can be chosen for y , as shown in Figure 6.

Hence the reduction between symmetric equilibria (x,x) of a symmetric game and Nash equilibria (x,y) of the corresponding imitation game stated in Proposition 3 does not preserve uniqueness if the game is degenerate.

2.4 Lemke paths and Lemke–Howson paths

Consider a labeled polytope P^ℓ as in (9). We assume throughout that P^ℓ is nondegenerate, that is, no point of P^ℓ has more than m labels. Therefore, P^ℓ is a simple polytope, and every tight inequality defines a separate facet (we can omit inequalities that are never tight), each of which has a label in $\{1, \dots, m\}$. The path-following methods described in this section can be extended to degenerate games and polytopes; for an exposition see von Stengel (2002).



Figure 6 Labeled mixed-strategy sets for the imitation game (I, C^\top) for the degenerate symmetric game (11) where the equilibria (x, y) are not unique.

A *Lemke path* is a path that starts at a completely labeled vertex of P^ℓ such as $\mathbf{0}$ and ends at another completely labeled vertex. It is defined by choosing one label k in $\{1, \dots, m\}$ that is allowed to be *missing*. After this choice of k , the path proceeds in a unique manner from the starting point. 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 , then the path 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 is repeated. The resulting path consists of a sequence of *k-almost complementary* edges and vertices (so defined by having all labels except possibly k , where k occurs only at the starting point and endpoint of the path). The path cannot revisit a vertex because this would offer a second way to proceed when that vertex is first encountered, which is not the case because P^ℓ is nondegenerate. Hence, the path terminates at another completely labeled vertex of P^ℓ (which is a Nash equilibrium of the corresponding unit vector game in Proposition 1 if the path starts at $\mathbf{0}$). Figure 7 shows an example.




Figure 7 Lemke path for missing label 1 for the polytope in Figure 3.

For a fixed missing label k , every completely labeled vertex of P^ℓ is a separate endpoint of a Lemke path. Because each path has two endpoints, there is an even number of them, and all of these except $\mathbf{0}$ are Nash equilibria of the unit vector game, so the number of Nash equilibria is odd.

This path-following method was first described by Lemke (1965) in order to find a solution to a linear complementarity problem (LCP); it is normally described for polyhedra, not for polytopes, so that termination requires additional assumptions (see Cottle, Pang, and Stone 1992). The standard description of an LCP assumes a square matrix B with labels $\ell(j) = j$ for $j = 1, \dots, m$. Allowing P^ℓ to have $m+n$ rather than $2m$ facets with individual labels $\ell(j)$ for the last n facets corresponds to a *generalized LCP* (sometimes also called “vertical LCP”), as studied in Cottle and Dantzig (1970). The term “Lemke paths” for polytopes is due to Morris (1994).

The algorithm by Lemke and Howson (1964) finds one Nash equilibrium of an $m \times n$ bimatrix game (A, B) . Let $C = \begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix}$ as in Proposition 2. Then one way to define a Lemke–Howson (LH) path for missing label k in $\{1, \dots, m+n\}$ is as a Lemke path for missing label k for the labeled polytope

$$R^\ell = \{z \in \mathbb{R}^{m+n} \mid z \geq \mathbf{0}, Cz \leq \mathbf{1}\} \quad (12)$$

where the $2(m+n)$ inequalities of R^ℓ have labels $1, \dots, m+n, 1, \dots, m+n$ (that is, $\ell(i) = i$ for $i = 1, \dots, m+n$).

The more conventional way to define the LH algorithm is to consider $P \times Q$ with P and Q as in (4). Clearly, with $z = (x, y)$, R^ℓ in (12) is equal to $P \times Q$. Starting from $(\mathbf{0}, \mathbf{0})$, the chosen missing label k is a pure strategy of player 1 (for $1 \leq k \leq m$) or of player 2 (for $m+1 \leq k \leq m+n$). Instead of a single point z that moves on the graph (of vertices and edges) of R^ℓ , the pair (x, y) (which equals z) moves on $P \times Q$ by alternately moving x on the graph of P and y on the graph of Q . This alternate move of a pair of “tokens” can be nicely shown for 3×3 games on the two mixed strategy sets X and Y subdivided into best-response regions as in Figure 1, extended with the origin $\mathbf{0}$ (as done by Shapley 1974). This is obviously more accessible than a path on a six-dimensional polytope, but requires keeping track of the alternating tokens.




Figure 8 Lemke–Howson path for missing label 1 for the game (2).

Figure 8 illustrates this for the game in (2). The pair of tokens starts on $\mathbf{0}, \mathbf{0}$, which is identified by the pair of label sets 123,456. Let 1 be the missing label, which means moving (in the left-hand diagram in Figure 8) from $(0,0,0)$ with labels 123 to the vertex $(1,0,0)$ of X with labels 236. The new pair has labels 236,456 with duplicate label 6, so the next move is in the right-hand diagram from $(0,0,0)$ with labels 456 to the vertex $(0,0,1)$ of Y with labels 345. The label that is picked up is 3 which is now duplicate, so the next move is in X from 236

to 256. Then 5 is duplicate, with a move in Y from 345 to 234. With 2 duplicate, the next move in X is from 256 to 356. Then 3 is duplicate, moving in Y from 234 to 246. Then 6 is duplicate, moving in X from 356 to 345, which is the point $x = (\frac{1}{3}, \frac{2}{3}, 0)$. Then 4 is duplicate, moving in Y from 246 to 126, which is the point $y = (\frac{1}{2}, \frac{1}{2}, 0)$ which has the missing label 1. This terminates the LH path for missing label 1 at the Nash equilibrium (x, y) .

The two diagrams in Figure 8 show two separate paths on P and Q , respectively (represented by X and Y subdivided into best-response regions). These paths are traversed in alternate steps and define a single path on the product polytope $P \times Q$. In general, a simple path on $P \times Q$ (that is, a path that does not revisit a vertex) may not “project” to simple paths on P and Q . However, for LH paths this is the case, as stated in the following proposition (Lemma 2.3 of Savani 2006, and implicit in McLennan and Tourky 2010, Section 4).

Proposition 4 *Every LH path on $P \times Q$ induces a simple path in each polytope P and Q , that is, no vertex of P or Q is ever left and visited again on an LH path.*

Proof. Suppose to the contrary that a vertex x of P is left and visited again on an LH path. This means that there are three vertex pairs (x, y) , (x, y') , and (x, y'') of $P \times Q$, with pairwise distinct vertices y , y' , and y'' of Q , on an LH path with missing label k , say. All three pairs have all labels except possibly k . The m labels of x define $n - 1$ labels shared by y , y' , and y'' . However, this is impossible, since these $n - 1$ labels correspond to $n - 1$ equations in \mathbb{R}^n that define a line, which can only contain two vertices of Q . The same reasoning applies to a vertex y of Q that would be visited multiple times on an LH path. \square

Consider an $m \times n$ unit vector game (A, B) where $A = [e_{\ell(1)} \cdots e_{\ell(n)}]$. According to Proposition 1, the labeled polytope P^ℓ carries all information about the Nash equilibria of (A, B) . Recall that P^ℓ is the polytope P in (4) but where the labels $m + j$ for the strategies j of the column player, $1 \leq j \leq n$, are replaced by $\ell(j)$, that is, by the best responses of the row player to these columns. Replacing these labels in the left-hand diagram in Figure 1 gives the left-hand diagram in Figure 5, equivalent to P^ℓ in Figure 3,

We now establish the same correspondence with regard to the LH paths on $P \times Q$ for the game (A, B) , where the corresponding “projection” to P defines a Lemke path on P^ℓ . For example, the LH path projected to P shown in the left-hand diagram in Figure 8 is the same as the Lemke path on P^ℓ in Figure 7. Both paths are defined for the missing label 1. It seems natural that the LH path for missing label i in $\{1, \dots, m\}$ projects to the Lemke path for missing label i on P^ℓ . However, there are n additional LH paths for the game (A, B) for the missing labels $m + j$ for j in $\{1, \dots, n\}$, which do not exist as labels of P^ℓ . The following theorem states that these project to the Lemke paths on P^ℓ for the missing label $\ell(j)$. This generalizes the corresponding assertion by McLennan and Tourky (2010, p. 9) and Savani and von Stengel (2006, Proposition 15) for imitation games where $\ell(j) = j$.

Theorem 5 *Consider an $m \times n$ unit vector game (A, B) where $A = [e_{\ell(1)} \cdots e_{\ell(n)}]$, with P^ℓ as in (9) and P and Q as in (4). Then the LH path on $P \times Q$ for this game for missing label k projects to a path on P that is the Lemke path on P^ℓ for missing label k if $1 \leq k \leq m$, and that is the Lemke path for missing label $\ell(j)$ if $k = m + j$ for $1 \leq j \leq n$.*

Proof. In Proposition 1 it was shown that the completely labeled pairs (x, y) of $P \times Q$ correspond to the completely labeled points x of P^ℓ . It is easy to see that if P is nondegenerate, as assumed here, then this correspondence is one-to-one, and x and y are vertices.

In the following, i is always an element of $\{1, \dots, m\}$, and j is always an element of $\{1, \dots, n\}$.

Consider a step of an LH path on $P \times Q$ that leaves or arrives at a vertex x of P , as part of a pair (x, y) . If the dropped label is i , then $x_i = 0$ changes to $x_i > 0$, and if the dropped label is $m + j$, then $(B^\top x)_j = 1$ changes to $(B^\top x)_j < 1$. If i is a label that is picked up, then $x_i > 0$ changes to $x_i = 0$, and if $m + j$ is a label that is picked up, then $(B^\top x)_j < 1$ changes to $(B^\top x)_j = 1$.

Similarly, consider a vertex y of Q . Because Q is a product of m simplices as in (8), for each i the following holds: either $y_j = 0$ for all $j \in N_i$, or for exactly one $j \in N_i$ we have $y_j = 1$ (which means $(Ay)_i = 1$ and y has label $\ell(j) = i$) and $y_h = 0$ for all $h \in N_i - \{j\}$. We can also describe precisely which label is picked up after moving away from y by dropping a label:

- (a) If the dropped label is i , then $y_j = 1$ (for some $j \in N_i$) changes to $y_j = 0$, so that $m + j$ is the label that is picked up.
- (b) If the dropped label is $m + j$, this is just the reverse step: $j \in N_i$ for a unique $i = \ell(j)$, so $y_j = 0$ changes to $y_j = 1$, which means $i = \ell(j)$ is the label that is picked up.

Consider now steps on an LH path with missing label k , and assume that any label that is picked up is not the missing label k , and therefore duplicate. Suppose label i is picked up in P , corresponding to the binding inequality $x_i = 0$. Label i is duplicate and therefore dropped in Q . By (a), this means that $m + j$ with $j \in N_i$ is picked up in Q , where $i = \ell(j)$. The duplicate label $m + j$ in P corresponds to the binding inequality $(B^\top x)_j = 1$. So the next step is to move away from this facet in P . In P^ℓ this same facet with $(B^\top x)_j = 1$ has label $\ell(j) = i$, and moving away from this facet is exactly the next step on the Lemke path on P^ℓ .

Similarly, suppose label $m + j$ is picked up in P , which corresponds to the facet $(B^\top x)_j = 1$ which in P^ℓ has label $i = \ell(j)$. On the LH path, the duplicate label $m + j$ in Q is dropped as in (b), where label i is picked up in Q and therefore duplicate. In P , the facet with this duplicate label is given by $x_i = 0$. The next step on the LH path is to move away from this facet, which is the same facet from which the Lemke path on P^ℓ moves away.

Similar considerations apply when the LH path is started or terminates. If the missing label is k in $\{1, \dots, m\}$, then the LH path starts by dropping k in P , and the Lemke path starts in the same way in P^ℓ . When the LH path terminates by picking up the missing label k in P , the Lemke path ends in the same way in P^ℓ . If it terminates by picking up the missing label k in Q , then by (b) this was preceded by dropping the previously duplicate label $m + j$ where $j \in N_k$, that is, after the path reached in P the facet defined by $(B^\top x)_j = 1$ which has label $\ell(j) = k$ in P^ℓ , so the Lemke path has already terminated on P^ℓ .

The LH path with missing label $k = m + j$ starts by dropping this label in Q . By (b), the label that is picked up in Q is $\ell(j)$, which is now duplicate, and the path proceeds by dropping this label in P which is the same as starting the Lemke path on P^ℓ with this missing label. The LH path terminates by picking up the missing label $k = m + j$ in P by reaching the facet defined by $(B^\top x)_j = 1$ which has label $\ell(j)$, so that the Lemke path on P^ℓ terminates. Alternatively, label $k = m + j$ is picked up in Q which by (a) was preceded by dropping label $i = \ell(j)$, which

was duplicate because it was picked up in P when encountering the facet $x_i = 0$, where i is the missing label $\ell(j)$ on the Lemke path that has therefore terminated on P^ℓ . \square

3 Hard-to-solve bimatrix games

With the help of Theorem 5, it suffices to construct suitable labeled polytopes with (exponentially) long Lemke paths in order to show that certain games have long LH paths. McLennan and Tourky (2010) (summarized in Savani and von Stengel 2006, Section 5) showed with the help of imitation games that the polytopes with long Lemke paths due to Morris (1994) can be used for this purpose. In this section we extend this construction, with the help of unit vector games, to games that are not square and that are hard to solve not only with the Lemke–Howson algorithm, but also with support enumeration methods.

In Section 3.1 we present a very simple model of random games that have very few Nash equilibria on average, unlike games where all payoffs are chosen at random. These games are unit vector games, and the result (Proposition 6) is joint work with Andy McLennan. We then describe in Section 3.2 dual cyclic polytopes, whose facets have a nice combinatorial structure, which have proved useful for the construction of games with many equilibria, and with long LH paths. Our main result, Theorem 11 in Section 3.3, describes unit vector games based on dual cyclic polytopes whose equilibria are hard to find not only with the LH algorithm, but also with support enumeration.

3.1 Permutation games

We present here a small “warmup” result that was found jointly with Andy McLennan. A *permutation game* is an $n \times n$ game (A, B) where A is the identity matrix and B is a permuted identity matrix, that is, the i th row of B is the unit vector $e_{\pi(i)}^\top$ for some permutation π of $\{1, \dots, n\}$ (so column $\pi(i)$ is the best response to row i , and, because $A = I$, the best response to column j is row j). Let I^π be this matrix B , so that the permutation game is (I, I^π) .

Because a permutation game (I, I^π) is an imitation game, the two strategies in an equilibrium have equal support. It is easy to see that any equilibrium of (I, I^π) is of the form (x, x) where x mixes uniformly over its support S where S is any nonempty subset of $\{1, \dots, n\}$ that is closed under π , that is, $i \in S$ implies $\pi(i) \in S$. In other words, S is any nonempty union of *cycles* of π .

A very simple model of a “random” game is to consider a permutation game (I, I^π) for a random permutation π .

Proposition 6 *A random $n \times n$ permutation game has in expectation n Nash equilibria.*

Proof. Consider a random permutation π of $\{1, \dots, n\}$. Let $E(n)$ be the expected number of Nash equilibria of (I, I^π) , where we want to prove that $E(n) = n$, which is true for $n = 1$. Let $n > 1$ and assume as inductive hypothesis that the claim is true for $n - 1$. With probability $\frac{1}{n}$ we have $\pi(n) = n$, in which case π defines also a random permutation of $\{1, \dots, n - 1\}$, and any equilibrium of (I, I^π) is either the pure strategy equilibrium where both players play n , or an equilibrium with a support S of a random $(n - 1) \times (n - 1)$ permutation game, or an equilibrium

with support $S \cup \{n\}$. Hence, in this case the number of equilibria of (I, I^π) is twice the number $E(n-1)$ of equilibria of a random $(n-1) \times (n-1)$ game plus one. Otherwise, with probability $\frac{n-1}{n}$, we have $\pi(n) \neq n$, so that π defines a random permutation of $\{1, \dots, n-1\}$ when removing n from the cycle of π that contains n . For any equilibrium of the $(n-1) \times (n-1)$ permutation game whose support contains this cycle, we add n back to the cycle to obtain the respective equilibrium of the $n \times n$ game. So in the case $\pi(n) \neq n$ the expected number of equilibria is $E(n-1)$. That is,

$$E(n) = \frac{1}{n}(1 + 2 \cdot E(n-1)) + \frac{n-1}{n}E(n-1) = \frac{1}{n}(1 + 2(n-1) + (n-1)(n-1)) = n,$$

which completes the induction. □

Random permutation games have very *few* equilibria, as Proposition 6 shows. In contrast, McLennan and Berg (2005) have shown that the expected number of equilibria of an $n \times n$ game with random payoffs is exponential in n . Bárány, Vempala, and Vetta (2007) show that such a game has with high probability an equilibrium with small support. A permutation game (I, I^π) , where the permutation π has k cycles, has $2^k - 1$ many equilibria, but a large number k of cycles is rare. In fact, there are $(n-1)!$ single-cycle permutations, so with probability $1/n = (n-1)!/n!$ the permutation game has only a single equilibrium with full support. For such games, an algorithm that enumerates all possible supports starting with those of small size takes exponential time. On the other hand, it is easy to see that the LH algorithm finds an equilibrium in the shortest possible time, because it just adds the strategies in a cycle of π to its current support.

However, a square game has only one full support, which is natural to test as to whether it defines a (completely mixed) equilibrium. The full support always defines an equilibrium in a permutation game. It also does for the square games described by Savani and von Stengel (2006) which have exponentially long LH paths. They therefore constructed also non-square games where support enumeration takes exponentially long time on average. It is an open question whether non-square games can be constructed from unit vectors as an extension of permutation games that are also hard to solve with support enumeration.

3.2 Cyclic polytopes and Gale evenness bitstrings

With the polytopes P and Q in (4), Nash equilibria of bimatrix games correspond to completely labeled points of $P \times Q$. The “dual cyclic polytopes” have the property that they have the maximal possible number of vertices for a given dimension and number of facets (see Ziegler 1995, or Grünbaum 2003). In addition, it is easy to describe each vertex by the facets it lies on. Using these polytopes, von Stengel (1999) constructed counterexamples for $n \geq 6$ to a conjecture by Quint and Shubik (1997) that a nondegenerate $n \times n$ game has at most $2^n - 1$ equilibria. McLennan and Park (1999) proved this conjecture for $n = 4$; the case $n = 5$ is still open. Morris (1994) gave a construction of labeled dual cyclic polytopes with exponentially long Lemke paths, which we extend in Theorem 11 below.

A standard way to define a *cyclic polytope* P' in dimension m with f vertices is as the convex hull of f points $\mu(t_j)$ on the *moment curve* $\mu: t \mapsto (t, t^2, \dots, t^m)^\top$ for $1 \leq j \leq f$. However, the

polytopes in (4) are defined by inequalities and not as convex hulls of points. In the *dual* (or “polar”) of a polytope, its vertices are reinterpreted 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 $\mu(t_j)$ from each such point. The resulting vectors $\mu(t_j) - \bar{\mu}$ then define the *dual cyclic polytope* in dimension m with f facets

$$C_f^m = \{x \in \mathbb{R}^m \mid (\mu(t_j) - \bar{\mu})^\top x \leq 1, 1 \leq j \leq f\}. \quad (13)$$

A suitable affine transformation of C_f^m (see von Stengel 1999, p. 560) gives a polytope P as in (4) or (9) so that the first m inequalities of P have the form $x \geq \mathbf{0}$. The last $n = f - m$ inequalities $B^\top x \leq \mathbf{1}$ of P then determine the $m \times n$ payoff matrix B . If the first m inequalities have labels $1, \dots, m$ and the last n inequalities have labels $\ell(1), \dots, \ell(n)$, then this defines a labeled polytope P^ℓ as in (9) and a unit vector game as in Proposition 1.

A vertex u of C_f^m is characterized by the *bitstring* $u_1 u_2 \dots u_f$ of length f , where the j th bit u_j indicates whether u is on the j th facet ($u_j = 1$) or not ($u_j = 0$). The polytope is simple, so exactly m bits are 1, and the other $f - m$ bits are 0. Assume (which is all that is needed) that $t_1 < t_2 < \dots < t_f$ when defining the j th facet of C_f^m by the binding inequality $(\mu(t_j) - \bar{\mu})^\top x = 1$ in (13). As shown by Gale (1963), the vertices of C_f^m are characterized by the bitstrings that fulfill the *Gale evenness* condition: A bitstring with exactly m 1s represents a vertex if and only if in any substring of the form $01^s 0$ the number s of 1s is even, so it has no odd-length substrings of the form 010 , 01110 , and so on (the reason is that the two zeros $u_i = u_j = 0$ at the end of such an odd-length substring would represent two points $\mu(t_i)$ and $\mu(t_j)$ on the moment curve that are on opposite sides of the hyperplane through the points $\mu(t_k)$ for $u_k = 1$, so that this hyperplane cannot define a facet of the cyclic polytope that is the convex hull of all the points, and therefore does not correspond to a vertex of the dual cyclic polytope). Initial substrings $1^s 0$ and terminal substrings 01^t are allowed to have an odd number s or t of 1s. We only consider *even* dimensions m , where s and t can only be both odd and by a cyclic shift (“wrapping around”) of the bitstring define an even-length substring $01^t 1^s 0$, which shows the cyclic symmetry of the Gale evenness condition.

Consider, for even m , the bitstrings of length f with m 1s that fulfill Gale evenness, and as before let $n = f - m$. One such string is $1^m 0^n$, that is, m 1s followed by n 0s. For the corresponding vertex of C_f^m , the first m inequalities are tight, and if we label them with $1, \dots, m$, then this defines the completely labeled vertex that is mapped to $\mathbf{0}$ in the affine map from C_f^m to the polytope P , which will be a labeled polytope P^ℓ . The last n facets of P^ℓ correspond to the last n positions of the bitstring, and they have labels $\ell(1), \dots, \ell(n)$. If we view ℓ as a string $\ell(1) \dots \ell(n)$ of n labels, each of which is an element of $\{1, \dots, m\}$, then these labels specify a labeled polytope. A *completely labeled* vertex corresponds to a Gale evenness bitstring $u_1 \dots u_f$ with $f = m + n$ where the positions i so that $u_i = 1$ have all m labels, the label being i if $1 \leq i \leq m$, and $\ell(j)$ if $i = m + j$ for $1 \leq j \leq n$. We call the resulting polytope C_ℓ^m , so this is the dual cyclic polytope C_f^m where $f = m + n$ and n is the length of the string ℓ of the last n facet labels, mapped affinely to P^ℓ as in (9), with facet labels as described.

3.3 Triple Morris games

In the notation just introduced, Morris (1994) studied Lemke paths on the labeled dual cyclic polytope C_σ^m , which we call the *Morris polytope*, for a string σ of m labels defined as follows. Let τ be the string $\tau(1)\cdots\tau(m)$ of m labels, which is 1324 for $m = 4$, 132546 for $m = 6$, 13254768 for $m = 8$, and in general defined by

$$\tau(1) = 1, \quad \tau(i) = i + (-1)^i \quad (2 \leq i \leq m-1), \quad \tau(m) = m, \quad (14)$$

and let σ be the string τ in reverse order, that is,

$$\sigma(i) = \tau(m-i+1) \quad (1 \leq i \leq m), \quad (15)$$

so $\sigma = 4231$ for $m = 4$, $\sigma = 645231$ for $m = 6$, $\sigma = 86745231$ for $m = 8$, and so on. We define the *triple Morris polytope* as $C_{\sigma\tau\sigma}^m$, where the concatenated string $\sigma\tau\sigma$ is a string of $3m$ labels, for example 645231132546645231 if $m = 6$.




Figure 9 Lemke paths for missing label 1 on the Morris polytope C_6^6 (left), and on the triple Morris polytope $C_{\sigma\tau\sigma}^6$ (right).

The left-hand diagram of Figure 9 shows the Lemke path for missing label 1 on the Morris polytope C_σ^m for $m = 6$. The top row gives the labels, where the first m are the labels $1, \dots, m$ corresponding to the inequalities $x \geq \mathbf{0}$ in P^ℓ , followed by the labels 645231 of σ . The rows below show the vertices of C_σ^6 as bitstrings, where bit 1 is written in a different font and 0 as a dot \cdot to distinguish them better. The first string 111111000000 represents the starting vertex $\mathbf{0}$ of P^ℓ . A facet that is left by dropping a label, at first the missing label 1, has a small “v” underneath the bit 1, whereas the facet that is just encountered, with the corresponding label that is picked up, has the “v” above it. Dropping label 1 means the second vertex is 011111100000, where label 6 is picked up and duplicate. Because the previous facet with that label corresponds to

the second to last bit 1, it is dropped next, which gives the next vertex as 011110110000 where label 4 is picked up, and so on.

The right-hand diagram in Figure 9 shows the Lemke path for missing label 1 on the triple Morris polytope $C_{\sigma\tau\sigma}^6$. Because in this case the only affected bits are those with labels in the first substring σ of the entire label string $\sigma\tau\sigma$, the path is essentially the same as in the Morris polytope C_σ^6 on the left.




Figure 10 Lemke paths for missing label 4 on the Morris polytope C_σ^6 (left), and on the triple Morris polytope $C_{\sigma\tau\sigma}^6$ (right).

Figure 10 shows the Lemke paths for these two polytopes for the missing label 4. In this case, to preserve Gale evenness, the bitstring that follows the starting bitstring 111111000000 is 111011000001, which “wraps around” the left end to add a bit 1 in the rightmost position, which has label 1 that is picked up. The resulting path is the composition of two sub-paths. The first path moves away from the dropped label 4 to the left (and wrapping around), which behaves essentially like a path with dropped label 4 on a Morris polytope in dimension 4, until label 5 is picked up. This starts the second sub-path with the original label 5 being dropped, which is essentially a (rather short) path with dropped label 1 on a Morris polytope in dimension 2.

In general, the Lemke path on the Morris polytope C_σ^m for the missing label k , where k is even, is the composition of two sub-paths. The first sub-path is equivalent to a Lemke path for missing label k on a Morris polytope C_σ^k for missing label k , which, by symmetry (writing the label strings backwards), is the same as the Lemke path on C_σ^k for missing label 1. The second sub-path is equivalent to a Lemke path for missing label 1 on a Morris polytope C_σ^{m-k} . (If k is odd, then a similar consideration applies by symmetry.) In this way, the Lemke paths for any missing label k are described by considering Lemke paths for missing label 1 in dimension k or $m - k$, where clearly k or $m - k$ is at least $m/2$ (which is used in the second part of Theorem 8 below).

The right-hand diagram in Figure 10 shows that the same Lemke path for missing label 4 results in the triple Morris polytope $C_{\sigma\tau\sigma}^6$, because the two copies of the label string σ have the same effect as the single label string σ in C_σ^6 . Clearly, this correspondence holds for any missing label.

Proposition 7 *There is a one-to-one correspondence between the Lemke path for missing label k starting from the Gale evenness string $1^m 0^m$ (vertex $\mathbf{0}$) of the Morris polytope C_σ^m and the*

Lemke path for missing label k starting from the Gale evenness string $1^m 0^{3m}$ (vertex $\mathbf{0}$) of the triple Morris polytope $C_{\sigma\tau\sigma}^m$, for $1 \leq k \leq m$.

The length of the Lemke path for missing label 1 on C_{σ}^m is *exponential* in the dimension m . Essentially, this path composed of two such paths in dimension $m - 2$, with another such path in dimension $m - 4$ between them (Figure 9 gives an indication). Hence, if the length of the path is a_m , the recurrence $a_m = 2a_{m-2} + a_{m-4}$ implies that it grows from a_{m-2} to a_m by an approximate factor of $1 + \sqrt{2}$; for details see Morris (1994), and for similar arguments Savani and von Stengel (2006, Theorem 7). Recall that $\Theta(f(n))$ means bounded above and below by a constant times $f(n)$ for large n .

Theorem 8 (Morris 1994, Proposition 3.4) *The longest Lemke path on C_{σ}^m is for missing label 1 and has length $\Theta((1 + \sqrt{2})^{m/2})$. The shortest Lemke path on C_{σ}^m is for missing label $m/2$ and has length $\Theta((1 + \sqrt{2})^{m/4})$.*

Consequently, the Lemke paths on triple Morris polytopes are also exponentially long. Hence, these polytopes define unit vector games which by Theorem 5 have exponentially long LH paths. We consider these games because they are of dimension $m \times 3m$ rather than $m \times m$ for the unit vector game defined by the Morris polytope C_{σ}^m . The latter, square game has a single completely mixed equilibrium, which is easily found by support enumeration. We show next that the $m \times 3m$ game has multiple equilibria, each of them with full support for player 1 (for which we need the “middle” label string τ).

Proposition 9 *The $m \times 3m$ unit vector game that corresponds to the triple Morris polytope $C_{\sigma\tau\sigma}^m$ has $3^{m/2}$ Nash equilibria. Each of them has full support for player 1.*

	1	2	3	4	5	6	σ					τ						σ						
	1	2	3	4	5	6	6	4	5	2	3	1	1	3	2	5	4	6	6	4	5	2	3	1
						1	0											0	0					
					1	1	0	0								0	0	0	0	0				
				1	1	1	0	0	0							0	0	0	0	0	0			
			1	1	1	1	0	0	0	0			0	0	0	0	0	0	0	0	0	0		0
		1	1	1	1	1	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0
	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
(a)						0	0											1	0					
					0	0	0	0										1	1	0	0			
(b)				0	0	0	0	0	1							0	1	1	0	0	0			
			0	0	0	0	0	0	1	1						0	0	1	1	0	0	0		
(c)		0	0	0	0	0	0	0	1	1	0		1	0	0	1	1	0	0	0	0	0	0	
	0	0	0	0	0	0	0	0	1	1	0	0	1	1	0	0	1	1	0	0	0	0	0	

Figure 11 Illustration of the proof of Proposition 9 for $m = 6$. The top half shows the only completely labeled bitstring $u = 1^m 0^{3m}$ where $u_m = 1$, the bottom half one such string where $u_m = 0$. There are three choices in each of the $m/2$ lines (a), (b), (c).

Proof. With the label string $\sigma\tau\sigma$, we identify the completely labeled vertices of $C_{\sigma\tau\sigma}^m$ as completely labeled Gale evenness strings $u = u_1 \cdots u_{4m}$. First, we show that if $u_m = 1$, then $u = 1^m 0^{3m}$, which is the vertex $\mathbf{0}$. This is illustrated in the top part of Figure 11. Because $u_m = 1$ which has label m , the other positions j with label m have $u_j = 0$, which are $j = m + 1$ and $j = 3m + 1$ (the first positions in the two substrings σ of $\sigma\tau\sigma$) and $j = 3m$ (the last position of τ), shown in the first line of Figure 11. The substring 10 with bits u_m and u_{m-1} requires $u_{m-1} = 1$ by Gale evenness, with label $m - 1$, which now requires $u_j = 0$ for $j = m + 3, 3m - 2, 3m + 3$, as shown for label 5 in the second line of Figure 11. The single bit u_{m+2} (label 4 in the picture) must be 0 by Gale evenness, and similarly $u_{3m-1} = u_{3m+2} = 0$ and hence $u_{m-2} = 1$, as shown in the next line. Continuing in this manner, the only possible string is $u = 1^m 0^n$ as claimed.

Suppose now that $u_m = 0$, where we will show that the resulting completely labeled Gale evenness string is of the form $u = 0^m \beta$, which represents a Nash equilibrium (x, y) of the game with full support for player 1, that is, $x > \mathbf{0}$. Consider the lower part of Figure 11, where in the first line (a) we now have three choices where to put the label m , namely by setting $u_j = 1$ for exactly one j in $\{m + 1, 3m, 3m + 1\}$, corresponding to the first position in one of the σ s or the last position in τ (where we choose the latter in the picture). So $u_{3m} = 1$, which requires $u_{3m-1} = 1$ by Gale evenness because $u_{3m+1} = 0$. This next position always has label $m - 2$ (label 4 if $m = 6$), so that $u_{m-2} = 0$, and similarly in the other positions with that label. But then $u_{m-1} = 0$ by Gale evenness and we again have three choices, in one of the substrings σ , τ , σ , of where to set that next bit with label $m - 1$. In the picture, we choose it in line (b) in the first substring σ , that is, $u_{m+3} = 1$. Continuing in that manner, there are $m/2$ times where we can choose a pair 11 of two bits 1 in either substring σ , τ , σ to obtain a completely labeled Gale evenness string, making $3^{m/2}$ choices in total, as claimed. \square

The $m \times 3m$ game in Proposition 9 has an exponential number of equilibria, which define a certain set E of equilibrium supports of player 2. However, they form an exponentially small subset of all possible supports. An equilibrium is therefore hard to find with a support enumeration algorithm, even if that algorithm is restricted to testing only supports of size m for player 2.

Proposition 10 *Consider an $m \times 3m$ game where a pair of supports defines a Nash equilibrium if and only if both supports have size m , and player 2's support belongs to the set E , a set of m -sized subsets of $\{1, \dots, 3m\}$. A support enumeration algorithm that tests supports picked uniformly at random without replacement from the set U of all m -sized subsets of $\{1, \dots, 3m\}$ has to test an expected number of*

$$\frac{\binom{3m}{m} - |E|}{|E| + 1} + 1 \quad (16)$$

supports before finding an equilibrium support.

Proof. To find the expected number of guesses required to find an equilibrium we use a standard argument (Motwani and Raghavan 1995, p. 10). Consider a random enumeration of the elements of U . The elements of $U - E$, which we index by $i = 1, \dots, |U - E|$, correspond to

non-equilibrium supports. Let W_i be the *indicator variable* that takes value 1 if the i th element of $U - E$ precedes all members of E in the enumeration of U , and 0 otherwise. Then $W = \sum_{i=1}^{|U-E|} W_i$ is the random variable equal to the number of supports checked before the first equilibrium is found. For a single element of $U - E$, the probability that it is in front of all elements of E is $1/(|E| + 1)$. Hence, using the linearity of expectation,

$$\mathbb{E}(W) = \mathbb{E}\left(\sum_{i=1}^{|U-E|} W_i\right) = \sum_{i=1}^{|U-E|} \mathbb{E}(W_i) = \sum_{i=1}^{|U-E|} \frac{1}{|E|+1} = \frac{|U| - |E|}{|E|+1} = \frac{\binom{3m}{m} - |E|}{|E|+1}.$$

This shows that the expected number of support guesses until an equilibrium is found is given by (16), as claimed. \square

In Proposition 10, we assume that the algorithm does not identify any particular pattern as to which supports should be tested. One way to achieve this is to permute the columns of the game randomly (if one knows that the payoff matrix B of player 2 is derived from a dual cyclic polytope, then this random order can be identified with a specialized method, see Savani 2006, Section 3.6; this is not a general method for solving games so we do not consider it). However, unless one distorts the polytope Q in (8), this still leaves a payoff matrix of player 1 where each unit vector appears three times. In this case, even if the algorithm picks only columns where each unit vector appears once, there would be 3^m possible supports which define a set U of size 3^m rather than $\binom{3m}{m}$ in Proposition 10. Such a set U is still exponentially large compared to the set E of $3^{m/2}$ supports that define a Nash equilibrium. In that case the expected time for the support-testing algorithm in the following theorem is $(\sqrt{3})^m \approx 1.732^m$.

Theorem 11 *Finding a Nash equilibrium of the $m \times 3m$ unit-vector game that corresponds to the triple Morris polytope $C_{\sigma\tau\sigma}^m$ takes at least time $\Theta((1 + \sqrt{2})^{m/4}) \approx \Theta(1.246^m)$ with the Lemke–Howson algorithm, and on expectation time $\Theta((27/4\sqrt{3})^m / \sqrt{m}) \approx \Theta(3.897^m / \sqrt{m})$ with an algorithm that tests in random order arbitrary supports of size $m \times m$ of the game.*

Proof. The length of the LH paths follows from Theorem 8, Proposition 7, and Theorem 5. For the support-testing algorithm, we have $|E| = (\sqrt{3})^m$ in Proposition 10 by Proposition 9. Using Stirling’s formula $n! \sim \sqrt{2\pi n} \cdot (n/e)^n$, we have $\binom{3m}{m} \sim (\sqrt{3} \cdot 3^{3m}) / (2\sqrt{\pi m} \cdot 2^{2m})$, so that the expression in (16) is $\Theta((27/4\sqrt{3})^m / \sqrt{m})$. \square

To conclude, we note results on the following combinatorial problem: Let m be even and let ℓ be a string of n labels from $\{1, \dots, m\}$, and consider the set of Gale evenness bitstrings of length $m + n$ which encode the vertices of the labeled polytope C_ℓ^m . The problem is to find a second completely labeled Gale evenness string other than $1^m 0^n$. Casetti, Merschen, and von Stengel (2010) have shown that this is equivalent to finding a second perfect matching in the *Euler graph* with nodes $1, \dots, m$ and edges defined by the Euler tour $1, \dots, m, \ell(1), \dots, \ell(n), 1$. The edges in a perfect matching encode the pairs of 1s in a Gale evenness bitstring, which is completely labeled because the edges cover all nodes. Végé and von Stengel (2015, Theorem 12) give a near-linear time algorithm that finds such a second perfect matching that, in addition, has opposite *sign*, which corresponds to a Nash equilibrium of positive index as it would be found by a Lemke path (which, however, can be exponentially long). So this combinatorial problem is simpler than the problem of finding a Nash equilibrium of a bimatrix game, even though it gives rise to games that are hard to solve by the standard methods considered in Theorem 11.

Acknowledgements

The first author was supported by EPSRC grant EP/L011018/1. We thank the referees for helpful comments.

References

- Balthasar, A. V. (2009), *Geometry and Equilibria in Bimatrix Games*, PhD Thesis, London School of Economics.
- Bárány, I., S. Vempala, and A. Vetta (2007), “Nash equilibria in random games,” *Random Structures and Algorithms* 31, 391–405.
- Casetti, M. M., J. Merschen, and B. von Stengel (2010), “Finding Gale strings,” *Electronic Notes in Discrete Mathematics* 36, 1065–1072.
- Chen, X., and X. Deng (2006), “Settling the complexity of two-player Nash equilibrium,” *Proc. 47th Symp. Foundations of Computer Science (FOCS)*, 261–272.
- Cottle, R. W., and G. B. Dantzig (1970), “A generalization of the linear complementarity problem,” *J. Combinatorial Theory* 8, 79–90.
- Cottle, R. W., J.-S. Pang, and R. E. Stone (1992), *The Linear Complementarity Problem*. Academic Press, San Diego.
- Daskalakis, C., P. W. Goldberg, and C. H. Papadimitriou (2009), “The complexity of computing a Nash equilibrium,” *SIAM Journal on Computing* 39, 195–259.
- Dickhaut, J., and T. Kaplan (1991), “A program for finding Nash equilibria,” *The Mathematica Journal* 1, Issue 4, 87–93.
- Gale, D. (1963), “Neighborly and cyclic polytopes,” V. Klee, ed., *Convexity, Proc. Seventh Symposium in Pure Mathematics, Vol. 7*, 225–232, Providence, Rhode Island: American Mathematical Society.
- Gale, D., H. W. Kuhn, and A. W. Tucker (1950), “On symmetric games,” H. W. Kuhn and A. W. Tucker, eds., *Contributions to the Theory of Games I, Annals of Mathematics Studies 24*, 81–87, Princeton: Princeton University Press.
- Goldberg, P. W., C. H. Papadimitriou, and R. Savani (2013), “The complexity of the homotopy method, equilibrium selection, and Lemke–Howson solutions,” *ACM Transactions on Economics and Computation* 1, Article 9.
- Grünbaum, B. (2003), *Convex Polytopes, 2nd ed.*, Graduate Texts in Mathematics, Vol. 221, New York: Springer.
- Lemke, C. E. (1965), “Bimatrix equilibrium points and mathematical programming,” *Management Science* 11, 681–689.
- 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.
- McLennan, A., and J. Berg (2005), “Asymptotic expected number of Nash equilibria of two-player normal form games,” *Games and Economic Behavior* 51, 264–295.
- McLennan, A., and I.-U. Park (1999), “Generic 4×4 two person games have at most 15 Nash equilibria,” *Games and Economic Behavior* 26, 111–130.
- McLennan, A., and R. Tourky (2010), “Imitation games and computation,” *Games and Economic Behavior* 70, 4–11.
- Morris, W. D., Jr. (1994), “Lemke paths on simple polytopes,” *Mathematics of Operations Research* 19, 780–789.
- Motwani, R., and P. Raghavan (1995), *Randomized Algorithms*, Cambridge: Cambridge University Press.

- 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.
- Quint, T., and M. Shubik (1997), “A theorem on the number of Nash equilibria in a bimatrix game,” *International Journal of Game Theory* 26, 353–359.
- Savani, R. (2006), *Finding Nash Equilibria of Bimatrix Games*, PhD Thesis, London School of Economics.
- Savani, R., and B. von Stengel (2006), “Hard-to-solve bimatrix games,” *Econometrica* 74, 397–429.
- Shapley, L. S. (1974), “A note on the Lemke–Howson algorithm,” *Mathematical Programming Study 1: Pivoting and Extensions*, 175–189.
- Végh, L. A., and B. von Stengel (2015), “Oriented Euler complexes and signed perfect matchings,” *Mathematical Programming Series B* 150, 153–178.
- von Stengel, B. (1999), “New maximal numbers of equilibria in bimatrix games,” *Discrete and Computational Geometry* 21, 557–568.
- von Stengel, B. (2002), “Computing equilibria for two-person games,” R. J. Aumann and S. Hart, eds., *Handbook of Game Theory, Vol. 3*, 1723–1759, Amsterdam: North-Holland.
- Ziegler, G. M. (1995), *Lectures on Polytopes*. Graduate Texts in Mathematics, Vol. 152, New York: Springer.