A Stable-Set Bound and Maximal Numbers of Nash Equilibria in Bimatrix Games

Constantin Ickstadt* Thorsten Theobald*
Bernhard von Stengel†

22 September 2025

Abstract

Quint and Shubik (1997) conjectured that a non-degenerate $n \times n$ game has at most $2^n - 1$ Nash equilibria in mixed strategies. The conjecture is true for $n \leq 4$ but false for $n \geq 6$. We answer it positively for the remaining case $n = 5$, which had been open since 1999. The problem can be translated to a combinatorial question about the vertices of a pair of simple n-polytopes with $2n$ facets. We introduce a novel obstruction based on the index of an equilibrium, which states that equilibrium vertices belong to two equal-sized disjoint stable sets of the graph of the polytope. This bound is verified directly using the known classification of the 159,375 combinatorial types of dual neighborly polytopes in dimension 5 with 10 facets. Non-neighborly polytopes are analyzed with additional combinatorial techniques where the bound is used for their disjoint facets.

1 Introduction

A bimatrix game is a two-player game in strategic form, a basic model of game theory. Its central solution concept is *Nash equilibrium*, which is a pair of randomized strategies that are optimal against each other. Nash equilibria of bimatrix games are best understood as certain combinatorial properties of two polytopes derived from the payoff matrices of the two players, called *best-response polytopes*. The possible *number* of Nash equilibria is of interest for algorithms that find one or all Nash equilibria of a game. Most insights in this area have been found using properties and constructions of polytopes. The present paper continues in this tradition with

*FB 12 – Institut für Mathematik, Goethe-Universität, Robert-Mayer-Str. 10, 60325 Frankfurt am Main, Germany. Email: ickstadt@math.uni-frankfurt.de, theobald@math.uni-frankfurt.de

†Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom. Email: b.von-stengel@lse.ac.uk

additional new techniques. We answer a long-open problem: We show that a non-degenerate 5×5 bimatrix game has at most 31 Nash equilibria (degenerate games are non-generic and may have infinitely many Nash equilibria). Throughout this paper, all the games that we talk about are assumed to be non-degenerate, which implies that their best-response polytopes are simple (defined by generic linear inequalities).

Quint and Shubik (1997) conjectured that the maximal number of Nash equilibria in an $n \times n$ game is $2^n - 1$. Their evidence for this, not based on polytopes, was that this is the number of Nash equilibria in the "coordination game" where each player's payoff matrix is the identity matrix. The corresponding best-response polytope is the n-cube. For the non-degenerate games that we consider, every equilibrium strategy of a player is a separate polytope vertex, so that the number of vertices is an upper bound on the number of Nash equilibria. This implies the Quint-Shubik conjecture for $n \leq 3$, and it was proved for $n = 4$ by Keiding (1997) and McLennan and Park (1999). However, the conjecture is false for $n \geq 6$. Using dual cyclic polytopes, which for $n \times n$ games have proportional to about $2.6^n/\sqrt{n}$ many vertices, von Stengel (1999) constructed games with about $2.4^n/\sqrt{n}$ many Nash equilibria (slightly fewer for odd n), and a 6×6 game with 75 Nash equilibria as the first counterexample to the Quint-Shubik conjecture. The remaining case $n = 5$ had since been open, which we solve in this paper.

We develop new techniques to obtain our result. Briefly, they are based on existing lists of combinatorial types of simple polytopes in dimensions 4 and 5 with 8, 9, and (for dual-neighborly polytopes) 10 facets, and a novel *obstruction* that limits the number of equilibrium vertices in certain face configurations. A basic version of this obstruction is that a triangle as a polytope face can have at most two equilibrium vertices, which "loses" one possible equilibrium vertex per triangle. This basic obstruction has been used by Keiding (1997) to show that 4×4 games have at most 15 Nash equilibria, using the list by Grünbaum and Sreedharan (1967) of the 37 combinatorial types of simple polytopes in dimension 4 with up to 8 facets (3 of these 37 polytopes are *dual-neighborly*, that is, any two facets intersect). For 5×5 games, no such list is known. We make use of two lists by Firsching (2017), namely of the 159,375 dual-neighborly simple polytopes in dimension 5 with 10 facets, and of all 1,142 simple polytopes in dimension 4 with 9 facets, but need further considerations.

Our novel obstruction is based on the concept of an equilibrium *index* due to Shapley (1974). It is convenient to consider as an equilibrium also the "artificial equilibrium", given by the pair of all-zero vectors that are vertices of the two best-response polytopes and fulfill the equilibrium condition, but cannot be rescaled as randomized strategies to represent a Nash equilibrium. The important *parity argument*, based on the algorithm by Lemke and Howson (1964) and already exploited by Keiding (1997), states that the number of equilibria is even (and hence the number of Nash equilibria is odd). The index property strengthens this further: Half of the equilibria have index +1, and the other half (including the artificial equilibrium) have index −1. In a best-response polytope, equilibrium vertices connected by an edge have opposite index (Lemma 5 below). Hence, in the graph of each best-response polytope (as defined by its vertices and edges), any two index +1 equilibrium vertices are non-adjacent and form a *stable set* (also called independent set), and the same applies to the index −1 equilibrium vertices. The maximum size of two disjoint stable sets of equal size is therefore an upper bound on the number of equilibria (Theorem 3), which is our new *stable-set bound*. Because the polytope graphs are small, this bound on equilibrium numbers can be computed quickly.

For dual-neighborly polytopes, the bound of 32 equilibria (and hence 31 Nash equilibria) for a 5×5 game can then be verified by computer calculations using the list by Firsching (2017) of these 159,375 polytopes. In fact, this result was already obtained in an earlier unpublished attempt by Vissarion Fisikopoulos and Bernhard von Stengel using the list of 159,750 neighborly oriented matroids by Miyata and Padrol (2015), and using the simpler obstruction of disjoint simplices as faces, each of which can only contain two equilibrium vertices. Oriented matroids are more general than polytopes (and hence the bound is valid), but only slightly so; Firsching (2017) showed that all but 375 of them are realizable as polytopes.

The challenge solved in this paper is when a best-response polytope in a 5×5 game is not dual-neighborly and therefore has two disjoint facets. Such a polytope has at most 40 vertices (dual-neighborly polytopes have 42 vertices). There is as yet no list of all 5-dimensional simple polytopes with 10 facets. The two disjoint facets are 4-dimensional simple polytopes with up to 8 facets. Each of these facets can be seen as a 4×4 game, which has at most 16 equilibria. However, this does not complete the proof because there may be further equilibrium vertices on neither facet.

A crucial observation is that the stable-set bound applies also to the equilibrium vertices on a *facet* of a best-response polytope (Theorem 4, proved in a nicely combined use of polytopes and game theory). Each of the two disjoint facets has one of the 37 combinatorial types known from Grünbaum and Sreedharan (1967). By the new facet-stable-set bound, 35 of these types have at least 4 non-equilibrium vertices, and the entire polytope has most 40 vertices. The remaining two combinatorial types are the 4-cube and what we call the *semi-cube* (Figure 1). In order to have more than 32 equilibria, at least one of the two disjoint facets is therefore a 4-cube or semi-cube; furthermore, the corresponding facet of the other best-response polytope needs 9 facets as a polytope in dimension 4. Again, an exhaustive computer search of these 1,142 simple polytopes, and the stable-set bound, give combinatorial properties (Theorem 7) that we use to show that more than 32 equilibria are not possible. This completes the proof that 5×5 games have no more than 31 Nash equilibria.

For the history of using polyhedra for studying bimatrix games see von Stengel (2002). Cyclic polytopes have been used by Savani and von Stengel (2006, 2016) to show that Lemke-Howson paths (see Section 4) may be exponentially long. The number of Nash equilibria in N-player games has been studied by McKelvey and McLennan (1997) and recently Vujić (2022). The more restrictive two-player rank-1 games were introduced by Kannan and Theobald (2010) and shown to have exponentially many equilibria by Adsul, Garg, Mehta, Sohoni, and von Stengel (2021). For semi-definite games, equilibrium numbers were considered by Ickstadt, Theobald, and Tsigaridas (2024) and Ickstadt, Theobald, Tsigaridas, and Varvitsiotis (2023), which are relevant for quantum games (Bostanci and Watrous (2021)). The Quint-Shubik bound does hold for tropical games (Allamigeon, Gaubert, and Meunier (2023)).

Section 2 states preliminaries on bimatrix games and polytopes. Best-response polytopes are recalled in Section 3. A concise definition of Lemke-Howson paths and the equilibrium index for our purposes is given in Section 4. The stable-set bound is proved in Section 5. Based on these concepts, Section 6 gives the main proof that non-degenerate 5×5 games have at most 31 Nash equilibria.

2 Preliminaries on bimatrix games and polytopes

All matrices have real entries. The transpose of a matrix B is denoted by $B^\top$. All vectors are column vectors. We treat vectors and scalars as matrices, so scalars are multiplied from the right with a column vector and from the left with a row vector. The ith component of a vector x is denoted by x_i. The all-zero vector is denoted by $\mathbf{0}$ and the all-one vector by $\mathbf{1}$, their dimension depending on the context. Let e_i be the i-th unit vector. Inequalities between vectors such as $x \geq \mathbf{0}$ are assumed to hold for all components. We let $[k] = \{1, \ldots, k\}$ for any positive integer k.

An $m \times n$ *bimatrix game* is a pair of $m \times n$ matrices (A, B). The m rows and n columns are the *pure strategies* of player 1 and 2, respectively. The players simultaneously each choose a pure strategy, with the resulting entry of A as payoff to player 1 and of B to player 2. A *mixed strategy* is a probability vector over the player's pure strategies. Players are interested in maximizing their expected payoff, given by $x^\top Ay$ to player 1 and $x^\top By$ to player 2 if their mixed strategies are x and y. A mixed strategy is a *best response* to the other player's mixed strategy if it maximizes the player's expected payoff over all possible mixed strategies. A mixed-strategy pair of mutual best responses is called a *Nash equilibrium*.

For a mixed strategy, say x of player 1, its *support* $\mathsf{supp}(x)$ is the set of pure strategies that are played with positive probability. It is easy to see (Nash (1951), p. 287) that x is a best response to a mixed strategy y of player 2 if and only if every pure strategy in $\mathsf{supp}(x)$ is a best response to y, that is, if the following *best-response condition* holds:

$$x_i > 0 \quad \Rightarrow \quad (Ay)_i = u = \max_{k\in[m]} (Ay)_k \qquad (i \in [m]). \tag{1}$$

If (1) holds, then u is also the best-response payoff $x^\top Ay$ of the mixed strategy x against the mixed strategy y. The analogous condition holds for a mixed strategy y of player 2 being a best response to a mixed strategy x of player 1.

A bimatrix game is *non-degenerate* if for any mixed strategy z of a player there are at most $|\mathsf{supp}(z)|$ many pure strategies that are best responses to z. For a comprehensive list of equivalent definitions see von Stengel (2021, thm. 14). Generic games (A, B) (where the entries of A and B are chosen from some continuous distributions) are non-degenerate. We only consider non-degenerate games.

We use the following concepts from convex geometry (see also Grünbaum (2003); Ziegler (1995); Joswig and Theobald (2013)). An *affine combination* of points $z^1, \dots, z^k$ in some Euclidean space is of the form $\sum_{i=1}^k z^i\lambda_i$ where $\lambda_1, \dots, \lambda_k$ are reals with $\sum_{i=1}^k \lambda_i = 1$. It is called a *convex combination* if $\lambda_i \geq 0$ for all i. A set of points is *convex* if it is closed under forming convex combinations. The *convex hull* of a set of points is the smallest convex set that contains all these points. Given points are *affinely independent* if none of these points is an affine combination of the others. A convex set has *dimension* d if and only if it has $d+1$, but no more, affinely independent points.

A *polyhedron* P is a subset of $\mathbb{R}^d$ defined by finitely many linear inequalities, that is, $P = \{z \in \mathbb{R}^d \mid Cz \leq q\}$ for some C in $\mathbb{R}^{\ell\times d}$ and some q in $\mathbb{R}^\ell$. A *face* F of the polyhedron P is obtained by converting some of its inequalities into equalities, that is, if for some $S \subseteq \{1, \dots, \ell\}$

$$F = \{z \in \mathbb{R}^d \mid Cz \leq q,\ (Cz)_i = q_i \text{ for } i \in S\}\,. \tag{2}$$

A *polytope* is a bounded polyhedron. A *simplex* is the convex hull of affinely independent points.

Suppose the polytope P has dimension d (also called a d-polytope; as a subset of $\mathbb{R}^d$ it is then called *full-dimensional*). Then a face F of P is called a *facet* if it has dimension $d-1$, a *ridge* if it has dimension $d-2$, an *edge* if it has dimension 1, and a *vertex* if it has dimension 0. A vertex is a singleton $\{v\}$ and also identified with v.

A d-polytope P is called *simple* if every vertex of P belongs to exactly d facets. For the following more general statement see Ziegler (1995, prop. 2.16(iv)).

Lemma 1. *In a simple d-polytope P, every non-empty face of P of dimension k is a subset of exactly $d-k$ facets.*

The faces of a polytope, partially ordered by inclusion, form a lattice (see Ziegler (1995)). Two polytopes have the same *combinatorial type* if their face lattices are isomorphic. We denote the set of combinatorial types of d-dimensional simple polytopes with k facets by $\mathcal{P}_k^d$.

A simple d-polytope is *dual-neighborly* if any $\lfloor d/2 \rfloor$ of its facets have a non-empty intersection. The dimension d of the polytopes that we consider is 4 or 5, so they are dual-neighorly if any two facets have a non-empty intersection, which is then a ridge by Lemma 1.

Lemma 2. *For a given dimension and number of facets, only simple polytopes that are dual-neighborly have the largest numbers of vertices. The maximal number of vertices of a polytope in* $\mathcal{P}^4_8, \mathcal{P}^4_9, \mathcal{P}^5_9, \mathcal{P}^5_{10}$ *is 20, 27, 30, 42, respectively.*

Proof. See Grünbaum (2003, p. 174) and (3.2) in Joswig and Theobald (2013). □

The *graph* of a polytope is defined by the vertices and edges of the polytope, where in the graph an edge is considered as the unordered pair of its endpoints, written as uv if the endpoints are u and v; then u and v are also called *adjacent*. The *degree* of a vertex in a graph is the number of vertices it is adjacent to.

3 Best-response polytopes

For more details on the following construction of "best-response polytopes" see von Stengel (1999, 2002, 2022). Let (A, B) be an $m \times n$ bimatrix game and consider the polyhedra

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

We assume that P and Q are polytopes, which holds if and only if the best-response payoff to any mixed strategy of the other player is always positive (von Stengel, 2022, lemma 9.9). A sufficient condition for this is that A and $B^\top$ are nonnegative and have no zero column. This is not restrictive since adding a constant to every entry of a payoff matrix does not change best responses.

With the exception of $\mathbf{0}$, every x in P represents a mixed strategy of player 1 after re-scaling it to $x\frac{1}{\mathbf{1}^\top x}$ as a probability distribution. Similarly, every y in $Q \setminus \{\mathbf{0}\}$ represents a mixed strategy $y\frac{1}{\mathbf{1}^\top y}$ of player 2. We always consider x in $P \setminus \{\mathbf{0}\}$ and y in $Q \setminus \{\mathbf{0}\}$ as "mixed strategies" with this implied re-scaling.

Suppose that $j \in [n]$ and consider the last n inequalities $B^\top x \leq \mathbf{1}$ in the definition of P in (3). If the jth inequality is *binding* (holds as equality), that is, $(B^\top x)_j = 1$, then j is a pure best response of player 2 to x. That is, the right-hand side 1 in every inequality in $B^\top x \leq \mathbf{1}$ represents the best-response payoff to x, re-scaled as 1, if at least one of these inequalities is binding; the actual payoff is the re-scaling factor $\frac{1}{\mathbf{1}^\top x}$. Similarly, a binding inequality $(Ay)_i = 1$ for $i \in [m]$ among the inequalities $Ay \leq \mathbf{1}$ in the definition of Q means that i is a pure best response to y.

If $x = \mathbf{0}$ then $B^\top x < \mathbf{1}$ and if $y = \mathbf{0}$ then $Ay < \mathbf{1}$. Hence, P and Q contain the affinely independent vectors $\mathbf{0}$ and $e_i\varepsilon$ for the unit vectors e_i (in $\mathbb{R}^m$ respectively $\mathbb{R}^n$) for sufficiently small positive ε and are full-dimensional.

We introduce *labels* in $[m+n]$ to uniquely identify the pure strategies of player 1 as $1, \ldots, m$ and of player 2 as $m+1, \ldots, m+n$. The inequalities in (3) are written in the order of these labels. For a point $x \in P$ or $y \in Q$, *its labels* are the labels of the *binding* inequalities for that point. That is, for a pure strategy i in $[m]$ of player 1, the point x has label i if $x_i = 0$, and y has label i if i is a best response to the mixed strategy y (if $y = \mathbf{0}$ then $Ay < \mathbf{1}$ and y cannot have label i). Similarly, for a pure strategy j in $[n]$ of player 2, the mixed strategy x has label $m+j$ if j is a best response to x, and y has label $m+j$ if $y_j = 0$. We also consider these labels if $x = \mathbf{0}$ (then x has labels $1, \ldots, m$) or $y = \mathbf{0}$ (then y has labels $m+1, \ldots, m+n$).

By the best-response condition (1), (x, y) is a Nash equilibrium of (A, B) if and only if every pure strategy of a player is a best response or played with probability zero (or both). That is, every label in $[m+n]$ has to appear as a label of x or y.

Definition 1. Consider a bimatrix game (A, B) and assume that P and Q in (3) are polytopes, called the *best-response polytopes* for the game. Let $(x, y) \in P \times Q$ be such that every element of $[m+n]$ appears as a label of x or y. Then (x, y) is called *completely labeled* and an *equilibrium* of the game. This includes the *artificial equilibrium* $(\mathbf{0}, \mathbf{0})$. Any other equilibrium, with x and y re-scaled as mixed strategies, is a Nash equilibrium of (A, B).

Note that if $x \in P \setminus \{\mathbf{0}\}$, then x is missing at least one label $i \in [m]$, which in an equilibrium (x, y) has to appear as a label of y and therefore $y \neq \mathbf{0}$. Hence $(\mathbf{0}, \mathbf{0})$ is the only equilibrium (x, y) where $x = \mathbf{0}$ or $y = \mathbf{0}$.

The Nash equilibria (x, y) of a bimatrix game are exactly the completely labeled pairs of points in $P \times Q \setminus \{(\mathbf{0}, \mathbf{0})\}$, even for a degenerate game. For a non-degenerate game, they are always pairs of vertices.

Lemma 3. *For a bimatrix game (A, B) with best-response polytopes P and Q, the following are equivalent:*

(a) *The game is non-degenerate.*

(b) *No point in P has more than m labels and no point in Q has more than n labels.*

(c) *P and Q are simple polytopes, and for both polytopes any redundant inequality (which can be omitted without changing the polytope) is never binding.*

Furthermore, assume the game is non-degenerate and let $(x, y) \in P \times Q$ be an equilibrium.

(d) *Then x is a vertex of P and y is a vertex Q, and x has exactly m labels and y has exactly the other n labels.*

Proof. Conditions (b) and (c) are theorem 14 (b) and (h) of von Stengel (2021). Because there are only $m+n$ labels in total, (b) implies (d). □

In Lemma 3(c), redundant inequalities refer to the description of the polytope, which may not affect whether the polytope as a set is simple. It implies, for example, that there are no duplicate inequalities. A never-binding inequality represents a pure strategy that is never a best response; this does not affect non-degeneracy and may exist in a generic game, but we can omit such a strategy from the game. In the non-degenerate games that we consider, every binding inequality defines a facet that has its own label.

The transition from bimatrix games to polytopes can also be reversed. For projective transformations see Ziegler (1995), of which (4) is a simple form.

Lemma 4. *Consider a simple m-polytope P and a simple n-polytope Q, each with $m+n$ facets labeled $1,\ldots,m+n$. Let (x_0,y_0) be a completely labeled vertex pair of $P\times Q$. Then there is a game (A',B') with best-response polytopes P' and Q' such that:*

(a) *There are affine bijections $P\to P'$ and $Q\to Q'$ with a corresponding permutation of $[m+n]$ that maps the facet labels of P and Q to those of P' and Q', with (x_0,y_0) mapped to $(\mathbf{0},\mathbf{0})$ in $P'\times Q'$, such that all completely labeled vertex pairs of $P\times Q$ are in bijection to those of $P'\times Q'$.*

(b) *By applying, if necessary, bijective projective transformations to the polytopes P' and Q' that leave their face lattices unchanged, all entries of A' and B' are positive.*

Proof. For (a) see von Stengel (1999, prop. 2.1). For (b), note that e.g. Q', given by $\{y'\in\mathbb{R}^n \mid A'y'\le\mathbf{1},\ y'\ge\mathbf{0}\}$, may be a polytope even when A' has non-positive entries. In that case add a sufficiently positive constant α to all entries of A' to yield an all-positive matrix $\hat{A}=A'+\mathbf{1}\alpha\mathbf{1}^\top$, and let $\hat{Q}=\{\hat{y}\in\mathbb{R}^n \mid \hat{A}\hat{y}\le\mathbf{1},\ \hat{y}\ge\mathbf{0}\}$. We claim that

$$y'\mapsto\hat{y}=y'\frac{1}{1+\alpha\,\mathbf{1}^\top y'} \tag{4}$$

is a bijective projective transformation $Q'\to\hat{Q}$ that preserves all binding inequalities and therefore the face lattices of Q' and $\hat{Q}$. In (4), $\mathbf{0}$ in Q' is mapped to $\mathbf{0}$ in $\hat{Q}$, and for $y'\ne\mathbf{0}$ the mapping in (4) is obtained by mapping y' to $(\overline{y},u)$ with a mixed strategy $\overline{y}=y'u$ where $u=\frac{1}{\mathbf{1}^\top y'}$, then $(\overline{y},u)$ to $(\overline{y},u+\alpha)$, and this is mapped to $\hat{y}=\overline{y}\frac{1}{u+\alpha}$. To see that binding inequalities are preserved in Q' and $\hat{Q}$, clearly $y'_j=0$ if and only if $\hat{y}_j=0$ for $j\in[n]$. Furthermore, $\hat{A}\hat{y}\le\mathbf{1}$ if and only if $\hat{A}y'\le\mathbf{1}(1+\alpha\mathbf{1}^\top y')$, which for the ith row $a_i^\top$ of A' means $(a_i^\top+\alpha\mathbf{1}^\top)y'\le 1+\alpha\mathbf{1}^\top y'$ or equivalently $a_i^\top y'\le 1$, which is the ith inequality in $A'y'\le\mathbf{1}$.

A similar transformation changes B' to an all-positive matrix $\hat{B}$ and P' to a "projectively equivalent" polytope $\hat{P}$. ☐

Lemma 4 allows changing the labels of P and Q such that a Nash equilibrium and the artificial equilibrium exchange roles, which we will make use of in Theorem 4.

4 Equilibria come in pairs of opposite index

Our "stable-set bound" is based on the fact that equilibria come in pairs of opposite *index*, which we explain in this section. The algorithm by Lemke and Howson (1964) starts at one equilibrium by following a path of alternating edges in the best-response polytopes P and Q in (3) where one of the labels is allowed to be *missing*, until it terminates at another equilibrium. The endpoints of the Lemke-Howson paths are equilibria of opposite index. The most concise description, given here, considers the product polytope $Z = P \times Q$. For more details see, for example, von Stengel (2002, 2021, 2022).

Definition 2. A *labeled polytope* Z is a simple d-polytope of the form

$$Z = \{z \in \mathbb{R}^d \mid -z \leq \mathbf{0},\ Cz \leq \mathbf{1}\} \tag{5}$$

for some $k \times d$ matrix C without redundant inequalities where each binding inequality has a label in $[d]$, such that: For $i \in [d]$, the ith inequality $-z_i \leq 0$ in $-z \leq \mathbf{0}$ has label i, and for $j \in [k]$, the jth inequality $(Cz)_j \leq 1$ in $Cz \leq \mathbf{1}$ has some given label $\ell(j)$ in $[d]$. A point z in Z that has all labels in $[d]$ is called *completely labeled*, and *h-almost completely labeled* if it has all labels in $[d] \setminus \{h\}$, for some $h \in [d]$.

We will apply this to $Z = P \times Q$ where in Definition 2 we take $d = m + n$, $z = (x, y)$, and

$$C = \begin{bmatrix} 0 & A \\ B^\top & 0 \end{bmatrix}, \quad k = d, \quad \text{and} \quad \ell(j) = j \quad \text{for } j \in [k]\,. \tag{6}$$

Theorem 1. *Any labeled polytope has an even number of completely labeled points.*

Proof. Let Z be the labeled polytope as in Definition 2. Because no inequality is redundant, every facet is defined by a binding inequality and has a unique label. Every vertex of Z belongs to exactly d facets and is incident to exactly d edges, obtained by allowing one of the binding inequalities to be non-binding (visualized as "moving away" from the respective facet).

Consider a fixed label h in $[d]$ that is allowed to be "missing", and all the points in Z that are h-almost completely labeled. Those points that have each label in $[d] \setminus \{h\}$ exactly once define edges of Z except for the endpoints of these edges. Their endpoints, which are vertices of Z, are either completely labeled, or are missing label h and therefore have a *duplicate* label. Any such vertex with the duplicate label belongs to two facets that have that label, and is incident to *two* h-almost completely labeled edges, which belong to only one of the two facets.

Hence, the h-almost completely labeled points of Z are vertices and edges of Z that define a subgraph of the graph of Z where every vertex has either degree two (with a duplicate label but missing label h) or degree one (if the vertex is completely labeled). Such a graph consists of paths and cycles, and every path has two endpoints, which are the completely labeled vertices, so their number is even. □

Completely labeled vertices (except **0**) of a labeled polytope also represent Nash equilibria of a game, namely of the $d \times k$ game $(U, C^\top)$ where $U = [e_{\ell(1)} \cdots e_{\ell(k)}]$ is composed of the unit vectors for the labels of the rows of C. This simplifies certain constructions of games by using only a single labeled polytope, see Savani and von Stengel (2016). Applied to C in (6), U is the $d \times d$ identity matrix, whose Nash equilibria (z, w) correspond to the *symmetric* Nash equilibria (z, z) of the symmetric game $(C, C^\top)$, which we seek when setting $z = (x, y)$ to find the Nash equilibria (x, y) of (A, B).

Definition 3. Consider a labeled d-polytope Z as in Definition 2 and a completely labeled vertex z with its d binding inequalities written as $a_i^\top z = \beta_i$ for each label $i \in [d]$. That is, if the ith inequality in $-z \leq \mathbf{0}$ is binding then $a_i^\top z = \beta_i$ stands for $-z_i = 0$ with a_i as the negative $-e_i$ of the ith unit vector e_i, and if $i = \ell(j)$ then the jth inequality in $Cz \leq \mathbf{1}$ has label i and $a_i^\top z = \beta_i$ stands for $(Cz)_j = 1$ with $a_i^\top$ as the jth row of C. Writing the facet normal vectors $a_1, \ldots, a_d$ in the order of their labels as a matrix with determinant $|a_1 \cdots a_d|$, we define the *index* of z as the sign of this determinant (times -1 if d is even),

$$\mathsf{index}(z) = (-1)^{d+1} \mathsf{sign}(|a_1 \cdots a_d|). \tag{7}$$

Because the polytope Z is simple, the normal vectors a_i of the incident facets of a vertex are linearly independent and the determinant is non-zero. In (7), the minus sign for even d has the purpose of ensuring that the artificial equilibrium **0** has index -1, because the determinant of the negative of the $d \times d$ identity matrix is $(-1)^d$.

The following theorem is due to Shapley (1974). For accessibility in the present context, we outline a streamlined version of its proof due to von Stengel (2021, thm. 13).

Theorem 2. *For a labeled polytope, the endpoints of any h-almost completely labeled path are completely labeled vertices of opposite index. Hence, half of the completely labeled vertices have index $+1$ and the other half, including* **0**, *have index -1.*

Proof. The proof extends that of Theorem 1 by *orienting* each h-almost completely labeled edge to give a consistent orientation of the resulting paths. The orientation of the edge is from its endpoint with negative index to the endpoint with positive index, where the definition of index is extended to h-almost completely labeled vertices as follows.

Suppose for simplicity that $h = 1$ and that d is odd so that the signs of index and determinant coincide (for even d they are opposite and the same argument applies). Consider a completely labeled vertex z of negative index with normal facet vectors $a_1, \ldots, a_d$ in the order of their labels, that is, $|a_1 \cdots a_d| < 0$. Let the 1-almost completely labeled edge connect z to w where the facet normal vector a_1 is replaced by b_g with label g in $[d]$. Then $|b_g\, a_2 \cdots a_d| > 0$, because the determinant of the endpoints of any edge of a simple polytope have opposite sign when replacing the changed facet normal vector and leaving the unchanged facet normal vectors in place (von Stengel, 2021, lemma 12).

If g is the missing label 1 then we are done, that is, w is the endpoint of the path and has positive determinant. Otherwise, label g is duplicate. We now exchange b_g with a_g in writing down the determinant, which changes sign, that is, $|a_g\, a_2 \cdots b_g \cdots a_d| < 0$. Then a_g is again the normal vector of the facet that the next edge on the path moves away from, and we proceed as before. In that way, all the h-almost completely labeled edges are oriented in the same direction on the path. □

5 The stable-set bound

In the previous section, we have considered a general labeled polytope Z to explain equilibria as endpoints of oriented paths on Z defined by the points of Z that have all labels except a missing label h. For bounds on the number of equilibria of an $m \times n$ game (A, B), it is better to consider the product structure $Z = P \times Q$ with the m-polytope P and n-polytope Q in (3), because they have much fewer vertices than a general simple polytope Z of dimension $m + n$. We assume throughout that P and Q are the simple polytopes in (3), with their $m + n$ inequalities having the respective labels in $[m + n]$ when they are binding and thus defining a facet of the polytope.

By Lemma 3(d), any equilibrium is a vertex pair (x, y) of $P \times Q$. We then call x an *equilibrium vertex* of P and the vertex y of Q its *partner* (and vice versa). The partner y of x is *unique* because it is defined by the unique set of labels in $[m + n]$ (and the corresponding facets of Q) that are missing from x, by Lemma 3(d).

The index of an equilibrium (x, y) cannot be told from the equilibrium vertex x and its position in P alone (other than by uniquely identifying its partner y). However, the following lemma applies.

Lemma 5. *Any two equilibrium vertices of P that are connected by an edge belong to equilibria of opposite index.*

Proof. Let x and x' be two equilibrium vertices of P with partners y and y' in Q, respectively, and let xx' be an edge in the graph of P. Then x and x' differ by exactly one label, where the set of labels of x is $\{h\} \cup K$ and of x' is $\{g\} \cup K$, with $|K| = m - 1$ and $h \neq g$. Because (x, y) and (x', y') are completely labeled, the set of labels of y is therefore $\{g\} \cup L$ and of y' is $\{h\} \cup L$, with $L = [m + n] \setminus (K \cup \{h, g\})$.

Hence, yy' is also an edge of the graph of Q. In the graph of the polytope $P \times Q$ (which is the "product graph" of the graphs of P and Q), the two equilibria are therefore the endpoints of the h-almost completely labeled path given by the sequence of vertex pairs $(x, y) - (x', y) - (x', y')$ and therefore have opposite index by Theorem 2. □

For any undirected graph G, a *stable set* is a set of vertices no two of which are adjacent.

Theorem 3 (Stable-set bound)**.** *Let G be the graph of P and S_1 and S_2 be two disjoint stable sets of G of equal size $|S_1|$, and let this size be maximal among all such pairs. Then the game has at most $2|S_1|$ equilibria.*

Proof. By Theorem 2, half of all equilibria have positive index. No two of their equilibrium vertices in P are adjacent in G by Lemma 5, so they form a stable set S of G. The same applies to the equilibria of negative index, whose equilibrium vertices in P form a stable set of G disjoint from S of the same size. □

A *clique* of a graph is a set of vertices every two of which are adjacent. Disjoint cliques of size three or larger provide a bound that is weaker than the stable-set bound, but may match it (in particular when reduced by 1 to obtain an even number using Theorem 1).

Corollary 1 (Disjoint-clique bound)**.** *Suppose the graph of P has V vertices and k pairwise disjoint cliques of sizes $c_1, \ldots, c_k$. Then every clique contains at most two equilibrium vertices, that is, the game has at most $V - \sum_{i=1}^{k}(c_i - 2)$ many equilibria. This bound is at least as large as the stable-set bound.*

Proof. As in the proof of Theorem 3, no two equilibrium vertices for equilibria of the same index can belong to a clique because they would be adjacent, in contradiction to Lemma 5. □

Corollary 1 has been known and used earlier by Keiding (1997) for 4×4 games and triangles as cliques. It can be proved without the concept of an equilibrium index, as follows, here for general $m \times n$ games.

Alternative Proof of Corollary 1. Consider a clique of the graph of P of size c. The convex hull F of its vertices is a simplex of dimension $c - 1$ and therefore the intersection of $m - c + 1$ facets of P with as many labels, by Lemma 1. Considered as a $(c - 1)$-polytope, F has c facets, each of which is obtained by intersecting F with an additional facet of P that provides an additional label. The labels of all vertices of F belong therefore to a set of size $m + 1$. Of these, any equilibrium vertex needs a partner in Q that has the missing $m + n - (m + 1)$, that is, $n - 1$ labels, all of which belong therefore to a face of Q of dimension 1, which is an edge. An edge has only two endpoints, so at most two vertices in the clique can be equilibrium vertices. □

The stable-set bound in Theorem 3 is our main tool for proving that 5×5 games have at most 32 equilibria. It is enhanced by the corresponding property for the equilibrium vertices on a facet of P.

Theorem 4 (Facet-stable-set bound). *Let $m \geq 2$ and consider a non-degenerate $m \times n$ game (A, B) with positive payoff matrices and best-response polytopes P and Q. Consider a facet F with label ℓ of P and the graph G_F of F. Then the equilibrium vertices in G_F are contained in two disjoint stable sets in G_F of equal size.*

Proof. We can assume that F has at least one equilibrium vertex x_0 with partner y_0 in Q. By assumption, x_0 has label ℓ. If $\ell \notin [m]$ then $\ell = m + j$ for some pure strategy $j \in [n]$ of player 2. Then we use the bijections in Lemma 4 to map (x_0, y_0) to a new polytope pair $P' \times Q'$ such that the Nash equilibrium (x_0, y_0) of (A, B) maps to the artificial equilibrium $(\mathbf{0}, \mathbf{0})$ of a new game (A', B') and therefore ℓ is in bijection to a strategy of player 1. Furthermore, A' and B' have positive entries by Lemma 4(b). We then consider (A', B') instead of (A, B), with the same number of equilibria for the two games, and ℓ (replaced by its image under the bijection) now denoting a strategy of player 1. The polytopes P and Q and facet F of P are also replaced accordingly.

In this (if needed, new) game, consider all the equilibria (x, y) such that x is a vertex of G_F. Because $x \in F$, these equilibria (including the artificial equilibrium $(\mathbf{0}, \mathbf{0})$) have label ℓ in $[m]$ with $x_\ell = 0$, that is, row ℓ is not played. All the Nash equilibria among them are also Nash equilibria of the smaller $(m - 1) \times n$ game obtained from (A, B) by deleting row ℓ from both A and B. (The smaller game may have additional Nash equilibria.) The smaller game has F as its best-response polytope for player 1, and Q with inequality ℓ omitted as the best-response polytope Q' for player 2.

The facet F is simple because P is simple. The polyhedron Q' is a polytope because all entries of A are positive. However, Q' may not be simple and then the smaller game may be degenerate. However, F has label ℓ and therefore the partners of the equilibrium vertices in F are not affected by omitting the inequality with label ℓ in Q and are the same vertices in Q'. (Alternatively, one could perturb A without changing the combinatorial structure of Q to create a simple polytope Q'.)

The edges in G_F are exactly the edges of the graph of P with endpoints in F. Theorem 3 applies to the smaller game with F instead of P, which proves the claim. □

Theorem 4 is not trivial because it is conceivable that there are, say, $k + 1$ equilibrium vertices in F belonging to equilibria of (A, B) of positive index, and k for negative index, but with a facet-stable-set bound of $2k$. The theorem precludes this possibility.

6 Five-by-five games

We show that the Quint-Shubik conjecture holds for 5×5 games.

Theorem 5. *Any non-degenerate* 5×5 *game has at most 31 Nash equilibria.*

The proof uses some case distinctions and relies on computer calculations.

In the following, we let $m = n = 5$ and consider the two best-response polytopes P and Q in (3), both of which are full-dimensional as shown in Section 3 and therefore of dimension 5. They are simple by Lemma 3. By their definition in (3), each of P and Q has at most $5 + 5 = 10$ facets. We can assume that P and Q have ten facets each. If, say, P has at most nine facets, then P has at most 30 vertices by Lemma 2 and Theorem 5 holds immediately.

We first deal with the case that at least one of the polytopes P or Q is dual-neighborly. Recall that we also consider the artificial equilibrium $(\mathbf{0}, \mathbf{0})$ in $P \times Q$ as an equilibrium.

Theorem 6. *Consider the best-response polytopes P and Q for a* 5×5 *game. If P or Q is dual-neighborly, then the game has at most* 32 *equilibria.*

Proof. Assume that P is dual-neighborly. The 159,375 combinatorial types of dual-neighborly simple 5-polytopes with 10 facets have been classified by Firsching (2017). For each such combinatorial type we have verified the stable-set bound of Theorem 3 as not exceeding 32 by computer calculations. Indeed, for each of those polytopes, the size of a largest stable set is at most 16, which implies that the stable-set bound is at most 32. The calculations were carried out in `SageMath` using the coordinate descriptions of the polytopes from the data in Firsching (2017). We have documented the data of our stable set computations with links to the respective Sage programs and input files in Ickstadt, Theobald, and von Stengel (2025). ☐

As mentioned in the introduction, Theorem 6 was already found earlier with the help of Vissarion Fisikopoulos using the list of 159,750 neighborly oriented matroids by Miyata and Padrol (2015) with the disjoint-clique bound of Corollary 1.

From now on we assume that P and Q are *non-neighborly*, that is, not dual-neighborly. Currently, no list of all polytopes (not only the dual-neighborly ones) in $\mathcal{P}^5_{10}$ is available. Even if that list existed, it is unclear if finding the stable-set bounds would be computationally feasible and if they would produce the desired bound.

Instead we study the combinatorial structure of non-neighborly polytopes P and Q. Because simple 5-polytopes with 10 facets are dual-neighborly if and only if they have 42 vertices (by Lemma 2), P and Q have at most 41 vertices. In fact, as simple 5-polytopes, P and Q have an even number of vertices, because every vertex in their graph has degree 5, and the number of odd-degree vertices in a graph is even (because the sum of degrees is twice the number of edges). The interesting numbers of vertices of P and Q for proving Theorem 5 are therefore 40, 38, 36, and 34. Because the number of equilibria is even by Theorem 1, it suffices to know that seven vertices in P cannot be equilibrium vertices. We call such vertices *obstruction vertices*.

Because P is non-neighborly, P has two facets that are disjoint. We study the facet-stable-set bound stated in Theorem 4 for these two facets. For the dual-neighborly 5-polytopes in Theorem 6, it was sufficient to find maximal single stable sets to bound the equilibrium numbers. For the four-dimensional facets, we need to consider pairs of disjoint stable sets as in Theorem 3.

Lemma 6. *The stable-set bound for pairs of disjoint stable sets is the optimal value of an integer linear program (ILP) in* $2V$ *binary variables if* P *has* V *vertices.*

Proof. Let G be the graph of the polytope P with, for simplicity, vertex set $[V]$ and edge set E. The maximum size of a stable set in G can be computed by the integer linear program

$$\begin{array}{lrcll} \text{maximize} & \sum_{i=1}^{V} x_i & & & \\ \text{subject to} & x_i + x_j & \leq & 1 & \text{for all } ij \in E\,, \\ & x & \in & \{0,1\}^V\,, & \end{array}$$

where the vector x is the incidence vector of a stable set. In order to compute the stable set bound, we consider the modified ILP in $2V$ variables

$$\begin{array}{lrcll} \text{maximize} & \sum_{i=1}^{V} x_i + \sum_{i=1}^{V} y_i & & & \\ \text{subject to} & x_i + x_j & \leq & 1 & \text{for all } ij \in E\,, \\ & y_i + y_j & \leq & 1 & \text{for all } ij \in E\,, \\ & x_i + y_i & \leq & 1 & \text{for all } i \in [V]\,, \\ & \sum_{i=1}^{V} x_i & = & \sum_{i=1}^{V} y_i\,, & \\ & x, y & \in & \{0,1\}^V\,. & \end{array}$$

Both x and y provide incidence vectors of stable sets. The conditions $x_i + y_i \leq 1$ enforce the disjointness of these stable sets and the condition $\sum_{i=1}^{V} x_i = \sum_{i=1}^{V} y_i$ ensures that the stable sets specified by x and y have the same size. Then the objective function gives the sum of the sizes of the two stable sets, or, equivalently, twice of size of either of them. □

As in the proof of Theorem 4, it does not matter if the labels of the two disjoint facets of P denote pure strategies of player 1 or player 2 or of both players. Hence, we assume that the two disjoint facets of P have labels 1 and 2 (representing the first two strategies of player 1, but this is irrelevant), and call the facets P_1 and P_2. Similarly, let Q_1 and Q_2 be the (not necessarily disjoint) facets of Q with labels 1 and 2, respectively.

Each facet P_1 and P_2 of P is a four-dimensional polytope with at most eight facets. Every other facet of P, as a 4-polytope, may have up to nine (three-dimensional) facets, obtained as intersections with the other nine facets of P.

A computer calculation of the facet-stable-set bound of Theorem 4 shows that 35 out of the 37 combinatorial types of polytopes in $\mathcal{P}_8^4$ have at least four obstruction vertices (which cannot be equilibrium vertices). Hence if both facets P_1 and P_2 belong to these 35 types, we "lose" at least four equilibria in each facet, and hence there are at most $40 - 4 - 4 = 32$ equilibria. The two other combinatorial types are:

(i) the 4-cube, and

(ii) a type with 17 vertices that has only 3 obstruction vertices, which we call the *semi-cube*.

These two types play a prominent role in the subsequent investigations. While the 4-cube is well-known, the semi-cube deserves additional explanations. Three of its facets are 3-cubes. The graph of the semi-cube is shown in Figure 1, where one can identify two disjoint stable sets of size 7. Each stable set contains two opposite corners from each of the of the square faces. The semi-cube has also three disjoint triangles, so the stable-set bound is 14 by the disjoint-clique bound of Corollary 1. In the classification of Grünbaum and Sreedharan (1967, p. 462), the semi-cube appears as number 26, whereas in the data by Firsching (2017) the semi-cube appears as number 23.

The following is a list of the 37 combinatorial types of polytopes in $\mathcal{P}_8^4$, with the numbering of Firsching (2017). Each pair (V, b) states the number V of vertices of the polytope and its stable-set bound b. The number $V - b$ of obstruction vertices is at least 4 except for the underlined semi-cube number 23 and cube number 24.

$$\begin{array}{lllll}
1:(17,12) & 2:(18,12) & 3:(18,14) & 4:(18,12) & 5:(19,14)\\
6:(17,12) & 7:(19,14) & 8:(16,10) & 9:(16,10) & 10:(16,12)\\
11:(17,12) & 12:(19,14) & 13:(18,12) & 14:(16,10) & 15:(15,10)\\
16:(20,14) & 17:(18,14) & 18:(19,14) & 19:(17,12) & 20:(20,16)\\
21:(20,16) & 22:(18,12) & \underline{23:(17,14)} & \underline{24:(16,16)} & 25:(17,12)\\
26:(17,12) & 27:(17,12) & 28:(15,10) & 29:(16,10) & 30:(16,12)\\
31:(15,10) & 32:(15,10) & 33:(15,8) & 34:(16,12) & 35:(14,8)\\
36:(14,8) & 37:(14,8) & & &
\end{array} \tag{8}$$

We also have to consider disjoint facets P_1 and P_2 of P that as 4-polytopes have fewer than eight facets. There are only five combinatorial types of polytopes in $\mathcal{P}_7^4$. Similar to the pairs (V, b) in (8), their numbers V of vertices and stable set bounds b are (via a computer calculation, using data from Firsching (2017))

$$(14,10), \quad (13,8), \quad (12,8), \quad (12,8), \quad (11,6)\,, \tag{9}$$

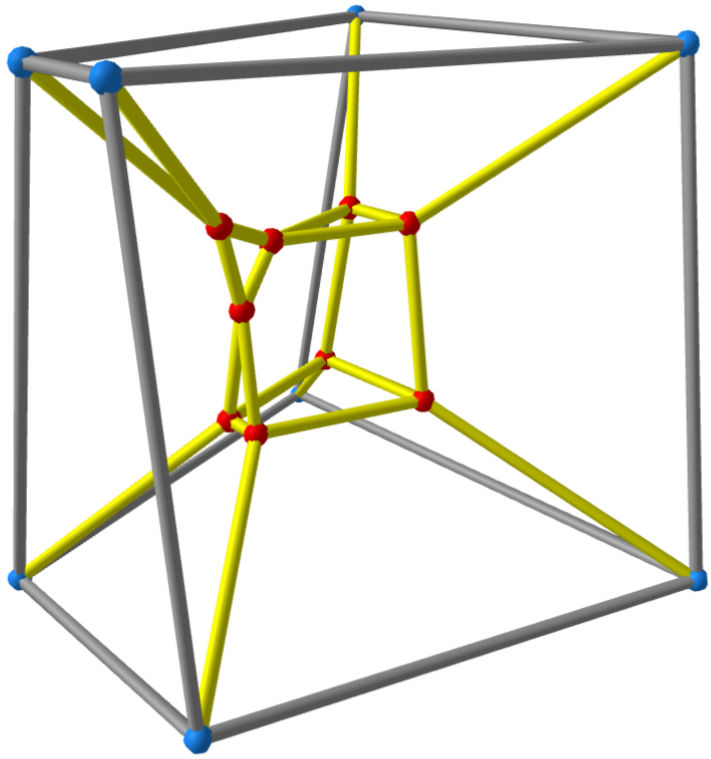

Figure 1: Schlegel diagram (see Ziegler (1995)) and graph of the semi-cube.

so each of them has also at least four obstruction vertices. A polytope in $\mathcal{P}^4_6$ has at most nine vertices and there are two possible types, with (V,b) given by $(9,6)$ and $(8,4)$. The only polytope in $\mathcal{P}^4_5$ is the 4-dimensional simplex with $(V,b) = (5,2)$.

Hence, only the following polytopes P are candidates for obtaining more than 32 equilibria:

(i) P_1 or P_2 is a 4-cube.

(ii) P_1 or P_2 is a semi-cube.

(iii) P_1 or P_2 is a 4-polytope with only six facets and $(V,b) = (9,6)$, or a simplex.

In every other case the polytope P has at least seven obstruction vertices and therefore no more than 32 equilibrium vertices.

The following theorem is useful for understanding where equilibrium vertices in Q can be found, which will be applied to facets Q_i of Q using the facet-stable-set bound of Theorem 4.

Theorem 7. *Computer calculations show the following for every polytope Q_i in $\mathcal{P}^4_9$.*

(a) *The stable-set bound of Q_i is at most 20.*

(b) *Q_i has at least three obstruction vertices.*

(c) *If Q_i has stable-set bound 20, then Q_i does not have a 3-cube as a facet.*

(d) *If Q_i has stable-set bound 20, then Q_i has at least five obstruction vertices.*

Proof. The computer calculations are based on the list of the 1,142 combinatorial types of polytopes in $\mathcal{P}_9^4$ of Firsching (2017). The stable-set bound is computed using Lemma 6. For (a), there are 39 types where the stable-set bound is 20. The calculations are documented in Ickstadt, Theobald, and von Stengel (2025). □

We will use these properties to prove Theorem 5. Case (iii) above is ruled out by the following lemma, because P_1 or P_2 do not have enough vertices.

Lemma 7. *Suppose P is non-neighborly and has at least one facet P_i with a stable-set bound of less than 14. Then the game has at most 32 equilibria.*

Proof. Let P_i have label i, and let Q_i be the facet of Q with label i. Every equilibrium (x, y) needs to have label i, where by assumption there are fewer than 14 equilibrium vertices x. In order to obtain 34 equilibria in total, the remaining equilibria must have more than 20 equilibrium vertices y in Q_i. By Theorem 7(a), such a polytope Q_i does not exist in $\mathcal{P}_9^4$. The same applies if the 4-polytope Q_i has eight or fewer facets, because then it has at most 20 vertices. □

We next treat case (ii) above.

Lemma 8. *Let P and Q be non-neighborly, with disjoint facets P_1 and P_2 of P, and let P_1 be a semi-cube. Then the game has at most 32 equilibria.*

Proof. Assume that the game has 34 or more equilibria (their number is even by Theorem 1), which will lead to a contradiction. The stable-set set bound for P_1 is 14 by (8), and P_1 needs to have 14 equilibrium vertices, by Lemma 7. The remaining 20 equilibrium vertices with label 1 belong to Q_1, which is in $\mathcal{P}_9^4$ by Theorem 7(a) (by (8), no polytope in $\mathcal{P}_8^4$ has 20 equilibrium vertices). Hence, Q_1 intersects with every other facet of Q, where the resulting ridge (a three-dimensional face both of Q and of any facet of Q) is never a 3-cube by Theorem 7(c). Hence, none of the facets of Q is a 4-cube. In order to obtain more than 32 equilibria, the two disjoint facets of Q are therefore semi-cubes, because combining a semi-cube with any of the other combinatorial types in (8) would give at least $3 + 4$ obstruction vertices. This argument with the polytopes exchanged implies that P_2 is also a semi-cube. Hence, also $Q_2 \in \mathcal{P}_9^4$.

By Theorem 7(d), both Q_1 and Q_2 have 20 equilibrium vertices and at least five obstruction vertices. By Theorem 7(d) and (b) and (8), each of the other eight facets of Q has at least three obstruction vertices. Counted by multiplicity over all ten facets (each vertex belongs to five facets), the number of obstruction vertices is therefore at least $(2 \cdot 5 + 8 \cdot 3)/5 = 34/5 > 6$. Hence, Q has at least seven obstruction vertices, but at most 40 vertices, and therefore at most 32 equilibria. □

Due to the symmetry of the role of the two polytopes, case (ii) and the previous lemma cover the case that at least one of the disjoint facet pairs of P or Q contains a semi-cube. The remaining case (i) is covered by the following lemma.

Lemma 9. *Let P and Q be non-neighborly and assume that one of the disjoint facets P_1, P_2 is 4-cube, and that one facet of Q is a 4-cube. Then the game has at most 32 equilibria.*

Proof. Suppose P_1 is a 4-cube. Assume that the game has 34 or more equilibria, which will lead to a contradiction. If P_1 had 15 or fewer equilibrium vertices, then the remaining equilibria have label 1 with their 19 equilibrium vertices in Q_1, which needs to have a stable-set bound of 20 but then has no 3-cube as a facet by Theorem 7(c). However, then Q has no 3-cube as a ridge, which contradicts Q having a 4-cube as a facet. Hence, *all* 16 vertices of P_1 are equilibrium vertices.

The 4-cube P_1 has a special combinatorial structure, which is isomorphic to that of the unit 4-cube $W = \{z \in \mathbb{R}^4 \mid 0 \leq z_i \leq 1 \text{ for } i \in [4]\}$. Each of the eight inequalities of W corresponds to a facet of P_1 with a label in $\{3, \dots, 10\}$. These facets of P_1 come in four pairs with labels g_i, h_i that correspond to the binding inequalities $z_i = 0$ and $z_i = 1$ in W for $i \in [4]$. Every vertex v of P_1 has exactly one label g_i or h_i for $i \in [4]$, and has a *complementary* vertex (diagonally opposite in W) defined by exactly the other labels from each pair. Then the partner of v in Q has these four complementary labels, as well as label 2, so all the partners of P_1 are in Q_2. The complementary vertices in P_1 themselves have also their partners in Q_2. The resulting labels show that the partners of all the vertices of P_1 represent a 4-cube in Q_2. In particular, all of them have four neighbors (via their edges) in Q_2. Because Q_2 is a simple 4-polytope, it does not have any further vertices because they could not be connected to these 16 vertices. Hence, Q_2 is itself a 4-cube and has exactly 16 equilibrium vertices.

Every equilibrium (x, y) has to have label 2. If y has label 2 then $y \in Q_2$. The other equilibria (x, y) require that $x \in P_2$, where $P_2 \in \mathcal{P}^4_8$, but P_2 has at most 16 equilibrium vertices by (8). ☐

Lemmas 7, 8, and 9 cover the cases (iii), (ii), (i) of the pairs P_1, P_2 of disjoint facets of P that together have fewer than seven obstruction vertices. This completes the proof of Theorem 5.

Acknowledgments

The third author is indebted to Vissarion Fisikopoulous for our first attempt at this problem in 2015. That attempt provided an independent proof of Theorem 6 but was incomplete for proving Theorem 5, because we could not resolve the case of vertices that do not belong to either of the disjoint facets in a non-neighborly polytope in $\mathcal{P}^5_{10}$. We also thank Moritz Firsching for helpful communication on his classification of 5-polytopes.

This work was supported by the DFG Priority Program "Combinatorial Synergies" (grant no. 539847176).

References

Adsul, B., J. Garg, R. Mehta, M. Sohoni, and B. von Stengel (2021). Fast algorithms for rank-1 bimatrix games. *Oper. Res.* 69(2), 613–631.

Allamigeon, X., S. Gaubert, and F. Meunier (2023). Tropical complementarity problems and Nash equilibria. *SIAM J. Discrete Math.* 37(3), 1645–1665.

Bostanci, J. and J. Watrous (2021). Quantum game theory and the complexity of approximating quantum Nash equilibria. *Quantum* 6, 882.

Firsching, M. (2017). Realizability and inscribability for simplicial polytopes via nonlinear optimization. *Math. Program.* 166(1-2), 273–295. Classification data available at https://firsching.ch/polytopes.

Grünbaum, B. (2003). *Convex Polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer, New York, second edition.

Grünbaum, B. and V. P. Sreedharan (1967). An enumeration of simplicial 4-polytopes with 8 vertices. *J. Comb. Theory* 2, 437–465.

Ickstadt, C., T. Theobald, and E. Tsigaridas (2024). Semidefinite games. *Int. J. Game Theory* 53(3), 827–857.

Ickstadt, C., T. Theobald, E. Tsigaridas, and A. Varvitsiotis (2023). Semidefinite network games: Multiplayer minimax and complementarity problems. Preprint, arXiv:2310.20333.

Ickstadt, C., T. Theobald, and B. von Stengel (2025). Stable-set bounds for 5x5-games [Data set]. Zenodo. https://doi.org/10.5281/zenodo.17064079.

Joswig, M. and T. Theobald (2013). *Polyhedral and Algebraic Methods in Computational Geometry*. Universitext. Springer, London.

Kannan, R. and T. Theobald (2010). Games of fixed rank: A hierarchy of bimatrix games. *Econ. Theory* 42(1), 157–173.

Keiding, H. (1997). On the maximal number of Nash equilibria in an $n \times n$ bimatrix game. *Games Econ. Behav.* 21(1-2), 148–160.

Lemke, C. E. and J. T. Howson, Jr (1964). Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.* 12(2), 413–423.

McKelvey, R. D. and A. McLennan (1997). The maximal number of regular totally mixed Nash equilibria. *J. Econ. Theory* 72(2), 411–425.

McLennan, A. and I.-U. Park (1999). Generic 4×4 two person games have at most 15 Nash equilibria. *Games Econ. Behav.* 26(1), 111–130.

Miyata, H. and A. Padrol (2015). Enumerating neighborly polytopes and oriented matroids. *Experimental Mathematics* 24(4), 489–505.

Nash, J. (1951). Non-cooperative games. *Ann. Math.* 54(2), 286–295.

Quint, T. and M. Shubik (1997). A theorem on the number of Nash equilibria in a bimatrix game. *Int. J. Game Theory* 26(3), 353–359.

Savani, R. and B. von Stengel (2006). Hard-to-solve bimatrix games. *Econometrica* 74(2), 397–429.

Savani, R. and B. von Stengel (2016). Unit vector games. *Int. J. Econ. Theory* 12(1), 7–27.

Shapley, L. (1974). A note on the Lemke–Howson algorithm. *Mathematical Programming Study 1: Pivoting and Extensions* 175–189.

von Stengel, B. (1999). New maximal numbers of equilibria in bimatrix games. *Disc. Comp. Geom.* 21(4), 557–568.

von Stengel, B. (2002). Computing equilibria for two-person games. In: *Handbook of Game Theory with Economic Applications*, edited by R. J. Aumann and S. Hart, volume 3, 1723–1759. North-Holland, Amsterdam.

von Stengel, B. (2021). Finding Nash equilibria of two-player games. Preprint, arXiv:2102.04580.

von Stengel, B. (2022). *Game Theory Basics*. Cambridge University Press.

Vujić, M. (2022). *Geometry of the Sets of Nash Equilibria in Mixed Extensions of Finite Games*. Ph.D. thesis, Univ. Mannheim.

Ziegler, G. M. (1995). *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, New York.