Automated Equilibrium Analysis of $2 \times 2 \times 2$ Games

Sahar Jahani and Bernhard von Stengel

London School of Economics and Political Science, London WC2A 2AE, UK
{s.jahani,b.von-stengel}@lse.ac.uk

Abstract. The set of all Nash equilibria of a non-cooperative game with more than two players is defined by equations and inequalities between nonlinear polynomials, which makes it challenging to compute. This paper presents an algorithm that computes this set for the simplest game with more than two players with arbitrary (possibly non-generic) payoffs, which has not been done before. We give new elegant formulas for completely mixed equilibria, and compute visual representations of the best-response correspondences and their intersections, which define the Nash equilibrium set. These have been implemented in Python and will be part of a public web-based software for automated equilibrium analysis. For small games, which are often studied in economic models, a complete Nash equilibrium analysis is desirable and should be feasible. This project demonstrates the difficulties of this task and offers pathways for extensions to larger games.

Keywords: Equilibrium enumeration · Nash equilibrium · Three-player game.

1 Introduction

Game theory provides mathematical models for multiagent interactions. The primary solution concept is Nash equilibrium and its refinements (e.g., perfect equilibrium, [12]) or generalizations such as correlated equilibrium (which arises from regret-based learning algorithms). Already for two-player games, finding just one Nash equilibrium is PPAD-hard [3,5]. However, this "intractability" of the Nash equilibrium concept applies to *large* games. Many games that are used as economic models are small, with less than a few dozen payoff parameters, and often given in extensive form as game trees. It would be desirable to have a *complete* analysis of *all* Nash equilibria of such a game, in order to study the implications of the model. Such a complete analysis is known for two-player games. Their Nash equilibria can be represented as unions of "maximal Nash subsets" [14]. These are maximally "exchangeable" Nash equilibrium sets, that is, products of two polytopes of mixed strategies that are mutual best responses. Their non-disjoint unions form the topologically connected components of Nash equilibria, and are computed by the *lrsnash* algorithm of Avis, Rosenberg, Savani, and von Stengel [1], which works well for games with up to about twenty strategies per player.

Jahani, S., von Stengel, B. (2022). Automated Equilibrium Analysis of 2x2x2 Games. In: Kanellopoulos, P., Kyropoulou, M., Voudouris, A. (eds) Algorithmic Game Theory. SAGT 2022. Lecture Notes in Computer Science, vol 13584, pages 223-237. Springer, Cham.
https://doi.org/10.1007/978-3-031-15714-1_13

For games with more than two players, the set of all Nash equilibria cannot be described in such a way, because it is determined by equations and inequalities between nonlinear polynomials. The Gambit software package [10] provides access to polynomial solvers in order to compute Nash equilibria for generic games. "Generic" means that the payoffs do not represent edge cases, for example when (11) below reads as "$0 = 0$". The edge cases can be encoded as the zeros of a suitable polynomial in the game parameters and form a set of measure zero. Generic games have only finitely many equilibrium points. Non-generic games can have infinite set of equilibria.

However, rather remarkably, there is to our knowledge no algorithm that computes (in some description) the entire set of Nash equilibria for even the simplest game with more than two players if the game is non-generic, which naturally occurs for games in extensive form, such as "Selten's horse" [12]; see Section 4.3 below.

This paper describes an algorithm that computes the entire set of Nash equilibria for arbitrary $2 \times 2 \times 2$-games, that is, three-player games where every player has two strategies. These are the simplest games with more than two players that do not have a special structure (such as being a polymatrix game arising from pairwise interactions, see [8]). While this seems like a straightforward task, it is already challenging in its complexity.

One contribution of this paper is to reduce this complexity by carefully preserving the symmetry among the players, and a judicious use of intermediate parameters (equation (6) in Section 4) derived from the payoffs. We determine a quadratic equation (see (11)) that has a regular structure using determinants (not known to us before), which also implies that a generic $2 \times 2 \times 2$ game has at most two completely mixed equilibria (shown much more simply than in [4] or [9]). The standard approach to manipulating such complicated algebraic expressions is to use a computer algebra system [6].

As a "binary" game with only two pure strategies per player, the equilibria of a $2 \times 2 \times 2$ game can be visualized in a cube, but this needs some 3D graphics to be accessible (our graphics can be "moved in 3D"). We think that good visualizations of the geometry of equilibrium solutions of a game are important for understanding them, and their possible structure (both for applications of and research in game theory).

We present our algorithm in two parts: Identifying partially mixed equilibria (on the faces or edges of the cube) which arise from two-player equilibria where the third player plays a pure strategy that remains optimal; this part has a straightforward generalization to larger numbers of strategies for the three players, and may be practically very useful, certainly for a preliminary analysis. The second part is to look for completely mixed equilibria, which is challenging and does not generalize straightforwardly. A substantial part of the code, which we cannot describe in full because it involves a large number of case distinctions, deals systematically with the degenerate cases (which do arise in game trees even when payoffs are generic).

2 General form of the game

The following table describes the general form of a three-player game in which each player has two strategies:

Front: $1-r$

I \ II (III)	Left: $1-q$	Right: q
Up: $1-p$	0, 0, 0	0, a, 0
Down: p	A, 0, 0	B, c, 0

Back: r

I \ II (III)	Left: $1-q$	Right: q
Up: $1-p$	0, 0, α	0, b, γ
Down: p	C, 0, β	D, d, δ

(1)

This game is played by player I, II, III choosing (simultaneously) their second strategy with probability p, q, r, respectively. Player I chooses a row, either Up or Down (abbreviated U and D), player II chooses a column, either Left or Right (abbreviated L and R), and player III chooses a panel, either Front or Back (abbreviated F and B). The strategy names are also chosen to remember the six faces of the three-dimensional unit cube of mixed-strategy profiles (p, q, r), shown in Figure 1.

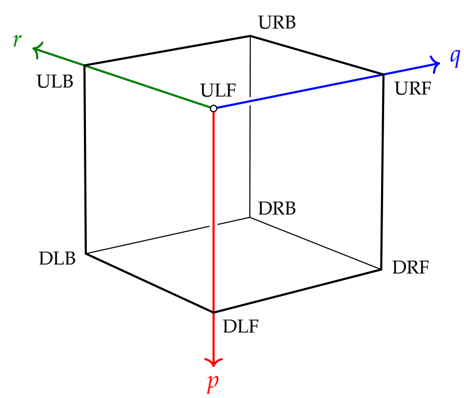

Fig. 1: Cube of mixed-strategy probabilities (p, q, r) drawn as in (1) down, right, and backwards.

Each of the eight cells in (1) has a payoff triple (T, t, τ) to the three players, with the payoffs to player I, II, III in upper case, lower case, and Greek letters, respectively. The payoffs in (1) are staggered and shown in color to distinguish them more easily between the players.

The payoffs have been normalized so that each player's first pure strategy has payoff zero throughout. (Zero is the natural "first" number, as in 0 and 1 for the two strategies of each player, or for the payoffs.) This normalization is obtained by subtracting a suitable constant from the player's payoffs for each combination of opponent strategies (e.g., each column for player I). This does not affect best responses [13, p. 239f]. With this normalization, the first strategy of each player gives always expected payoff zero.

For each player's second strategy, the expected payoffs are as follows:

$$\begin{array}{lrcl}
\text{player I}: & S(q,r) & = & (1-q)(1-r)\,A \;+\; q(1-r)\,B \;+\; (1-q)r\,C \;+\; qr\,D, \\
\text{player II}: & s(r,p) & = & (1-r)(1-p)\,a \;+\; r(1-p)\,b \;+\; (1-r)p\,c \;+\; rp\,d, \\
\text{player III}: & \sigma(p,q) & = & (1-p)(1-q)\,\alpha \;+\; p(1-q)\,\beta \;+\; (1-p)q\,\gamma \;+\; pq\,\delta,
\end{array} \tag{2}$$

so the three players can be treated symmetrically. The cyclic shift among p, q, r in (2), and corresponding choice of where to put b and c and β and γ in (1), will lead to more symmetric solutions.

The mixed-strategy profile (p, q, r) is a *mixed equilibrium* if each player's mixed strategy is a *best response* against the other players' strategies. That best response is a pure (deterministic) strategy, unless the two pure strategies have *equal* expected payoffs [11]. Hence, p is a best response of player I to (q, r) if the following conditions hold:

$$\begin{array}{lcl}
p = 0 & \Leftrightarrow & S(q,r) \le 0 \\
p \in [0,1] & \Leftrightarrow & S(q,r) = 0 \\
p = 1 & \Leftrightarrow & S(q,r) \ge 0\,.
\end{array} \tag{3}$$

Similarly, q is a best response of player II to (r, p) and r is a best response of player III to (p, q) if and only if

$$\begin{array}{lcl@{\qquad\qquad}lcl}
q = 0 & \Leftrightarrow & s(r,p) \le 0 & r = 0 & \Leftrightarrow & \sigma(p,q) \le 0 \\
q \in [0,1] & \Leftrightarrow & s(r,p) = 0 & r \in [0,1] & \Leftrightarrow & \sigma(p,q) = 0 \\
q = 1 & \Leftrightarrow & s(r,p) \ge 0 & r = 1 & \Leftrightarrow & \sigma(p,q) \ge 0\,.
\end{array} \tag{4}$$

For each player I, II, or III, the triples (p, q, r) that fulfill the respective conditions for p, q, or r in (3) or (4) define the *best-response correspondence* of that player, a subset of the cube $[0, 1]^3$. The set of Nash equilibria is the intersection of these three sets. The best-response correspondence for player I, for example, has one of the following forms, as shown in Figure 2:

(a) If $A = B = C = D = 0$, then $S(q, r) = 0$ for all $q, r \in [0, 1]$ and player I's best-response correspondence is the entire cube $[0, 1]^3$.

(b) If $A, B, C, D < 0$, then $S(q, r) < 0$ for all $(q, r) \in [0, 1]^2$ and strategy U strictly dominates D, so that player I will always play U, and the game reduces

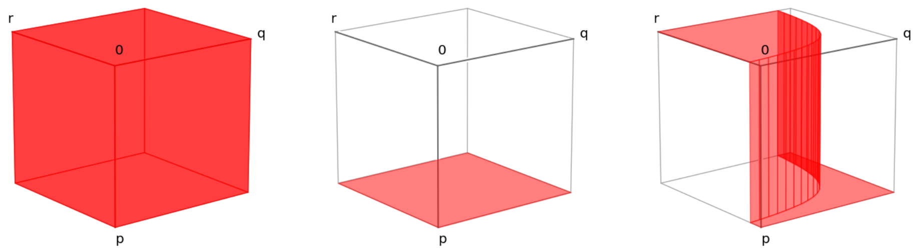

(a) (b) (c)

Fig. 2: Different forms of best-response correspondence.

to a two-player game between players II and III. The same happens when $A, B, C, D > 0$, in which case D strictly dominates U. In these two cases the best-response correspondence of player I is the upwards "U face" or downwards "D face" of the cube (as in Figure 2(b)), respectively.

(c) In all other cases, the best response of player I to (q, r) is sometimes U and sometimes D. The best-response correspondence of player I is then a surface that consists of subsets of the U or D face according to (3), which are connected by vertical parts, as in Figure 2(c) where player I is indifferent between U and D. Figure 3 shows a generic example.

The Nash equilibria of a game can be divided to two categories, based on which strategies are used:

- *partially mixed equilibria* in which at least one player plays a pure strategy (including pure equilibria where all the players play pure strategies). These equilibria are on the faces of the cube.
- *completely mixed equilibria* in which none of the players plays a pure strategy. These equilibria are in the interior of the cube.

In order to find all the equilibria in these games, we can divide the procedure into two parts:

(i) Find the *partially mixed* (including pure) equilibria.

(ii) Find the *completely mixed* equilibria.

We use different methods for each part. The union of the answers will be the set of all equilibria of the game.

3 Partially mixed equilibria

In an equilibrium, each player's strategy is a best response to the other players' strategies. In a partially mixed equilibrium, at least one player plays a pure strategy. All partially mixed equilibria are thus identified via six subgames. In each subgame we fix the strategy s_i of one player $i = 1, 2, 3$ to be 0 or 1. Fixing one player's strategy gives a 2×2 game for which we compute all equilibria. Then, for each equilibrium component of the subgame, given by the profile s_{-i} of strategies of the other two players, we check if it is the best response for the fixed player i and if it is, this means it is a partially mixed equilibrium (PMNE) of the game. Algorithm 1 gives a simplified pseudo-code.

Algorithm 1 Finding partially mixed equilibria

```
Input:   payoff matrix of a 2×2×2 game
Output: its set of partially mixed Nash equilibria
PMNE ← ∅
for each player i do
    for s_i ∈ {0,1} do
        SG ← 2×2 game when player i plays s_i
        cand ← all Nash equilibria of SG                  ▹ using lrsNash algorithm
        for each s_{-i} ∈ cand do                  ▹ strategy pair of the other two players
            if U_i(s_i, s_{-i}) ≥ U_i(1 − s_i, s_{-i}) then   ▹ U_i is the utility function for player i
                add (s_i, s_{-i}) to PMNE
return PMNE
```

Algorithm 1 also applies when cand is an infinite set of equilibria of the 2×2 subgame. Such a set can contain line segments or be equal to the whole facet of the cube. Then the output of Algorithm 1 is the intersection of the best-response surface of player i with cand. This intersection is computed as follows: Equilibrium segments of a 2×2 game can be only horizontal or vertical; then the strategy of one player is constant throughout the segment and just one variable changes. Furthermore, the intersection of an entire facet of the cube with the best-response surface has parameterized borders that are lines or hyperbola arcs (see Section 4.2).

4 Completely mixed equilibria

In this section, we assume that all the partially mixed equilibria are found using the previous algorithm. Here, we focus on finding the completely mixed equilibria. First, we find these equilibria algebraically by solving the best-response equations. Second, we display each player's best-response correspondence as a surface, and the intersection of these surfaces will also show the equilibria.

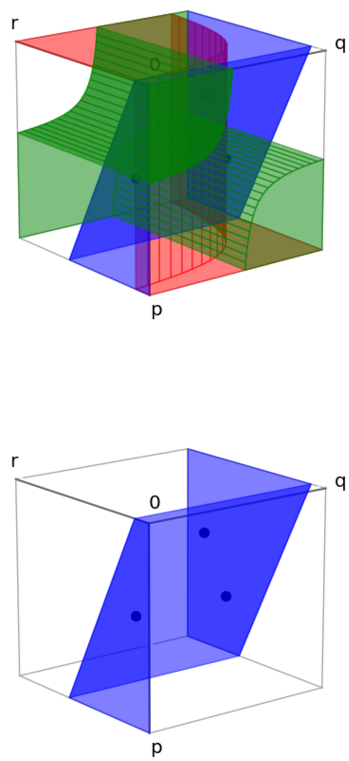

Fig. 3: Example of best-response surfaces of a game with two completely mixed equilibria and one partially mixed equilibrium, marked as black dots. (The actual display can be 3D-animated and handled interactively.)

4.1 Finding the completely mixed equilibria algebraically

We focus on player I using (3); the consideration for players II and III is analogous. We will show that the indifference equation $S(q, r) = 0$, which by (3) is necessary for player I to be able to mix ($0 < p < 1$), defines either a line or a (possibly degenerated) hyperbola, using possibly both branches.

Generically, the intersection of the three best-response surfaces is a finite set of points. However, certain kinds of degeneracy may occur, which leads to infinite components of Nash equilibria.

For our algebraic approach, we rewrite (2) as

$$\begin{aligned} S(q,r) &= A + Kq + Lr + Mqr \\ s(r,p) &= a + kr + lp + mrp \\ \sigma(p,q) &= \alpha + \kappa p + \lambda q + \mu pq \end{aligned} \tag{5}$$

with

$$\begin{array}{rclrclrcl} K &=& B-A, & L &=& C-A, & M &=& A-B-C+D, \\ k &=& b-a, & l &=& c-a, & m &=& a-b-c+d, \\ \kappa &=& \beta-\alpha, & \lambda &=& \gamma-\alpha, & \mu &=& \alpha-\beta-\gamma+\delta. \end{array} \tag{6}$$

The expressions in (5) are linear in each of p, q, r, and we consider when they are equal to zero, which defines the indifference surfaces:

$$\begin{array}{rcrcrcl} A &+& Kq &+& (L+Mq)r &=& 0 \\ a &+& lp &+& (k+mp)r &=& 0 \\ \alpha &+& \kappa p &+& (\lambda+\mu p)q &=& 0. \end{array} \tag{7}$$

We first eliminate r by multiplying the first equation in (7) with $(k+mp)$ and the second with $-(L+Mq)$ and adding them, which gives

$$(A+Kq)(k+mp)-(L+Mq)(a+lp)=0 \tag{8}$$

or, using determinants,

$$\begin{vmatrix} A & L \\ a & k \end{vmatrix} + \begin{vmatrix} A & L \\ l & m \end{vmatrix} p + \begin{vmatrix} K & M \\ a & k \end{vmatrix} q + \begin{vmatrix} K & M \\ l & m \end{vmatrix} pq = 0. \tag{9}$$

In the same way, we eliminate q by multiplying the last equation in (7) with $\begin{vmatrix} K & M \\ a & k \end{vmatrix} + \begin{vmatrix} K & M \\ l & m \end{vmatrix} p$ and (9) with $-(\lambda+\mu p)$ and addition, which gives

$$\left(\begin{vmatrix} K & M \\ a & k \end{vmatrix} + \begin{vmatrix} K & M \\ l & m \end{vmatrix} p\right)(\alpha+\kappa p) \;-\; \left(\begin{vmatrix} A & L \\ a & k \end{vmatrix} + \begin{vmatrix} A & L \\ l & m \end{vmatrix} p\right)(\lambda+\mu p) \;=\; 0 \tag{10}$$

or (verified by expanding each 3×3 determinant in the last column)

$$\begin{vmatrix} A & L & \alpha \\ K & M & \lambda \\ a & k & 0 \end{vmatrix} \;+\; \left(\begin{vmatrix} A & L & \alpha \\ K & M & \lambda \\ l & m & 0 \end{vmatrix} + \begin{vmatrix} A & L & \kappa \\ K & M & \mu \\ a & k & 0 \end{vmatrix}\right) p \;+\; \begin{vmatrix} A & L & \kappa \\ K & M & \mu \\ l & m & 0 \end{vmatrix} p^2 \;=\; 0. \tag{11}$$

Unless it states $0 = 0$, the quadratic equation (11) has at most two solutions for p, which have to belong to $[0,1]$ to represent a mixed equilibrium strategy of player I. Substituted into the linear equation (9) for q and the second equation in (7) for r, this then determines q and r uniquely unless one of the equations has no or infinitely many solutions. If q and r belong to $[0,1]$, these determine mixed equilibria. They are completely mixed if p, q, r are all strictly between 0 and 1. Moreover, a generic $2\times 2\times 2$ game has therefore at most two completely mixed equilibria (as proved in much more complicated ways by [4] and [9]).

The system (5) can be solved in exactly the same manner to derive a quadratic equation for q, where in (5), we only need to move the first equation into last position and change A, a, α to a, α, A respectively, and similarly for the other letters. Then (11) becomes

$$\begin{vmatrix} a & l & A \\ k & m & L \\ \alpha & \kappa & 0 \end{vmatrix} \;+\; \left(\begin{vmatrix} a & l & A \\ k & m & L \\ \lambda & \mu & 0 \end{vmatrix} + \begin{vmatrix} a & l & K \\ k & m & M \\ \alpha & \kappa & 0 \end{vmatrix}\right) q \;+\; \begin{vmatrix} a & l & K \\ k & m & M \\ \lambda & \mu & 0 \end{vmatrix} q^2 \;=\; 0. \tag{12}$$

Similarly, the quadratic equation for r states

$$\begin{vmatrix} \alpha & \lambda & a \\ \kappa & \mu & l \\ A & K & 0 \end{vmatrix} + \left(\begin{vmatrix} \alpha & \lambda & a \\ \kappa & \mu & l \\ L & M & 0 \end{vmatrix} + \begin{vmatrix} \alpha & \lambda & k \\ \kappa & \mu & m \\ A & K & 0 \end{vmatrix} \right) r + \begin{vmatrix} \alpha & \lambda & k \\ \kappa & \mu & m \\ L & M & 0 \end{vmatrix} r^2 = 0. \tag{13}$$

As before, in the generic case, any of the up to two solutions q to (12) determines r and p. Similarly, any of the up to two solutions r to (13) determines p and q.

If p, q, r are the solutions to (11), (12), (13) in a completely mixed equilibrium, they may be irrational numbers. They can be output as approximate floating-point numbers or symbolically with square roots as algebraic numbers (assuming rational payoffs as inputs).

The conditions (11), (12), (13) are all necessary when each player is required to be indifferent between his pure strategies. However, they may hold trivially in the form $0 = 0$, which may indicate infinite solution sets; an example is (11) for case (a) when $A = K = L = M = 0$. Furthermore, even if (11) has two real solutions p, say, then for one or both choices of p the third equation in (7) may state $0 = 0$ and then q is not determined; one would expect that this implies that (12) states $0 = 0$ as well. A further source of infinite solutions may be that some solutions for p, q, or r are 0 or 1, because then the respective player plays a pure strategy and does not have to be indifferent. This should come up in the analysis of the partially mixed equilibria in the previous section.

Other than these quadratic equations, we can acquire more information about the game by studying the relation between any two variables. Using (5), we can write each variable as a function of the other variable. So, from $S(q, r) = 0$ we will have:

$$q = \frac{-Lr - A}{Mr + K} = f_q(r), \qquad r = \frac{-Kq - A}{Mq + L} = f_r(q). \tag{14}$$

Similarly, we have four more equations derived from the other two equations. These equations will help us identify the mixed equilibria when the quadratic equations have infinite solutions and do not give us any information. An example for this case is the game shown in Figure 4. In this game, we have $0 = 0$ for (11) and (12). Also, (13) cannot be formed because we have to multiply by 0 to eliminate other variables; hence, we have to look at the relation between the variables. Here, r cannot be written as a function of other variables (division by 0), and for the other relations we have

$$p = f(q) = \tfrac{-8q+7}{-8q+8}, \qquad q = f(p) = \tfrac{-8p+7}{-8p+8},$$
$$p = f(r) = \tfrac{1}{2}, \qquad q = f(r) = \tfrac{3}{4}.$$

These relations show that there is a line of equilibria in the r direction where $p = \frac{1}{2}$ and $q = \frac{3}{4}$, as one can see on each player's best-response surface in Figure 4.

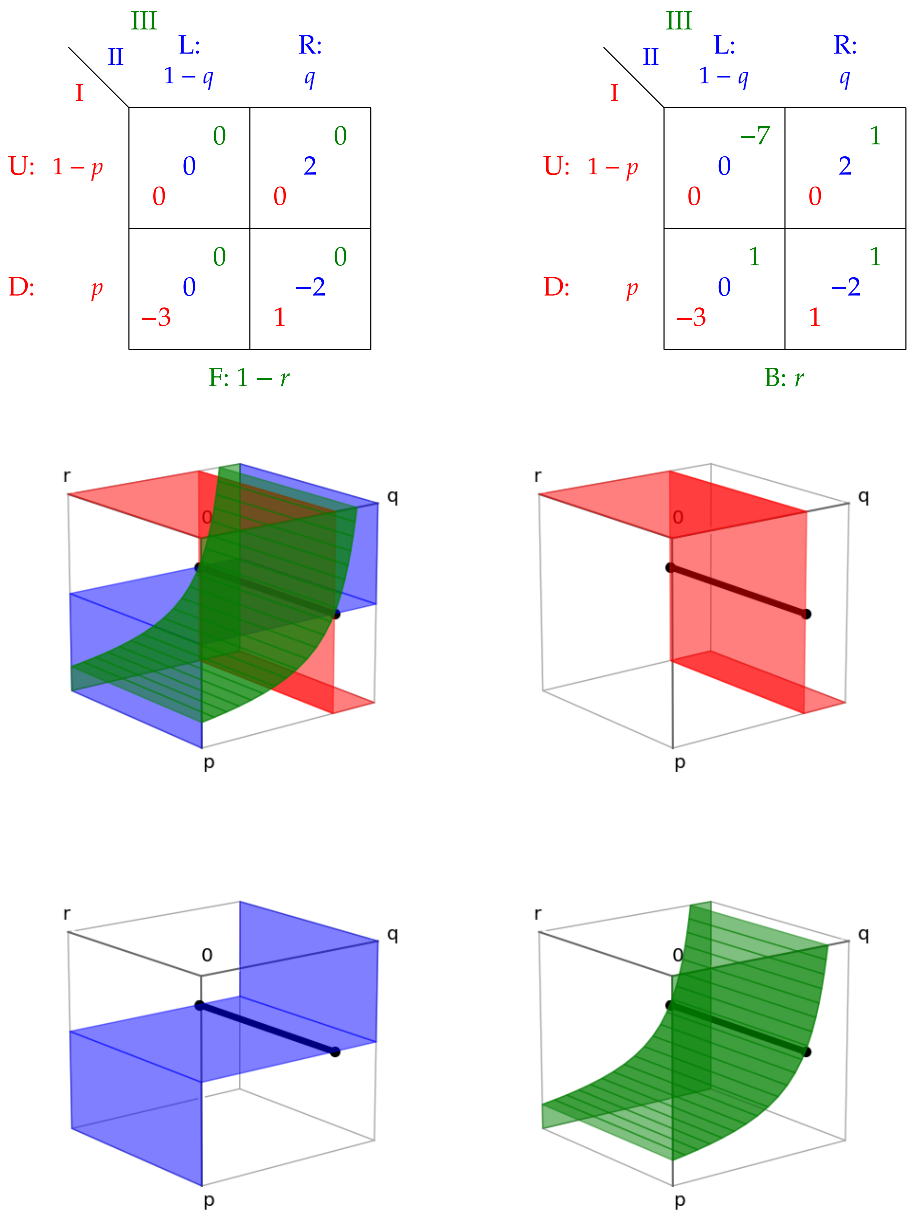

Fig. 4: Example of a game with a line of mixed equilibria when the quadratic equations have infinitely many solutions.

4.2 Displaying best-response surfaces

To see how the best-response surfaces look like, we focus on player I's expected payoff equation; for the other players it is similar. With (6) and (5), the condition $S(q,r) = 0$ states

$$S(q,r) = A + Kq + Lr + Mqr = 0\,. \tag{15}$$

(a) First, we exclude the case when $(A,B,C,D) = (0,0,0,0)$ because it means the player is completely indifferent between the two strategies in every point. Then every point in $[0,1] \times [0,1] \times [0,1]$ will be part of the best-response correspondence. In the next step we compute the intersection of the best-response correspondences of the other two players, so we do not need to take this first player into account.

We continue by studying different cases for $S(q,r) = 0$ when at least one of A,B,C,D is not 0.

(b) The *linear case* applies if $M = A - B - C + D = 0$, that is,

$$A + Kq + Lr = 0. \tag{16}$$

If $K = L = 0$ then $A = B = C = D$ and either $A = 0$ and (a) applies, or $A \neq 0$ and either U or D is dominant (and (16) has no solution), so assume $(K,L) \neq (0,0)$. If $K = 0$ then the line is defined by a constant for r, namely $r = -A/L$, and if $L = 0$ then the line is defined by a constant for q, namely $q = -A/K$. Otherwise, (16) expresses a standard linear relationship between q and r. In all cases, it is a line in the $q \times r$ plane which is extended vertically in the p-axis direction. According to (3), on this plane, Player I is indifferent between the first and second strategy. For the points on each side of the plane, (3) determines that the best response is $p = 0$ or $p = 1$.

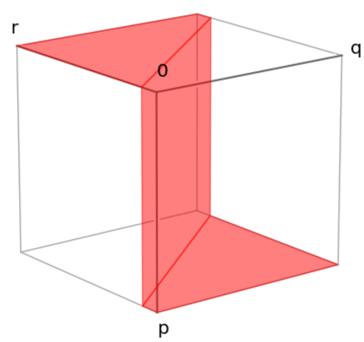

(c) Now, suppose $M \neq 0$. Then (15) is equivalent to

$$\frac{A}{M} + \frac{K}{M}q + \frac{L}{M}r + qr = 0\,. \tag{17}$$

Adding $\frac{KL}{M^2} - \frac{A}{M}$ on both sides of this equation and using (6) gives

$$\left(q + \frac{L}{M}\right)\left(r + \frac{K}{M}\right) = \frac{KL - AM}{M^2} = \frac{BC - AD}{M^2}. \tag{18}$$

If $BC - AD = 0$, then (18) states that $q = -\frac{L}{M}$ or $r = -\frac{K}{M}$. This defines two perpendicular lines, each similar to a line in case (b). This is a degenerate hyperbola, with a best-response surface like the blue surface in Figure 5 in the q direction, or as in this picture:

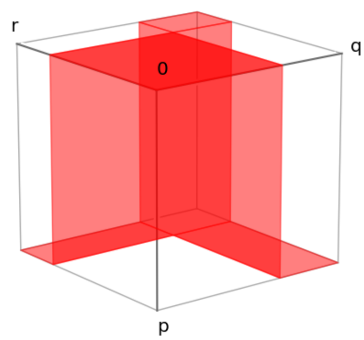

If $BC - AD \neq 0$, then these two lines are the asymptotes of a hyperbola defined by (18). Depending on the values of A, B, C, D, it is possible that the $[0,1] \times [0,1]$ rectangle contains two parts of the arcs of hyperbola or a part of one of the arcs (see the green and red best-response surface in Figure 3), or just a point on it, or none at all (but then the game has a dominated strategy). For the points (q, r) that are not located on the hyperbola, player I's pure best response is determined according to the inequalities $S(q, r) < 0$ and $S(q, r) > 0$ in (3). Note that when (15) with "<" or ">" instead of "=" is replaced by the corresponding inequality in (17), its direction is reversed if $M < 0$.

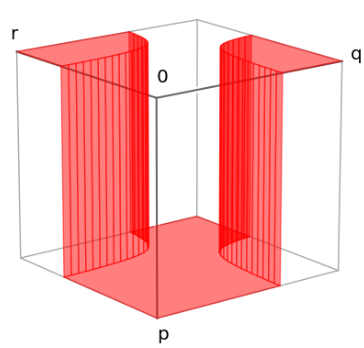

4.3 A well-known example

The extensive-form game in Figure 5 is a famous example from Selten [12, Fig. 1]. The game tree is in the shape of a horse, so this game is also known as "Selten's horse". The strategic form of this game is displayed on the right. It is known that this game has two segments of partially mixed equilibria and no completely mixed equilibria. Below that the best-response correspondences are displayed, with the Nash equilibria marked in black. There are two segments of equilibria, which include the pure equilibria (U, R, F) and (D, R, B).

5 Conclusions

The computational complexity of finding Nash equilibria is often concerned with asymptotic properties such as PPAD-hardness. Many concrete games are small and would profit from a complete analysis of all its Nash equilibria. This has been done for two-player games, but is significantly more difficult for general games with more than two players. Solvers based on solving polynomial equations and inequalities often fail in degenerate cases, which can have an infinite number of equilibria [10].

Our contribution is a "proof of principle" that the complete equilibrium set can be computed and displayed in full no matter how degenerate the game is. We apply it only to the simplest multiplayer game, namely three players with two strategies each, which had not been done before. We also streamlined the corresponding algebraic expressions using determinants in (11), (12), (13) and exploiting the symmetry of the setup. We did exploit the fact that the mixed-strategy profiles can be displayed in a three-dimensional cube. Our experience with the implementation is that one needs to deal with a large number of case distinctions for the possible degenerate cases.

Another insight is that computing partially mixed equilibria is a fruitful approach. For larger games, this means reducing the number of strategies for some players. This can already give information about equilibria with relatively little extra effort.

For larger games, it seems advisable to proceed incrementally in the same manner: Computing partially mixed equilibria, and using algebraic solvers such as done by Datta [6] under the assumption of nondegeneracy. The main question in this context is what kind of multiplayer games people really want to solve. Using models such as polymatrix games, which are based on pairwise interactions [8,7,2], may be the appropriate next step in this direction.

Acknowledgements We thank the anonymous referees for helpful comments.

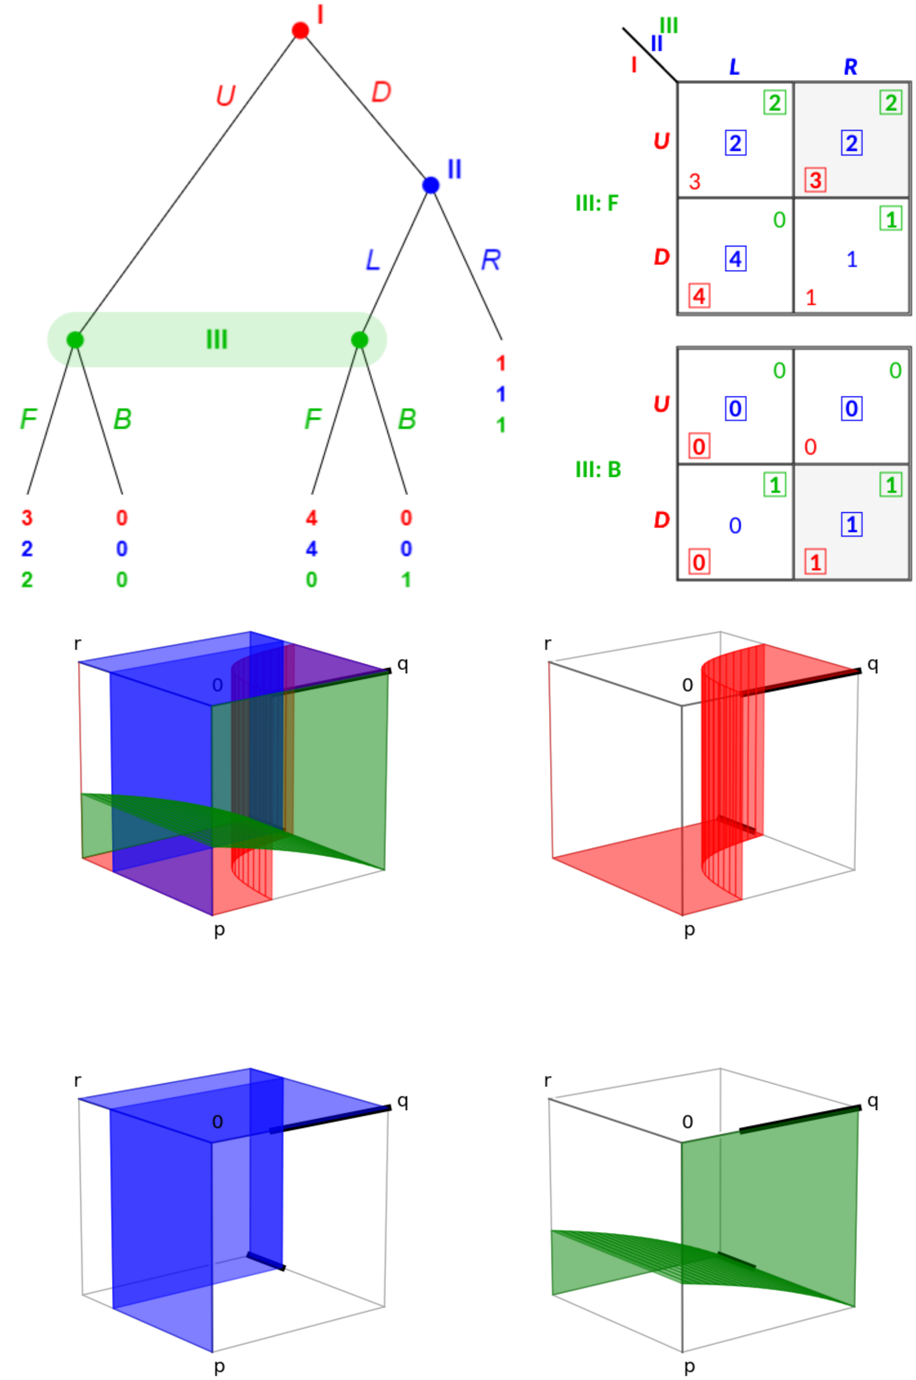

Fig. 5: A famous example: “Selten’s horse”, see Section 4.3.

References

1. Avis, D., Rosenberg, G.D., Savani, R., von Stengel, B.: Enumeration of Nash equilibria for two-player games. Economic Theory **42**(1), 9–37 (2010)
2. Cai, Y., Candogan, O., Daskalakis, C., Papadimitriou, C.: Zero-sum polymatrix games: A generalization of minmax. Mathematics of Operations Research **41**(2), 648–655 (2016)
3. Chen, X., Deng, X., Teng, S.H.: Settling the complexity of computing two-player Nash equilibria. Journal of the ACM **56**(3), Article 14 (2009)
4. Chin, H., Parthasarathy, T., Raghavan, T.: Structure of equilibria in N-person non-cooperative games. International Journal of Game Theory **3**(1), 1–19 (1974)
5. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SIAM Journal on Computing **39**(1), 195–259 (2009)
6. Datta, R.S.: Finding all Nash equilibria of a finite game using polynomial algebra. Economic Theory **42**(1), 55–96 (2010)
7. Govindan, S., Wilson, R.: Computing Nash equilibria by iterated polymatrix approximation. Journal of Economic Dynamics and Control **28**(7), 1229–1241 (2004)
8. Howson, Jr, J.T.: Equilibria of polymatrix games. Management Science **18**(5, Part I), 312–318 (1972)
9. McKelvey, R.D., McLennan, A.: The maximal number of regular totally mixed Nash equilibria. Journal of Economic Theory **72**(2), 411–425 (1997)
10. McKelvey, R.D., McLennan, A.M., Turocy, T.L.: Gambit: Software tools for game theory, version 16.0.1 (2016), `http://www.gambit-project.org`
11. Nash, J.: Non-cooperative games. The Annals of Mathematics **54**(2), 286–295 (1951)
12. Selten, R.: Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory **4**(1), 25–55 (1975)
13. von Stengel, B.: Game Theory Basics. Cambridge University Press, Cambridge, UK (2022)
14. Winkels, H.M.: An algorithm to determine all equilibrium points of a bimatrix game. In: Moeschlin, O., Pallaschke, D. (eds.) Game Theory and Related Topics, pp. 137–148. North-Holland, Amsterdam (1979)