A Canonical Game – 75 Years in the Making – Showing the Equivalence of Matrix Games and Linear Programming

Benjamin Brooks and Philip J. Reny*

April 28, 2023†

Abstract

According to Dantzig (1949), von Neumann was the first to observe that for any finite two-person zero-sum game, there is a feasible linear programming (LP) problem whose saddle points yield equilibria of the game, thus providing an immediate proof of the minimax theorem from the strong duality theorem. We provide an analogous construction going in the other direction. For any LP problem, we define a game and, with a brief and elementary proof, show that *every* equilibrium either yields a saddle point of the LP problem or certifies that one of the primal or dual programs is infeasible and the other is infeasible or unbounded. We thus obtain an immediate proof of the strong duality theorem from the minimax theorem. Taken together, von Neumann's and our results provide a succinct and elementary demonstration that matrix games and linear programming are "equivalent" in a classical sense.

Keywords: Matrix games, linear programming, equivalence.

JEL Codes: D00, C61, C72.

---

*Department of Economics, University of Chicago, babrooks@uchicago.edu, preny@uchicago.edu. We thank Songzi Du, Elliot Lipnowski, Roger Myerson, Bernhard von Stengel, and two anonymous referees for helpful comments. We are especially grateful to Ilan Adler and Rakesh Vohra for extensive discussions of the literature and our main result. Reny gratefully acknowledges financial support from the National Science Foundation (SES-1724747).

†First version June 2021.

1 Introduction

Since at least the early 1950's, starting with the work of Dantzig (1951), it was widely understood that the theory of two-person zero-sum games (matrix games) and the theory of linear programming (LP) had been shown to be "equivalent," both in the sense that a solution to any matrix game can be obtained by solving a suitably chosen LP problem and vice versa, as well as in the sense that their fundamental theorems, minimax and strong duality, each follow from the other. At least, this was the perceived state of affairs prior to 2013.

In an important paper, Adler (2013) points out that there are significant flaws in the classical arguments for equivalence which "...led to incomplete proofs of the relationship between the Minimax Theorem of game theory and the Strong Duality Theorem of linear programming" (p. 165). Adler then makes sufficient repairs to these flaws so as to correctly establish a certain equivalence between matrix games and linear programming.

As Adler notes, the flaws in the classical arguments occur only in one direction, namely when attempting to reduce an LP problem (i.e., a primal and dual pair of linear programs) to a game problem and when attempting to prove the strong duality theorem from the minimax theorem. The other direction, reducing games to LP problems and proving the minimax theorem from strong duality, was correctly settled very early. Indeed, according to Dantzig (1949, 1951, 1982), von Neumann observed in 1947 that any matrix game can be solved by solving a suitably chosen feasible LP problem. Since the chosen LP problem is feasible, this construction yields an immediate proof of the minimax theorem from the strong duality theorem.[1]

To correct the flaws in the classical arguments, Adler (2013) does two things.[2] First, for any LP problem, he constructs a matrix game with the property that any equilibrium of that game yields a saddle point of the LP problem whenever a saddle point exists, thereby correctly reducing an LP problem to a game problem. Second, and entirely separately, he provides a proof of the strong duality theorem from the minimax theorem by showing that the minimax theorem implies Ville's theorem implies Tucker's theorem implies Farkas' lemma implies strong duality.

[1]It is evident that von Neumann explained to Dantzig the reduction of a game to a linear program, Farkas lemma, and von Neumann's (1947) LP duality theorem, during a private meeting in 1947 that is recounted in Dantzig (1982, p. 45). So whenever we refer here to "von Neumann's construction/result," we are referring both to the reduction of a game to an LP problem whose primal program is as given in Dantzig (1951, pp. 330-331), and to the immediate implication that the minimax theorem follows from the strong duality theorem.

[2]Adler (2013) also obtains computability results which we will not discuss here.

A key point for our purposes here is that, while von Neumann's LP problem construction for solving games yields an immediate proof of the minimax theorem from the strong duality theorem, Adler's game construction for solving LP problems does not yield an immediate proof of the strong duality theorem from the minimax theorem. Indeed, in order to prove that his game solves LP problems, Adler must appeal to the strong duality theorem itself. Thus, even after Adler's (2013) corrections, the natural counterpart to von Neumann's result remained missing.

The purpose of the present paper is to provide the counterpart to von Neumann's result by constructing for any LP problem a matrix game with the property that every equilibrium of the game either provides a saddle point of the LP problem or provides what we call an *unbounded direction* of the LP problem. An unbounded direction is a pair of vectors, a primal direction and a dual direction, such that starting from a feasible solution to one of the programs, it is feasible to move in the corresponding direction and improve the objective without bound.[3] To see that a game with these properties achieves the desired goal, notice first that if the LP problem has a saddle point, then both the primal and dual programs are feasible and bounded, and so there can exist no unbounded direction. So every equilibrium of the game must yield a saddle point, and the LP problem has been reduced to a game problem. Additionally however—and this is the crucial distinction between our game construction and Adler's (2013)—a game with these properties provides an immediate proof of strong duality from minimax. Indeed, by the minimax theorem, the constructed game has an equilibrium. By the properties of the game, either, the equilibrium yields a saddle point, in which case a saddle point exists, or, the equilibrium yields an unbounded direction, in which case if either one of the constituent programs is feasible then it is unbounded. From this last fact, it follows (by the elementary weak duality inequality) that the other program is infeasible. Thus, from the minimax theorem, it immediately follows that either a saddle point exists or one of the constituent programs is infeasible and the other is infeasible or unbounded, which is precisely the conclusion of the strong duality theorem.

So analogous to von Neumann's construction, the game constructed here not only reduces any LP problem to a game problem, it also furnishes an immediate proof of the strong duality theorem from the minimax theorem. Furthermore, the proof that our game has these

[3]Unbounded directions are output by the simplex algorithm, in the event that no saddle point exists. This is how Dantzig (1963) proves the strong duality theorem. Another well-known proof of strong duality establishes the existence of either a saddle point or an unbounded direction by an application of Farkas' Lemma (see, e.g., Gale, 1960, Chapter 3). Indeed, an unbounded direction is a solution to the Farkas dual of the inequality system characterizing saddle points. Thus, in reducing LP problems to matrix games, it is natural to seek a game whose equilibria yield either saddle points or unbounded directions. Note that the proof that equilibria of our game have this property does not invoke Farkas' lemma or any other separation result.

properties is brief and elementary, using only the defining property of linearly independent vectors. Taken together, von Neumann's classical result and the present result provide a natural formalization of the very old idea that matrix games and linear programming are "equivalent."

2 Preliminaries

Let F be any ordered field with total order $>$, e.g., the reals or the rationals with their usual order.[4] Henceforth, all entries of any matrix or vector are in F. All vectors are column vectors (so any transposed vector, $x^\top$, is a row vector), and (x, y) denotes the concatenation of vectors x and y. The ith row (column) of A is denoted A_i (A^i). Given vectors x and y of the same length, we write $x \geq y$ to mean $x_i \geq y_i$ for every coordinate i, and $x \geq 0$ means that every coordinate of x is non-negative. We use "0" to denote either a matrix or a vector of zeroes, "1" to denote a vector of 1's, and I to denote the identity matrix. In each case, their sizes are those that are uniquely appropriate given the context.

A *linear programming (LP) problem* is a dual pair of constituent linear programs. Formally, for any $m \times n$ matrix A and vectors b and c of lengths m and n respectively, the ($m \times n$ dimensional) LP problem (A, b, c) is the pair of optimization problems:

$$\max_{x \geq 0} c^\top x \text{ s.t. } Ax \leq b; \tag{1a}$$

$$\min_{y \geq 0} b^\top y \text{ s.t. } A^\top y \geq c. \tag{1b}$$

We refer to (1a) as the *primal* program and to (1b) as the *dual* program. A *feasible solution* to (1a) is any non-negative $x \in F^n$ that satisfies $Ax \leq b$, and an *optimal solution* is a feasible solution that attains the maximum. These terms are defined analogously for (1b). It is well known that if x and y are feasible solutions for (1a) and (1b), respectively, then $cx \leq yb$ (a result known as *weak duality*), since

$$c^\top x \leq y^\top Ax \leq y^\top b. \tag{2}$$

[4]Because the minimax theorem, the strong duality theorem, and Farkas' lemma are valid in any ordered field (see, e.g., Gale, 1960), we have constructed our proofs so that they too are valid in any ordered field. In particular, we avoid the use of topological results such as the fact that bounded sequences have convergent subsequences, which are valid for the field of real numbers but not for, e.g., the field of rational numbers.

We say that (1) is *feasible* if both (1a) and (1b) have feasible solutions. A *saddle point* of (1) is a pair of feasible solutions that have the same value, i.e., a pair of non-negative vectors (x, y) such that $Ax \leq b$, $A^\top y \geq c$, and $c^\top x = b^\top y$. By the weak duality inequality (2), any saddle point consists of optimal solutions. The *strong duality theorem* states that for every LP problem, either there exists a saddle point, or one of the constituent programs is infeasible and the other is either infeasible or unbounded.

All definitions up to now have been standard. We next introduce a definition that will play an important role in our main result. Say that a pair of non-negative vectors (x, y) is an *unbounded direction* of the LP problem (A, b, c) if $Ax \leq 0$, $A^\top y \geq 0$, and $c^\top x > b^\top y$.

The sense in which an unbounded direction $(\bar{x}, \bar{y})$ of an LP problem is "unbounded" is that starting from a feasible solution to one of the constituent programs, moving in the corresponding direction, $\bar{x}$ for the primal, $\bar{y}$ for the dual, will improve the objective without bound. Indeed, suppose that the primal program is feasible. Then there is $x \geq 0$ such that $Ax \leq b$. Hence, $0 \leq (\bar{y}^\top A)x = \bar{y}^\top (Ax) \leq \bar{y}^\top b < \bar{x}^\top c$ and $A(x + \lambda\bar{x}) \leq b$ for every $\lambda \geq 0$. Consequently, the primal program is unbounded (above) because its value, $c^\top(x + \lambda\bar{x})$, at the feasible solution $x + \lambda\bar{x}$ increases without bound as $\lambda \to \infty$.

Of course, if the primal (dual) program is unbounded, then the dual (primal) program is infeasible by the weak duality inequality (2). Therefore, and importantly, the existence of an unbounded direction establishes that one of the constituent programs is infeasible and the other is infeasible or unbounded.

Consider any $m \times n$-dimensional LP problem (A, b, c). Define the $(m+n+1) \times (m+n)$ matrix $\hat{A}$ and the $(m+n+1)$-vector $\hat{b}$ by,

$$\hat{A} := \begin{bmatrix} 0 & -A^\top \\ A & 0 \\ -c^\top & b^\top \end{bmatrix} \quad \hat{b} := \begin{pmatrix} -c \\ b \\ 0 \end{pmatrix}, \tag{3}$$

and let $C := \begin{bmatrix} \hat{A} & I & -1 \end{bmatrix}$.[5] Say that $\alpha \in F$ is a *solution bound* for (A, b, c) if $\sum_j w_j < \alpha$ whenever $w \geq 0$, $Cw = \hat{b}$, and the columns j of C with $\hat{w}_j > 0$ are linearly independent, i.e., whenever w is a so-called *basic feasible* solution of $Cw = \hat{b}$.

Since for any set of linearly independent columns of C there is at most one solution $w \geq 0$ to $Cw = \hat{b}$ that places its positive weights on that particular set of columns, and because C has only finitely many sets of linearly independent columns, every LP problem has a solution

[5]Here I is the $(m+n+1)$-dimensional identity matrix, and the -1 is a column vector all of whose $(m+n+1)$ entries are equal to -1.

bound.[6] As we will show, solution bounds permit the search for saddle points to be confined to a bounded set such that either a saddle point exists within this set or the LP problem admits an unbounded direction.

Finally, a *matrix game* is simply a matrix P. For any positive integer k, let S^k be the set of non-negative elements of F^k whose entries sum to one. An *equilibrium* of the $(m \times n)$ matrix game P is a pair $(\hat{s}, \hat{t}) \in S^m \times S^n$ such that

$$\hat{s}^\top P \hat{t} = \max_{s \in S^m} s^\top P \hat{t} = \min_{t \in S^n} \hat{s}^\top P t.$$

If $v = \max_{s \in S^m} \min_{t \in S^n} s^\top P t = \min_{t \in S^n} \max_{s \in S^m} s^\top P t$ then v is called the *value* of the game P. If P has a value v, then $(\hat{s}, \hat{t})$ is an equilibrium if and only if $v = \hat{s}^\top P \hat{t} = \max_{i \in \{1,\ldots,m\}} P_i \hat{t} = \min_{j \in \{1,\ldots,n\}} \hat{s}^\top P^j$. The *minimax theorem* states that every matrix game has a value and an equilibrium. See von Neumann (1928) for the original result and Gale (1960) for the extension to any ordered field.

3 The Canonical Game

We can now state our main result.

Theorem 1. *Fix an $m \times n$ dimensional LP problem (A, b, c). Choose any positive $\alpha \in F$ that is a solution bound for (A, b, c). Then the matrix game P with $n + m + 1$ rows and columns defined by*

$$P := \alpha \begin{bmatrix} 0 & -A^\top & 0 \\ A & 0 & 0 \\ -c^\top & b^\top & 0 \end{bmatrix} + \begin{bmatrix} c & \cdots & c \\ -b & \cdots & -b \\ 0 & \cdots & 0 \end{bmatrix}. \tag{4}$$

has the following property: If $(x, y, z), (x^*, y^*, z^*) \in F^n \times F^m \times F$ are equilibrium strategies for the row and column players, respectively,[7] then either $(\alpha x^*, \alpha y^*)$ is a saddle point of (A, b, c), or (x, y) is an unbounded direction of (A, b, c).

[6] In fact, using the rank theorem (i.e., row-rank equals column-rank), which is not needed for the proof of our main result, it is easy to show that $\alpha = 1 + (n + m + 1)^2 \max\{||b||, ||c||\} \max_W ||W^{-1}||$, is a solution bound for (A, b, c), where $||\,\centerdot\,||$ is the maximum absolute value of the entries in the given vector or matrix, and where the second maximum is over all square invertible submatrices W of C.

[7] So, for example, x gives the weights placed by the row player on the first n rows of P, y on the next m rows, etc.

Proof. Suppose that $(s,t) = ((x,y,z),(x^*,y^*,z^*))$ is any equilibrium of P. Then the value is $v := s^\top Pt$. Letting $q^* := (\alpha x^*, \alpha y^*)$ and given $\hat{A}$ and $\hat{b}$ defined in (3), $Pt = \hat{A}q^* - \hat{b}$ and $v = s^\top(\hat{A}q^* - \hat{b})$. There are two cases, $v \le 0$ and $v > 0$.

Suppose first that $v \le 0$. Then $0 \ge v = \max_i P_i t = \max_i(\hat{A}_i q^* - \hat{b}_i)$, and so $\hat{A}q^* \le \hat{b}$, which implies that $A(\alpha x^*) \le b$, $A^\top(\alpha y^*) \ge c$, and $b^\top(\alpha y^*) \le c^\top(\alpha x^*)$. Weak duality (2) then implies that $b^\top(\alpha y^*) = c^\top(\alpha x^*)$. Hence, $(\alpha x^*, \alpha y^*)$ is a saddle point.

Next, suppose that $v > 0$. The proof will be complete if we show that (x,y) is an unbounded direction. To accomplish this, it suffices to show that $v = -s^\top\hat{b}$. Then, $-s^\top\hat{b} = v \le s^\top P^j = \alpha s^\top \hat{A}^j - s^\top\hat{b}$ for all $j \le n+m$, which implies that $s^\top\hat{A} \ge 0$ (since $\alpha > 0$). Hence, $y^\top A \ge zc^\top$ and $Ax \le zb$, and so $z(c^\top x - y^\top b) \le y^\top Ax - y^\top Ax = 0$. Since $c^\top x - y^\top b = -s^\top\hat{b} = v > 0$, it follows that $z = 0$ and $c^\top x > y^\top b$. Hence, $y^\top A \ge 0$, $Ax \le 0$ and (x,y) is an unbounded direction as desired. The remainder of the proof shows that $v = -s^\top\hat{b}$ by constructing for the column player a best reply to s that puts positive weight on the last column of P (and so $v = s^\top P^{m+n+1} = -s^\top\hat{b}$).

Since $v = \max_i(\hat{A}_i q^* - \hat{b}_i)$, we have $\hat{A}q^* \le \hat{b} + 1v$, and so there is a non-negative $u \in F^{m+n+1}$ such that $\hat{A}q^* + u = \hat{b} + 1v$. (So u_i is row's gain from optimally deviating from row i.) Choose a $(\hat{q},\hat{u},\hat{v}) \ge 0$ with the fewest number of positive coordinates such that $\hat{A}\hat{q} + \hat{u} = \hat{b} + 1\hat{v}$, $\hat{v} \le v$, and $\hat{u}_i = 0$ whenever $u_i = 0$. Observe that $s^\top(\hat{A}\hat{q} - \hat{b}) = \hat{v}$ because $s_i > 0$ implies that $u_i = 0$ (which in turn implies $\hat{u}_i = 0$). Let $\hat{w} := (\hat{q},\hat{v},\hat{u})$ and $C := \begin{bmatrix} \hat{A} & I & -1 \end{bmatrix}$. We claim that the columns j of C with $\hat{w}_j > 0$ are linearly independent. Otherwise,[8] there is $d \ne 0$ such that $Cd = 0$ and $d_j \ne 0$ implies $\hat{w}_j > 0$. We may assume that that, either, the last coordinate of d is positive, or, the last coordinate is zero and some other coordinate is positive (if not then replace d with $-d$). Then, setting $\lambda := \min_{j:d_j>0}(\hat{w}_j/d_j)$, the non-negative vector $(\tilde{q},\tilde{u},\tilde{v}) := (\hat{q},\hat{u},\hat{v}) - \lambda d$ has strictly fewer positive coordinates than $(\hat{q},\hat{u},\hat{v})$ which, since $\hat{A}\tilde{q} + \tilde{u} = \hat{b} + 1\tilde{v}$, $\tilde{v} \le \hat{v} \le v$, and $u_i = 0 \Rightarrow \hat{u}_i = 0 \Rightarrow \tilde{u}_i = 0$ for every i, contradicts the choice of $(\hat{q},\hat{u},\hat{v})$ and proves the claim.

Since $\hat{w} \ge 0$ satisfies $C\hat{w} = \hat{b}$ and the columns j of C with $\hat{w}_j > 0$ are linearly independent, we have $\sum_j \hat{w}_j < \alpha$ because α is a solution bound for (A,b,c). So since $\hat{w} := (\hat{q},\hat{v},\hat{u}) \ge 0$, we have $\sum_j \hat{q}_j < \alpha$. Hence, $\hat{t} := \left(\hat{q}/\alpha, 1 - \sum_j \hat{q}_j/\alpha\right)$ is in S^{m+n+1} and $\hat{t}_{m+n+1} = 1 - \sum_j \hat{q}_j/\alpha > 0$. Moreover, because $s^\top(\hat{A}\hat{q} - \hat{b}) = \hat{v} \le v = s^\top(\hat{A}q^* - \hat{b})$, we have

$$v \le s^\top P\hat{t} = s^\top(\hat{A}\hat{q} - \hat{b}) \le s^\top(\hat{A}q^* - \hat{b}) = s^\top Pt = v.$$

Therefore, $s^\top P\hat{t} = v$ and so $\hat{t}$, which puts positive weight on the last column of P, is a best reply to s as desired. □

[8] Note that by definition the empty set of vectors is linearly independent.

Remark 1. Theorem 1 is the result we desire because it delivers two key requirements.
I. *Every equilibrium of P provides a saddle point of* (1) *whenever one exists.* Indeed, given any equilibrium, Theorem 1 states that either (a) or (b) holds. If a saddle point exists, then both constituent programs are feasible and bounded and hence there can be no unbounded direction. Consequently, (b) cannot hold and so (a) must hold. Thus, the column player's strategy provides a saddle point whenever one exists.
II. *The properties of P yield an immediate proof of strong duality from the minimax theorem.* Indeed, for any LP problem (A, b, c), the minimax theorem ensures that P has an equilibrium. By Theorem 1, either (a) holds and the column player's strategy demonstrates that a saddle point exists, or (b) holds and the row player's strategy provides an unbounded direction, demonstrating that one of the constituent programs is infeasible and the other is infeasible or unbounded.

Remark 2. Note that Theorem 1 does not assume that P has an equilibrium. Theorem 1 merely states that *if P has an equilibrium*, then either (a) or (b) must hold for that equilibrium. In particular, Theorem 1 (and hence the equivalence described in the preceding remark) does not itself rely on the minimax theorem.

Remark 3. The game P is canonical in that it has a natural interpretation in terms of a pair of adversaries, one of whom is trying to exhibit a saddle point of (A, b, c), and the other of whom is trying to prove that no saddle point exists by exhibiting an unbounded direction. By the weak duality inequality (2), a saddle point is simply a non-negative pair $(\tilde{x}, \tilde{y})$ that satisfies $A\tilde{x} \leq b$, $A^\top \tilde{y} \geq c$, and $c^\top \tilde{x} \geq b^\top \tilde{y}$. Defining $\hat{A}$ and $\hat{b}$ as in the proof of Theorem 1, we can write this system of inequalities as $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b}$. As the proof of Theorem 1 shows, the value of P is non-positive if and only if a saddle point exists. Indeed, to force the maximizer's (row player's) payoff to be non-positive, any equilibrium strategy (x^*, y^*, z^*) for the minimizer (column player) in our game P must be such that $(\tilde{x}, \tilde{y}) := (\alpha x^*, \alpha y^*)$ is feasible for the system $(\hat{A}, \hat{b})$, meaning that $(\alpha x^*, \alpha y^*)$ is a saddle point. Even so, the maximizer can still obtain a payoff of zero by playing the last row and so the value is never negative. On the other hand, the maximizer wants to establish that the minimizer's proposed solution is infeasible by exhibiting a violated constraint. The proof of Theorem 1 shows that, in equilibrium, either the minimizer succeeds in identifying a saddle point, or the maximizer finds a weighted average of the constraints that shows that any candidate solution is infeasible. In the latter case, the maximizer's strategy (x, y, z) solves the Farkas alternative to $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b}$, i.e., $\hat{A}^\top(x, y, z) \geq 0$ and $\hat{b}^\top(x, y, z) < 0$. Since, as shown in our proof, these inequalities imply that $z = 0$, they reduce to our definition of (x, y) as an unbounded direction for the original LP problem.

Remark 4. As is well known, saddle points of an LP problem may be unbounded. In contrast, strategies in a matrix game (including in our game P) are necessarily bounded. So a key step in reducing an LP problem to a matrix game problem is to show that the search for a saddle point, i.e., a non-negative solution $(\tilde{x}, \tilde{y})$ to the system $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b}$ can be confined to a bounded set, which, after normalizing, may be identified with the strategies in a game. As is well-known, if this system has a feasible solution, then it has a so-called *basic* feasible solution, in which the submatrix defined by all of the columns of $\hat{A}$ that are given positive weight and by all of the rows of $\hat{A}$ corresponding to constraints that are satisfied with equality, has linearly independent columns.[9] So it would be natural when seeking to confine one's saddle point search to a bounded set, to choose a bound $\tilde{\alpha}$ that is just large enough to include each of the finitely many basic feasible solutions. Indeed, this is the approach taken by Adler (2013). But in fact such a bound is not large enough for our purpose. Indeed, when the value of the resulting game $\tilde{P}$ (with $\alpha = \tilde{\alpha}$) is positive, the bound is large enough to conclude that no saddle point exists, but it is not large enough for us to conclude that the row player's strategy yields an unbounded direction, and so we would be unable to conclude that the LP problem is infeasible/unbounded. Thus, the game $\tilde{P}$ would not prove strong duality from minimax. Consequently, the bound that we choose is larger and includes all $(\tilde{x}, \tilde{y})$ such that $(\tilde{x}, \tilde{y}, \tilde{v})$ is a basic feasible solution to $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b} + 1\tilde{v}$, which has a solution even when the original LP problem has no saddle point. Our solution bound is large enough to not only include all basic feasible solutions to $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b}$ when any exist, thereby admitting a strategy for the minimizer that makes the value of P zero, it is also large enough to ensure that when $\hat{A}(\tilde{x}, \tilde{y}) \leq \hat{b}$ has no non-negative solution, the maximizer can guarantee the value of P, which now is positive, only by using a strategy that yields an unbounded direction.

Remark 5. The proof of Theorem 1 provides an immediate proof of Farkas' lemma from the minimax theorem. Indeed, fix any $k \times l$ matrix $\hat{A}$ and any vector $\hat{b}$ of length k, where neither need be as in (3). The proof in fact shows that if $s \in F^k$ and $t = (q^*, z^*) \in F^l \times F$ are any equilibrium strategies for the row and column players, respectively, of the game[10]

$$\hat{P} := \left[\begin{array}{cc} \alpha\hat{A} & 0 \end{array} \right] + \left[\begin{array}{ccc} -\hat{b} & \cdots & -\hat{b} \end{array} \right],$$

[9]The type of argument that is used to establish the existence of a basic feasible solution is standard, being also the core part of the argument used to prove Caratheodory's theorem. We use this type of argument in our proof to establish the independence of the columns of $C := \left[\begin{array}{ccc} \hat{A} & I & -1 \end{array} \right]$ that are given positive weight by $\hat{w}$. We include this (elementary) argument explicitly rather than appealing to the properties of basic feasible solutions so as to make our proof entirely self-contained.

[10]Here, $0 < \alpha \in F$ must satisfy $\sum_j w_j < \alpha$ whenever $w \geq 0$ solves $Cw = \hat{b}$ and the columns j of $C := \left[\begin{array}{ccc} \hat{A} & I & -1 \end{array} \right]$ with $w_j > 0$ are linearly independent.

then exactly one of the following holds: (a) $\hat{A}q^* \leq \hat{b}$, or (b) $\hat{A}^\top s \geq 0$ and $\hat{b}^\top s < 0$. Consequently, Farkas' lemma (in its inequality form) follows immediately from the minimax theorem.

Remark 6. For a second immediate proof of Farkas' lemma from minimax, simply apply Theorem 1 to the LP problem (A, b, c) where $c = 0$. Then, by the minimax theorem, Theorem 1 implies that either, (A, b, c) has a saddle point (x, y) in which case, in particular, $Ax \leq b$ holds, or (A, b, c) has an unbounded direction (x, y) in which case $A^\top y \geq 0$ and $b^\top y < 0$ holds, because with $c = 0$ we cannot have $Ax \leq 0$ and $cx > 0$.

Remark 7. Just as von Neumann's construction is not unique, P is not the only matrix game that achieves our goals here. After the circulation of previous drafts of this paper, which included the game P and Theorem 1, Rakesh Vohra (private communication) and Bernard von Stengel (2022) each independently obtained distinct game constructions that yield the desired counterpart to von Neumann's result. Both constructions require finding a suitably large bound analogous to our α. The game P, however, may permit the most elementary, short, and direct self-contained proof.

4 Addendum: A Symmetric Perspective

An apparent asymmetry between matrix games and linear programs is that matrix games always have solutions, i.e., equilibria, while the constituent programs of an LP problem might not have any optimal solutions at all. We show here that this asymmetry can be eliminated with an appropriate modification of the meaning of a "solution" of an LP problem.

Say that a pair of non-negative vectors (x, y) is a *solution* of an LP problem (A, b, c) if (x, y) is either a saddle point or an unbounded direction.

With this definition, the relationship between matrix games and linear programs becomes completely symmetric. Indeed, we can now state the following.

(1) Every LP problem has a solution.

(2) Every matrix game has a solution.

(3) For every matrix game there is an LP problem whose solutions provide solutions to the game.

(4) For every LP problem there is a matrix game whose solutions provide solutions to the LP problem.

Statement (1) is a consequence of the minimax theorem combined with Theorem 1, and is a slight strengthening of the strong duality theorem;[11] (2) is the minimax theorem; (3) is von Neumann's result; and (4) is our main result, Theorem 1.

References

Adler, I. (2013): "The equivalence of linear programs and zero-sum games," *International Journal of Game Theory*, 42, 165–177.

Dantzig, G. B. (1949): "Programming of interdependent activities: II mathematical model," *Econometrica*, 200–211.

——— (1951): "A proof of the equivalence of the programming problem and the game problem," in *Activity analysis of production and allocation*, ed. by T. C. Koopmans, Cowles Commission for Research in Economics, 13, 330–338.

——— (1963): *Linear Programming and Extensions*, RAND Corporation.

——— (1982): "Reminiscences about the origins of linear programming," *Operations Research Letters*, 1, 43–48.

Gale, D. (1960): *The theory of linear economic models*, University of Chicago press.

von Neumann, J. (1928): "Zur theorie der gesellschaftsspiele," *Mathematische annalen*, 100, 295–320.

——— (1947): "Discussion of a maximum problem," in *Collected Works of John von Neumann*, ed. by A. H. Taub, Macmillan, New York, vol. VI, 89–95, 1963.

von Stengel, B. (2022): "Zero-Sum Games and Linear Programming Duality," Working paper, Department of Mathematics, London School of Economics.

[11] But as noted previously in footnote 3, the existence of a solution to an LP problem in this sense is demonstrated as part of the standard proof of the strong duality theorem using Farkas' lemma.