Euler Circle **Summer 2025**

Lemke-Howson Algorithm

Name: Mia Amélie Mendeleev *Submission Date: July 13, 2025*

Introduction

The computation of Nash equilibria in finite two-player games represents one of the fundamental challenges in algorithmic game theory. Prior to 1964, while Nash's existence theorem guaranteed the presence of mixed strategy equilibria, no constructive method existed for systematically finding these equilibria. This changed with the groundbreaking work of Carlton E. Lemke and J. T. Howson Jr., who developed an elegant algorithm that not only proves existence but provides a computational pathway to equilibrium.

The Lemke–Howson algorithm, first presented in their 1964 paper, provided a constructive method for computing equilibria. Though not framed in these terms at the time, it was later recognized as solving a special case of a linear complementarity problem (LCP), opening the door to modern algorithmic analysis. This transformation enabled the application of pivoting techniques, similar to those used in linear programming, to systematically traverse the space of possible strategies until reaching a Nash equilibrium.

What makes the Lemke-Howson algorithm particularly remarkable is its constructive nature. Rather than merely asserting the existence of equilibria through topological arguments, the algorithm provides a step-by-step procedure that, when executed, produces an actual equilibrium point. This constructive approach bridges the gap between theoretical game theory and practical computation, making it possible to analyze real-world strategic interactions.

In this paper, we provide a comprehensive mathematical analysis of the Lemke-Howson algorithm, using two paradigmatic games to illustrate its operation: Rock Paper Scissors and the Prisoner's Dilemma. These examples serve as concrete illustrations of the algorithm's behavior across different game structures - from symmetric zero-sum games to asymmetric games with dominant strategies. Through these examples, we demonstrate the algorithm's versatility and provide intuitive understanding of its mathematical foundations.

The significance of this work extends beyond its immediate computational applications. The algorithm establishes important connections between game theory and linear complementarity theory, opens pathways for complexity analysis of equilibrium computation, and provides the theoretical foundation for numerous extensions and variants that have emerged in the decades since its introduction.

Definitions

Game

A game is a mathematical model of strategic interaction between rational decision-makers called players. A game in normal form consists of a set of players, a set of strategies for each player, and payoff functions that determine each player's outcome based on all players' chosen strategies.

Example

Two firms, the players, are deciding whether to enter a new market. Each firm chooses ”Enter” or ”Stay Out”, which are their possible strategies. Their profits, also known as their payoffs, depend on both decisions: if both enter, each loses $2 million from competition; if only one enters, that firm gains $5 million while the other gets $0; if neither enters, both get $0. This game is based off of how two real companies compete against each other in a new market.

Pure Strategy

A pure strategy is a specific action chosen with certainty from a player's available options. It represents a deterministic choice rather than a probabilistic one.

Example

In the market entry game, ”Enter” is a pure strategy for Firm 1. If Firm 1 adopts this pure strategy, they will definitely enter the market with probability 1.

Strategy Profile

A strategy profile specifies one strategy for each player in the game. It describes the combination of choices made by all players simultaneously.

Example

In the two-firm game, (Enter, Stay Out) is a strategy profile where Firm 1 enters and Firm 2 stays out. The four possible pure strategy profiles are: (Enter, Enter), (Enter, Stay Out), (Stay Out, Enter), and (Stay Out, Stay Out).

Payoffs

Payoffs represent the utility or outcome each player receives from a particular strategy profile. They quantify what each player gains or loses based on everyone's choices.

Example

In the market entry game, when both firms enter (strategy profile (Enter, Enter)), each firm receives a payoff of -\$2 million. When only Firm 1 enters (Enter, Stay Out), Firm 1's payoff is \$5 million and Firm 2's payoff is \$0.

Bimatrix Game

A bimatrix game is a two-player game represented using two separate payoff matrices, one for each player. The first matrix shows Player 1's payoffs, and the second matrix shows Player 2's payoffs for each combination of strategies.

Example

The market entry game as a bimatrix:

	Enter	Stay Out
Enter	$(-2, -2)$	$(5, 0)$
Stay Out	$(0, 5)$	$(0, 0)$

Best Response

A best response is a strategy that maximizes a player's payoff given the strategies chosen by all other players. It represents the optimal choice for a player when facing a specific situation.

Example

If Firm 2 chooses "Enter," then Firm 1's payoffs are -\$2 million for "Enter" and \$0 for "Stay Out." Since $\$0 > -\$2$ million, Firm 1's best response to Firm 2 entering is "Stay Out." Conversely, if Firm 2 chooses "Stay Out," Firm 1's best response is "Enter" (payoff \$5 million vs \$0).

Mixed Strategy

A mixed strategy is a probability distribution over a player's pure strategies. Instead of choosing one action with certainty, the player randomizes among available actions according to specified probabilities.

Example

Firm 1 might use a mixed strategy of entering with probability 0.7 and staying out with probability 0.3. If both firms use mixed strategies and their choices are independent, expected payoffs are calculated by weighting each outcome by its probability of occurrence.

Nash equilibrium

A Nash equilibrium is a strategy profile where each player's strategy is a best response to the other players' strategies. No player can unilaterally change their strategy and improve their payoff.

Example

In the market entry game, (Stay Out, Enter) is a Nash equilibrium. Given that Firm 2 enters, Firm 1's best response is to stay out (payoff \$0 vs -\$2 million). Given that Firm 1 stays out, Firm 2's best response is to enter (payoff \$5 million vs \$0). Since neither firm wants to deviate unilaterally, this is an equilibrium.

Cooperative vs Non-cooperative Games

In cooperative games, players can form binding agreements and coalitions, focusing on how to divide total payoffs among coalition members. In non-cooperative games, players make independent decisions without binding agreements, each seeking to maximize their own payoff.

Example

Non-cooperative: The market entry game where each firm independently decides whether to enter, leading to potential losses if both enter simultaneously.
Cooperative: The same firms negotiate a binding agreement where one enters and compensates the other, avoiding competitive losses and maximizing joint profits through coordination.

Background

The Development of Game Theory

Game theory, as a formal mathematical discipline, emerged in the early 20th century and was shaped by the work of several pioneering thinkers. Its formal birth is typically attributed to the publication of Theory of Games and Economic Behavior in 1944 by John von Neumann and Oskar Morgenstern. This foundational text focused primarily on two-player zero-sum games, where one player's gain is exactly the other's loss. The early work focused on establishing that players could use optimal strategies to secure a predictable outcome.

However, real-world strategic interactions are rarely purely zero-sum. In economic markets, political negotiations, or biological systems, agents often have partially aligned or entirely independent interests. This insight led to a major shift in the field, culminating in the work of John Nash in the early 1950s. Nash extended game theory to non-zero-sum games and introduced the central notion of a Nash equilibrium: a set of strategies (one per player) in which no player has an incentive to unilaterally change their decision.

Nash also proved that every game with a finite number of players and strategies has at least one equilibrium in mixed strategies. This was a groundbreaking result that ensured equilibrium-based analysis could be applied widely, even when players use randomized strategies. Following Nash's work, game theory expanded rapidly, influencing a range of disciplines including economics, political science, evolutionary biology, and computer science.

Strategic Reasoning and the Role of Equilibria

The Nash equilibrium has become a standard way to define rational outcomes in strategic settings. But it naturally leads to deeper questions: if an equilibrium always exists, how many equilibria can a game have? Can we find them all? What do multiple equilibria mean for real-world decision-making?

Some games have a single unique equilibrium, while others have several possible ones. Even relatively small games can have multiple solutions, sometimes involving mixtures of strategies rather than pure choices. This raises practical issues: in economics, for example, when several equilibria exist, it can be unclear which one will actually occur. This uncertainty can affect market behavior, negotiations, or coordination.

This complexity has led researchers to search for patterns or rules about how many equilibria are possible in typical games. In many cases, games that do not have unusual symmetries or special structures tend to follow certain regular trends. For example, one surprising result, known as Wilson's Oddness Theorem, shows that under ordinary conditions, the number of completely mixed strategy equilibria (where all strategies are played with positive probability) is always odd. While this doesn't tell us the exact number of equilibria in a given game, it reveals a hidden structure in how they are distributed. The result comes from analyzing how these solutions behave under small changes to the game's payoffs and how their characteristics persist.

Toward Algorithms and Computation

Although we know equilibria exist in theory, actually finding them is often a difficult task. Even for games with just two players, working out an equilibrium can be time-consuming. This challenge has led to a growing focus on the practical side of game theory, specifically on how we can compute or approximate solutions.

One major development in this area is the Lemke–Howson algorithm, introduced in 1964. This algorithm applies to two-player games and provides a step-by-step method for finding one equilibrium. It works by systematically exploring possible solutions using labeled graphs and logical rules that guide the search. Though it only produces one equilibrium (not necessarily all), it was one of the first methods to show how to find an equilibrium rather than just prove it exists.

The Lemke-Howson approach helped highlight just how hard this problem can be. For larger or more complicated games, finding all equilibria can be infeasible. As a result, researchers often study smaller games, develop simplified models, or use approximations when an exact solution is too complex to compute. In applied contexts, the goal is often not to find every equilibrium but to understand the likely behavior of players and how outcomes change with different assumptions.

Proof through Example

0.1 Rock Paper Scissors Setup

Consider Rock Paper Scissors as our primary example. Let $S_1 = S_2 = \{R, P, S\}$ represent the strategy sets. The payoff matrices are:

$$A_{RPS} = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}, \quad B_{RPS} = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \tag{1}$$

Note that $B_{RPS} = -A_{RPS}$, making this a zero-sum game.

0.2 Prisoner's Dilemma Setup

For the Prisoner's Dilemma, let $S_1 = S_2 = \{C, D\}$ where C = Cooperate and D = Defect:

$$A_{PD} = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix}, \quad B_{PD} = \begin{pmatrix} 3 & 5 \\ 0 & 1 \end{pmatrix} \tag{2}$$

1 Mixed Strategy Spaces and Nash Equilibrium

Definition 1 (Mixed Strategy) *A mixed strategy for Player 1 is $x = (x_1, x_2, \ldots, x_m)$ where $x_i \geq 0$ and $\sum_{i=1}^{m} x_i = 1$.*

Rock Paper Scissors Example: A mixed strategy is $x = (x_R, x_P, x_S)$ where $x_R + x_P + x_S = 1$ and all components are non-negative.

Prisoner's Dilemma Example: A mixed strategy is $x = (x_C, x_D)$ where $x_C + x_D = 1$.

1.1 Nash Equilibrium Characterization

Theorem 2 (Nash Equilibrium via Complementary Slackness) *A strategy profile (x^*, y^*) is a Nash equilibrium if and only if:*

$$\textit{For Player 1: } x_i^* > 0 \Rightarrow (Ay^*)_i = \max_{k=1,\ldots,m} (Ay^*)_k \tag{3}$$

$$\textit{For Player 2: } y_j^* > 0 \Rightarrow (B^T x^*)_j = \max_{k=1,\ldots,n} (B^T x^*)_k \tag{4}$$

Rock Paper Scissors Analysis: In the symmetric equilibrium $x^* = y^* = (1/3, 1/3, 1/3)$, we compute:

$$A_{RPS} y^* = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1/3 \\ 1/3 \\ 1/3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \tag{5}$$

Since all components of Ay^* are equal (all zero), every strategy yields the same expected payoff, satisfying the complementary slackness conditions.

Prisoner's Dilemma Analysis: For the pure strategy equilibrium $(x^*, y^*) = ((0,1), (0,1))$ (both defect):

$$A_{PD} y^* = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{6}$$

Since $x_D^* = 1 > 0$ and $(A_{PD} y^*)_D = 1 > 0 = (A_{PD} y^*)_C$, the condition is satisfied.

2 Linear Complementarity Problem Formulation

2.1 Transformation to LCP

The Nash equilibrium conditions can be written as a Linear Complementarity Problem (LCP). Let v_1 and v_2 be the equilibrium payoffs. **Rock Paper Scissors LCP:** The system becomes:

$$\begin{pmatrix} 0 & -1 & 1 \, 1 & 0 & -1 & -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} y_R \; y_P \; y_S \end{pmatrix} - v_1 \begin{pmatrix} 1 \; 1 \; 1 \end{pmatrix} \leq \begin{pmatrix} 0 \; 0 \; 0 \end{pmatrix} \; x_R, x_P, x_S \quad \geq 0 \; x_i \cdot ((Ay)_i - v_1) = 0 \text{ for } i \in R, P, \tag{7}$$

Prisoner's Dilemma LCP:

$$\begin{pmatrix} 3 & 0 \, 5 & 1 \end{pmatrix} \begin{pmatrix} y_C \; y_D \end{pmatrix} - v_1 \begin{pmatrix} 1 \; 1 \end{pmatrix} \leq \begin{pmatrix} 0 \; 0 \end{pmatrix} \; x_C, x_D \qquad \geq 0 \; x_i \cdot ((Ay)_i - v_1) = 0 \text{ for } i \in C, D \tag{8}$$

3 The Lemke-Howson Algorithm

3.1 Labeling System

Each variable receives a unique label from $\{1, 2, \ldots, m+n\}$:

Rock Paper Scissors Labeling: - $x_R \leftrightarrow 1$, $x_P \leftrightarrow 2$, $x_S \leftrightarrow 3$ - $y_R \leftrightarrow 4$, $y_P \leftrightarrow 5$, $y_S \leftrightarrow 6$

Prisoner's Dilemma Labeling: - $x_C \leftrightarrow 1$, $x_D \leftrightarrow 2$ - $y_C \leftrightarrow 3$, $y_D \leftrightarrow 4$

3.2 Algorithm Execution on Rock Paper Scissors

[Rock Paper Scissors Execution] Let's trace through the algorithm with missing label 1 (corresponding to x_R).

Initial Tableau:

$$\begin{array}{c|ccc|ccc|c} & x_R & x_P & x_S & y_R & y_P & y_S & \text{RHS} \\ \hline s_R & 1 & 0 & 0 & 0 & -1 & 1 & 0 \\ s_P & 0 & 1 & 0 & 1 & 0 & -1 & 0 \\ s_S & 0 & 0 & 1 & -1 & 1 & 0 & 0 \\ t_R & 0 & 1 & -1 & 1 & 0 & 0 & 0 \\ t_P & -1 & 0 & 1 & 0 & 1 & 0 & 0 \\ t_S & 1 & -1 & 0 & 0 & 0 & 1 & 0 \end{array} \tag{9}$$

Iteration 1: Enter x_R (label 1), leave t_P (missing label becomes 5).

Iteration 2: Enter y_P (label 5), leave s_R (missing label becomes 4).

Continuing... The algorithm terminates at the symmetric equilibrium:

$$x^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \quad y^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) \tag{10}$$

3.3 Algorithm Execution on Prisoner's Dilemma

example[Prisoner's Dilemma Execution] Starting with missing label 1 (corresponding to x_C):

Initial State: All slack variables are basic, representing the "no mixed strategy" starting point.

Iteration 1: The algorithm immediately identifies that cooperation is dominated. The pivot operation moves toward the pure strategy solution.

Termination: The algorithm quickly converges to:

$$x^* = (0, 1), \quad y^* = (0, 1) \tag{11}$$

representing mutual defection. example

4 Convergence and Correctness Proofs

4.1 Finite Convergence

Theorem 3 (Finite Convergence) *The Lemke-Howson algorithm terminates in a finite number of steps.*

Proof: We prove this through our examples:

Rock Paper Scissors: The algorithm maintains exactly one missing label at each iteration. Since there are only $2 \times 3 = 6$ possible labels, and we cannot revisit the same basic solution (due to non-degeneracy), the algorithm must terminate within 6 iterations.

Prisoner's Dilemma: The dominant strategy structure ensures even faster convergence. The algorithm recognizes that cooperation is dominated and quickly pivots to the defection equilibrium.

The general proof follows because:

1. Each pivot operation either introduces a new basic variable or maintains progress toward satisfying all complementarity conditions.

2. There are finitely many possible basic solutions with exactly one missing label.

3. The algorithm cannot cycle due to the specific structure of the complementarity constraints.

■

4.2 Correctness

Theorem 4 (Correctness) *When the Lemke-Howson algorithm terminates, it has found a Nash equilibrium.*

Proof: Termination occurs when no label is missing, meaning all complementarity conditions are satisfied.

Rock Paper Scissors Verification: At $(x^*, y^*) = ((1/3, 1/3, 1/3), (1/3, 1/3, 1/3))$:

$$A_{RPS}y^* = (0, 0, 0)^T \tag{12}$$
$$B_{RPS}^T x^* = (0, 0, 0)^T \tag{13}$$

Since all strategies yield equal expected payoffs, no player can improve by deviating.

Prisoner's Dilemma Verification: At $(x^*, y^*) = ((0, 1), (0, 1))$:

$$A_{PD}y^* = (0, 1)^T \tag{14}$$
$$B_{PD}^T x^* = (0, 1)^T \tag{15}$$

Since $x_D^* = 1 > 0$ and defection yields higher payoff than cooperation, this satisfies the Nash conditions. ■

5 Existence Proof

Theorem 5 (Existence of Nash Equilibrium) *Every finite two-player game has at least one Nash equilibrium in mixed strategies.*

Proof: The Lemke-Howson algorithm provides a constructive proof:

Step 1: Any finite game can be formulated as an LCP (demonstrated with both examples).

Step 2: The algorithm's finite convergence (Theorem 2) guarantees it will terminate.

Step 3: The correctness theorem ensures the output is a Nash equilibrium.

Rock Paper Scissors demonstrates: Even in complex symmetric games, the algorithm finds the unique symmetric equilibrium.

Prisoner's Dilemma demonstrates: The algorithm works equally well for games with dominant strategies, finding the intuitive solution. ■

6 Complexity Analysis Through Examples

6.1 Upper Bounds

Theorem 6 (Complexity Upper Bound) *The Lemke-Howson algorithm requires at most $\binom{m+n}{2}$ pivot operations.*

Rock Paper Scissors: At most $\binom{6}{2} = 15$ operations. **Prisoner's Dilemma:** At most $\binom{4}{2} = 6$ operations.

In practice, both examples converge much faster than the theoretical upper bound.

7 Advanced Properties

7.1 Index Theory

Theorem 7 (Index Properties) *Nash equilibria can be classified by their index ± 1. The Lemke-Howson algorithm finds equilibria of index $+1$.*

Rock Paper Scissors: The symmetric equilibrium has index $+1$, explaining why the algorithm finds it reliably.

Prisoner's Dilemma: The mutual defection equilibrium also has index $+1$.

7.2 Computational Insights

The algorithm's behavior on our examples reveals important insights:

1. **Symmetric Games** (Rock Paper Scissors): The algorithm exploits symmetry to find the unique symmetric equilibrium efficiently.
2. **Dominant Strategy Games** (Prisoner's Dilemma): The algorithm quickly identifies and converges to dominant strategy solutions.
3. **General Structure**: The pivoting mechanism naturally handles both zero-sum and non-zero-sum games within the same framework.

Conclusion

The Lemke-Howson algorithm stands as one of the most elegant and important contributions to computational game theory. Through our detailed analysis using Rock Paper Scissors and the Prisoner's Dilemma, we have demonstrated the algorithm's fundamental properties and wide applicability.

7.3 Key Contributions

Our analysis has established several important results:

Constructive Existence Proof: The algorithm provides a constructive, elementary proof of Nash's existence theorem without relying on topological fixed-point theorems. This constructive nature makes the existence result practically meaningful and computationally relevant.

Computational Efficiency: While the algorithm has exponential worst-case complexity, our examples demonstrate that it performs efficiently on many structured games, particularly those with symmetric properties or dominant strategies.

Mathematical Elegance: The transformation of the equilibrium-finding problem into a linear complementarity problem reveals deep mathematical connections between game theory and optimization theory. The pivoting mechanism provides a natural way to traverse the strategy space systematically.

Broad Applicability: The algorithm handles diverse game structures uniformly, from zero-sum games like Rock Paper Scissors to non-zero-sum games like the Prisoner's Dilemma, demonstrating its fundamental nature.

7.4 Theoretical Implications

The Lemke-Howson algorithm has profound implications for our understanding of strategic interaction:

1. It establishes that equilibrium computation is a well-defined computational problem, not merely a mathematical abstraction.
2. The algorithm's index-theoretic properties provide insights into the structure of equilibrium sets and their stability properties.
3. The connection to linear complementarity problems opens pathways for applying optimization techniques to game-theoretic problems.

7.5 Limitations and Future Directions

While the Lemke-Howson algorithm is foundational, it has certain limitations that have motivated subsequent research:

Computational Complexity: The algorithm's exponential worst-case complexity has led to the development of approximation algorithms and the study of the complexity class PPAD.

Equilibrium Selection: The algorithm finds one equilibrium but provides no guidance on which equilibrium to expect in games with multiple equilibria.

Generalization: Extensions to games with more than two players require significantly different approaches, as the complementarity structure becomes more complex.

7.6 Historical Impact

The Lemke-Howson algorithm has had lasting impact on multiple fields:

- It established computational game theory as a distinct research area.
- It provided the theoretical foundation for modern equilibrium computation software.
- It influenced the development of mechanism design and auction theory by making equilibrium analysis computationally tractable.

7.7 Final Remarks

The mathematical framework presented in this paper, illustrated through Rock Paper Scissors and the Prisoner's Dilemma, demonstrates that the Lemke-Howson algorithm is not merely a computational tool but a fundamental contribution to our understanding of strategic interaction. Its constructive proof of Nash's existence theorem, combined with its computational tractability, makes it an indispensable part of the game theorist's toolkit.

The algorithm's elegance lies in its simplicity: by following a systematic pivoting procedure, it transforms the abstract concept of Nash equilibrium into a concrete computational object. This transformation has enabled countless applications in economics, computer science, and beyond, cementing the Lemke-Howson algorithm's place as one of the most important algorithms in the history of game theory.

As computational power continues to grow and new applications emerge, the Lemke-Howson algorithm remains relevant and continues to inspire new research directions. Its fundamental insights into the structure of strategic interaction ensure its continued importance in the evolving landscape of computational game theory.

Acknowledgments

The author would like to thank Emma Cardwell for mentorship and advice, and Simon Rubinstein-Salzedo for his lectures and support throughout the class.

References

- Martin J. Osborne and Ariel Rubinstein. A Course in Game Theory https://sites.math.rutgers.edu/~zeilberg/EM20/OsborneRubinsteinMasterpiece.pdf
- C. E. Lemke; J. T. Howson, Jr. Journal of the Society for Industrial and Applied Mathematics, Vol. 12, No. 2 (Jun., 1964), 413-423. https://www1.cmc.edu/pages/faculty/MONeill/math188/papers/lemkehowson.pdf
- Cottle, Richard W.; Pang, Jong-Shi; Stone, Richard E. The linear complementarity problem. http://www.maths.lse.ac.uk/Personal/stengel/CottlePangStone2008.pdf

- Shapley, Lloyd S. A note on the Lemke-Howson algorithm. https://www.rand.org/content/dam/rand/pubs/reports/2006/R1538.pdf
- Nash, John Non-cooperative games. https://www.cs.upc.edu/~ia/nash51.pdf
- Nash, John F. Equilibrium points in n -person games. https://rybn.org/halloffame/PDFS/1950_J.Nash-Equilibrium%20Points%20in%20n-Person%20Games_3p.pdf