Scilight

Bulletin of Computational Intelligence
https://www.sciltp.com/journals/bci



Article

Check for updates

Probabilistic Bounds on the Separation of Exact and Approximate Nash Equilibria in Solid Games

Markos Z. Tsoukalas

Computational Mechanics Laboratory, School of Pedagogical and Technological Education, 141 21 Athens, Greece; tsoukalas6967@gmail.com

How To Cite: Tsoukalas, M.Z. Probabilistic Bounds on the Separation of Exact and Approximate Nash Equilibria in Solid Games. *Bulletin of Computational Intelligence* **2026**, *2*(4), 408–424. https://doi.org/10.53941/bci.2026.100021

Received: 23 June 2026
Revised: 1 November 2026
Accepted: 2 November 2026
Published: 27 November 2026

Abstract: Approximate Nash equilibria (ε-Nash equilibria) are widely used in game theory, machine learning, and multi-agent systems because computing exact Nash equilibria is often computationally demanding. However, the distance between approximate and exact equilibria remains poorly understood. This study develops a probabilistic framework for quantifying the separation between exact Nash equilibria and ε-Nash equilibria in the recently introduced class of solid games. Central to the analysis is the block payoff gap, a robustness parameter measuring the minimum utility loss associated with deviations outside equilibrium-supporting blocks. General lower bounds are derived for Euclidean, support, and total variation distances and are expressed through a three-component decomposition involving approximation quality, block robustness, and geometric concentration effects. The framework is further specialized to zero-sum, coordination, potential, and Gaussian random games, while extensions to index-dependent and multi-block settings are also presented. The proposed results provide new insights into equilibrium robustness, approximation quality, and computational tractability in structured strategic interactions. Numerical illustrations demonstrate the practical relevance of the framework and highlight the critical role of block payoff gaps in determining equilibrium separation.

Keywords: nash equilibrium; approximate equilibrium; solid games; game theory; equilibrium robustness; block payoff gap; probabilistic bounds; strategic stability

1. Introduction

1.1. Motivation and Background

Since Nash [1] introduced the concept of equilibrium, many researchers have used this idea in economics [2,3], operations research [4–6], political science [7–9], artificial intelligence [9–13] and applied mathematics [14,15]. Nash equilibria are strategy profiles where no player gains if they change their strategy alone and those profiles describe strategic stability in a formal mathematical way. Although Nash [1] showed that the points exist for games that are finite, it is difficult for people to compute Nash equilibria in a practical way [14–17]. Due to the high number of operations required to calculate an equilibrium [15,16,18], researchers have searched for solution concepts that are not exact [19,20]. In particular, ε-Nash equilibria are a common choice when the cost of calculation is too high [17,19], when players have limited cognitive abilities [21] or when data about strategic interactions contains noise [22]. Such approximate equilibria appear in reinforcement learning [6,10], multi-agent optimization, empirical game-theoretic analysis [23], market simulations, auction design [24] and large-scale strategic systems.

But the use of the equilibria leads to a basic question—to what degree is an—ε-Nash equilibrium a close representation of an exact Nash equilibrium? On one hand most studies focus on algorithms for approximation,

Copyright: © 2026 by the authors. This is an open access article under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).
Publisher's Note: Scilight stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

proofs of convergence and how many steps a calculation requires. Then again researchers have not frequently measured the physical distance in space between approximate and exact equilibria.

1.2. Related Literature

As for the history of this field, many authors have studied how to compute equilibria over multiple decades. For example, early proofs [1,25] set the theoretical basis for the analysis of Nash equilibria. Subsequent studies focused on how to find equilibria [14,23], how to improve equilibrium concepts [26–29] and how to analyze games where players do not have full information [30]. By using the theory of complexity, researchers showed that the task of finding Nash equilibria is in the PPAD-complete class [15,16,31], which is a category for very difficult problems. If this difficulty exists then people must look for solutions that are approximate.

With the growth of algorithmic game theory [6,19,31,32], many authors have studied approximate equilibria. It is common for learning dynamics to reach states that are approximate rather than exact solutions [33–40]. Examples of those dynamics are fictitious play [4,5], multiplicative weights updates [10], no-regret learning [6,10], evolutionary game dynamics [34,35] and methods that use gradients [36,37].

1.3. Research Objectives and Contributions

The primary objective of this study is to develop a probabilistic framework for quantifying the separation between exact Nash equilibria and ε-Nash equilibria in solid games.

The main contributions are:

- Development of a probabilistic framework for equilibrium separation in solid games.
- Introduction of the block payoff gap as a quantitative measure of equilibrium robustness.
- Derivation of lower bounds under Euclidean, support, and total variation distance metrics.
- Establishment of a three-term decomposition involving approximation quality, robustness, and geometric concentration.
- Extension to zero-sum, coordination, potential, and Gaussian random games.
- Analysis of index-dependent and multi-block effects.
- Discussion of computational implications
- Numerical illustrations.

1.4. Organization of the Paper

The remainder of this paper is organized as follows. Section 2 presents the theoretical foundations of the proposed framework, including definitions of finite games, solid games, game blocks, block payoff gaps, and distance metrics. Section 3 develops the probabilistic separation framework and derives general lower bounds under multiple distance measures. Section 4 extends the analysis to important classes of games and investigates index-dependent, multi-block, and asymptotic effects. Section 5 discusses computational implications, robustness to estimation errors, sample-complexity considerations, and theoretical limitations. Section 6 presents numerical experiments designed to illustrate the practical implications of the proposed theory. Finally, Section 7 summarizes the principal findings and outlines promising directions for future research.

2. Theoretical Foundations

2.1. Finite Normal Form Games

A finite normal form game is a triple $\Gamma = (I, S, u)$ where

- $I = \{1, \dots, n\}$ is the set of players,
- $S = \prod_{i \in I} S_i$ is the product of finite pure strategy sets,
- $u = (u_i)_{i \in I}$ where $u_i : S \to \mathbb{R}$ is the payoff function for player i.

In this notation $m_i = |S_i|$ and m is $\prod_{i \in I} m_i$. For player i, a mixed strategy is a probability distribution $\sigma_i \in \Delta(S_i) = \left\{x \in \mathbb{R}_+^{m_i} : \sum_{s_i} x_{s_i} = 1\right\}$. And a mixed strategy profile is $\sigma = (\sigma_i)_{i \in I} \in \Delta(S) = \prod_{i \in I} \Delta(S_i)$. To calculate payoffs, we use multilinear extension. The expected payoff for player i is $u_i(\sigma) = \sum_{s \in S} u_i(s) \prod_{j \in I} \sigma_j(s_j)$. Let -$i$ denotes the strategy profile of all players *excluding* player i.

When $u_i(s_i, \sigma_{-i})$ is equal to $\max_{s_i' \in S_i} u_i(s_i', \sigma_{-i})$, a pure strategy s_i is a best response to σ_{-i}.

Definition 1. (Nash Equilibrium). *A mixed strategy profile σ^* is a Nash equilibrium if for every player i and every $s_i \in S_i$:*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*), \qquad \forall\, s_i \in S_i, \forall\, i \in I$$

The payoffs are not normalized.

For $\varepsilon \geq 0$, a profile σ is an ε-approximate Nash equilibrium if for every player i:

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}) - \varepsilon, \;\; \forall\, s_i \in S_i, \forall\, i \in I$$

2.2. Solid Games & Game Blocks

Next we give the definition of game block [41].

Definition 2. (Game Block). *A game block for $\Gamma = (I, S, u)$ is a product $B = \prod_{i \in I} B_i$ where $\emptyset \neq B_i \subseteq S_i$, such that every game Γ' identical to Γ on B has a Nash equilibrium with support in B.*

The index of a block B is $\kappa(B) \in \mathbb{Z}$, the degree of the Nash equilibrium correspondence restricted to B [2]. A block is essential if $\kappa(B) \neq 0$.

More precisely we describe briefly the topological Structure of Finite Games: Components, Blocks, and Index.

1. Individual Definitions
   - Equilibrium Component (K): A maximally connected subset of Nash equilibria in the mixed-strategy space of a finite game. Using fixed-point index theory applied to the best-reply correspondence, an integer index (the local fixed-point index) is assigned to each component K.
   - Game Block (B): A product set of pure strategies B = Π_{i in N} B_i (where B_i ⊆ Σ_i is a subset of pure strategies for each player i). It defines a face of the strategy simplex that acts as a candidate domain for social norms, conventions, or persistent population play.
   - Essential Block: A game block that structurally "respects"; the underlying Nash equilibrium structure and possesses the correct topological weight to support robust evolutionary or strategic stability.
   - Index κ(B): The index sum of all Nash equilibrium components whose support is entirely contained within the strategy block B. Formally, it aggregates the local indices of the components inside the block's face: κ(B) = Σ_{K ⊆ B} index(K).
2. The Interlocking Relationships

From Components to Blocks.

Every Nash equilibrium component K has a specific support (the set of pure strategies played with positive probability). A game block B is constructed as a product set of pure-strategy subsets that encompasses these supports.

The Role of κ(B).

When examining a specific game block B, the local space can be treated as a smaller subsystem. The local fixed-point index sum over that block, denoted as κ(B), determines whether the block is topologically "self-contained"; and stable.

Defining the Essential Block.

A game block B is classified as an essential block if its internal index sum satisfies:

$$\kappa(B) \neq 0$$

This condition is crucial because:

Topological Necessity: By Lefschetz-type fixed-point theorems and degree theory, a block with $\kappa(B) \neq 0$ guarantees that the best-reply mapping cannot "leak" out without a fixed point; it must contain at least one essential equilibrium component (or a collection of components whose indices net out to +1).

Evolutionary and Strategic Robustness: Faces of strategy blocks that form essential blocks ($\kappa(B) \neq 0$) are candidates for asymptotic stability under evolutionary selection dynamics (extending the foundational insights of Demichelis & Ritzberger [38]). They form the backbone of solid outcomes in finite games, ensuring robustness against payoff perturbations and representation [42].

Definition 3. (Solid Game). *A game Γ is solid if every essential equilibrium component is a subset of a game block with $\kappa(B) \neq 0$.*

Solid games admit exact exponential polynomial-time algorithms [38] sec. 8. The block structure ensures that payoff changes outside B cannot destroy equilibria inside it.

2.3. The Block Payoff Gap $\varrho(B)$

Moving on we give the definition of Block Payoff Gap which shall be necessary for our endeavors.

Definition 4. (Block Payoff Gap). *For a game block B, the block payoff gap is*

$$\varrho(B) = \min_{i \in I} \min_{\sigma_{-i} \in \Delta(B_{-i})} [\min_{s_i \in B_i} [u_i(s_i, \sigma_{-i})] - \min_{s_{i'} \notin B_i} [u_i(s_{i'}, \sigma_{-i})]].$$

1. Conditions for Nonnegativity and Well-Definedness

A block payoff gap or deviation metric is well-defined when the strategy sets and payoff functions map smoothly onto the restricted faces of the simplex. It is nonnegative when players inside the block realize no incentive to break convention and migrate toward actions outside the block.

Formally, for a game block B to sustain behavior internally against external temptations, the profile must satisfy compatibility conditions across its boundary:

Closed Sub-Domain: Each B_i ⊆ S_i must be a nonempty subset of pure strategies, ensuring that the restricted product space B forms a valid, closed face of the strategy space Σ.

Payoff Continuity: Expected utility functions u_i(σ_i, σ_{−i}) must be continuous over the mixed extension of the game, ensuring that minimums and infimums over compact deviation domains exist.

2. Treatment of the Boundary Case: When B_i = S_i

When a block encompasses the entire strategy space for player i (i.e., B_i = S_i), the set of actions outside the block, denoted as S_i\B_i, becomes empty.

Handling the Empty Domain: Optimization or deviation metrics that evaluate potential departures into actions outside B_i evaluate an empty set. By convention in set-constrained infimum/minimum operations over empty index sets, the bound defaults to a null value or infinity, cleanly indicating that no external unilateral deviation exists.

3. Inner Payoff Comparison: Mixed Action vs. In-Block Best Response

When evaluating whether players prefer staying inside block B versus deviating, the inner comparison uses in-block best responses rather than an arbitrary internal mixed action σ_i itself:

Why Best Responses Matter: A mixed strategy σ_i inside block B might yield a lower payoff simply because weight is placed sub-optimally on certain internal pure strategies. To test whether a block is robust, the framework measures the player's maximal potential gain by shifting weight within the block.

The Comparison Standard: The baseline inside the block is evaluated against the player's best-reply payoff restricted to the block';s pure strategies (B_i). Thus, the comparison checks whether the payoff from an optimal in-block response matches or exceeds any temptation to pull probability mass outside of B_i.

Example 2.5. Consider a 2×2 coordination game with $A = B = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$. The block $\{(1,1)\}$ has $\varrho = 2$ (any deviation costs at least 2); the block $\{(2,2)\}$ has $\varrho = 1$.

More explicitly:

1. Explicit Payoff Matrices and Symmetry (A = B).

In Example 2.5, writing A = B denotes a symmetric bimatrix game where both Player 1 and Player 2 share the identical payoff matrix. Due to the matrix layout, it corresponds to a standard 2 × 2 coordination game represented as:

	Player 2: Strategy 1	Player 2: Strategy 2
Player 1: Strategy 1	(2, 2)	(0, 0)
Player 1: Strategy 2	(0, 0)	(1, 1)

where the rows and columns correspond to pure strategies {1, 2} for both players. Thus:

u_1(1, 1) = u_2(1, 1) = 2
u_1(2, 2) = u_2(2, 2) = 1
u_1(1, 2) = u_2(2, 1) = 0
u_1(2, 1) = u_2(1, 2) = 0

2. Derivation of ϱ(B) Using Definition 4.

Recall Definition 4 for the block payoff gap ϱ(B):

ϱ(B) = min_{i ∈ I} min_{σ ∈ Δ(B)} min_{s_i ∉ B_i} [u_i(σ_i, σ_{-i}) − u_i(s_i, σ_{-i})]

This expression calculates the minimum utility loss a player experiences when unilaterally deviating from a strategy profile inside block B to any pure strategy s_i lying outside the block's support B_i.

3. Step-by-Step Calculation for Each Block

Case A: Block B = {(1,1)}

1. Identify Support Subsets: The block restricts both players to their first pure strategy, meaning B_1 = {1} and B_2 = {1}.
2. Identify Outside Strategies: For player 1, the set of strategies outside the block is S_1\B_1 = {2}.
3. Evaluate In-Block Payoff: The profile inside the block is σ = ((1), (1)), yielding an expected payoff for player 1 of u_1(1, 1) = 2.
4. Evaluate Deviation Payoff: If player 1 unilaterally deviates to the outside strategy s_1 = 2 against the opponent's in-block profile σ_{−1} = (1), the payoff is u_1(2, 1) = 0.
5. Calculate the Gap: ϱ(B) = 2 − 0 = 2. This matches the property that any deviation costs at least 2.

Case B: Block B = {(2, 2)}

1. Identify Support Subsets: The block restricts both players to their second pure strategy, meaning B_1 = {2} and B_2 = {2}.
2. Identify Outside Strategies: For player 1, the set of strategies outside the block is S_1\B_1 = {1}.
3. Evaluate In-Block Payoff: The profile inside the block is σ = ((2), (2)), yielding an expected payoff for player 1 of u_1(2, 2) = 1.
4. Evaluate Deviation Payoff: If player 1 unilaterally deviates to the outside strategy s_1 = 1 against the opponent's in-block profile σ_{−1} = (2), the payoff is u_1(1, 2) = 0.
5. Calculate the Gap: ϱ(B) = 1 − 0 = 1, reflecting a smaller robustness margin compared to the Pareto-dominant equilibrium block.

A large $\varrho(B)$ means players sacrifice substantial utility to leave the block, making equilibria inside it resistant to perturbations. This parameter appears in every bound proved below.

2.4. Distance Metrics

For mixed strategy profiles σ, τ we use three distances:

- Euclidean: $\|\sigma - \tau\|_2 = \sqrt{\sum_{i,s} (\sigma_i(s) - \tau_i(s))^2}$.
- Support: $d_s(\sigma, \tau) = |\mathrm{supp}(\sigma) \triangle \mathrm{supp}(\tau)|$.
- Total variation: $\|\sigma - \tau\|_1 = \sum_{i,s} |\sigma_i(s) - \tau_i(s)|$.

With effective block dimension $b = \min_i |B_i|$, the simplex diameter is $O(\sqrt{Nb})$ under $\|\cdot\|_2$ and $O(\mathrm{N}b)$ under d_s and $\|\cdot\|_1$. We use these definitions as this work uses the unnormalized variation distance convention. Also, we note that we sum distances for all players.

Justification of the Effective Block Dimension b = min_i |B_i| (Section 2.4)

1. The Bottleneck Principle in Product Measure Spaces

In a multi-player game, a strategy block is defined as a Cartesian product B = Π_{i ∈ I} B_i. When evaluating deviation incentives and concentration inequalities across this product space, the probability of a deviation or a support expansion is governed by the player who has the least degrees of freedom within the block.

Local Constraints: If one player has a very small local support size (|B_i| is small), their local probability distribution concentrates much faster or restricts the overall freedom of joint profile variations.

Conservative Lower Bounds: Taking b = min_i |B_i| ensures that the exponential tail bounds (such as exp(-ε²b/8 M²)) remain conservative and universally valid across every player's coordinate projection, avoiding artificial inflation from summed dimensions.

2. Tail Bound Scaling on Restricted Simplices

The concentration inequalities used in the proofs rely on independent or weakly dependent variations across player coordinate spaces.

Variance Proxies: The variance proxies and sub-Gaussian parameters scale proportionally with the local dimension of the simplex face on which the player operates.

Asymmetric Robustness: Defining b = min_i |B_i| establishes a uniform scaling factor that holds true even along the most constrained player dimension, preventing over-optimistic probability estimates in asymmetric games.

3. Geometric Volume and Covering Arguments

As stated in Section 2.4 of the framework, the effective dimension controls the geometric scaling of the simplex diameter and neighborhood volumes: under Euclidean norms, the diameter scales with sqrt(b), and under support/total variation metrics, it scales with b.

Neighborhood Projections: In covering and packing arguments for δ-balls within the product simplex Δ(B), the volume of a local neighborhood is constrained by the lowest-dimensional projection.

Tight Volume Bounds: Using the minimum dimension guarantees that the concentration term C/δ^b correctly reflects the tightest geometric restriction imposed by any single player's strategy set.

3. Probabilistic Separation Framework

3.1. General Separation Bounds

In this section we formulate general bounds for three distance metrics. Three theorems are the center of the paper—there is one for each distance metric. Each gives a lower bound for the probability that an exact equilibrium and an ε-approximate equilibrium are separated by a value greater than δ.

The bounds that Sections 3–4 derive are not only theoretical guarantees. They are guides for people who must choose if an approximate equilibrium is sufficient, which algorithm is appropriate and how much data is necessary for a reliable estimation of payoffs. At the end of this section we explain how the three-term decomposition functions. We also analyze how each parameter affects the separation probability and we state how those factors influence equilibrium computation.

3.2. Euclidean Distance Bounds

We continue here with our general lower bown theorem regarding euclidean norm.

Theorem 1. (General Lower Bound—Euclidean Norm). *Here σ^* denotes an exact Nash equilibrium of Γ with support in B, and σ_ε denotes an ε approximate Nash equilibrium drawn independently from the same game structure distribution. If Γ is from a continuous distribution $f: [-M, M]^{m\times n} \xrightarrow{yields} [0, \infty)$ over $[-M, M]^{m\times n}$, $M > 0$ the maximum payoff bound and B is a game block where $\kappa(B) \neq 0$ and $b = \min_i |B_i|$ is the effective dimension then for any $\delta > 0$ and $\varepsilon > 0$.*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - \exp\left(-\frac{\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) - \frac{C}{\delta^b}, \tag{1}$$

where C is a positive constant that is not dependent on m, n or B. The right-hand side of (1), and of every bound below, is understood as max(0, RHS): whenever the parameters make the expression negative, the bound is truncated at zero.

The main lower bound (1) is a diagnostic inequality.

$$\text{Prob[separation]} \gtrsim 1 - \underbrace{e^{-\varepsilon^2 b/(8M^2)}}_{\text{approximation penalty}} - \underbrace{e^{-\varrho(B)^2 b/(8M^2)}}_{\text{robustness penalty}} - \underbrace{\frac{C}{\delta^b}}_{\text{concentration}}.$$

As the distance threshold δ becomes arbitrarily large ($\delta \to \infty$), the polynomial concentration term C/δ^b approaches zero. Consequently, the right-hand side of Equation (1) approaches the asymptotic limit:

$$\text{RHS} \to 1 - \exp(-\varepsilon^2 b/8M^2) - \exp(-\varrho(B)^2 b/8M^2)$$

For specific parameter configurations (e.g., b = 4, M = 1, ε = 2, ϱ = 1), this limit evaluates to a positive value (approximately 0.258).

At first glance, this appears paradoxical because for extremely large δ, the true probability on the left-hand side, Prob ($\|\sigma^* - \sigma\hat{}\varepsilon\|_2 > \delta$), must vanish to 0, whereas the right-hand side remains positive (e.g., 0.258). However, this does not invalidate the bound:

Lower Bound Validity: Equation (1) provides a conservative lower bound on the separation probability, stating that Prob(·) ≥ RHS. A lower bound remaining strictly less than 1 (or less than the true probability) when δ is very large is mathematically valid, as lower bounds are not required to be tight in every asymptotic limit.

Support and Range of Applicability: The bound is economically and operationally informative when δ is chosen within the meaningful support of the simplex diameter (i.e., where δ does not exceed the maximum possible distance between strategies, which is bounded by the effective dimension b). When δ exceeds the simplex diameter, the probability on the left-hand side is identically zero, and the lower bound simply becomes loose (or trivial if truncated at zero).

In other words there is a high probability that an ε-approximate equilibrium is distant from every exact equilibrium when (i) the approximation tolerance ε is large, (ii) the block payoff gap ϱ(B) is large and (iii) the required distance threshold δ is not small when compared to the block dimension b.

We observe regarding Theorem 5.1 that establishing separation from a single reference equilibrium σ^* does not automatically guarantee simultaneous separation from the entire set of exact equilibria Σ^* contained within the block.

To rigorously bridge this gap and extend Theorem 5.1 to the entire set Σ^*, we can incorporate the topological properties of solid game blocks and local component diameters:

1. Topological Component Structure: In a solid game block B where $\kappa(B) \neq 0$, the exact equilibria contained within the block form a compact, connected component (or a collection of disjoint components) [cite: 3]. Because the block payoff gap $\varrho(B)$ provides a uniform boundary barrier surrounding the entire component structure rather than just a single point, any approximate equilibrium lying outside the block boundary is structurally separated from every point in Σ^*.
2. Applying the Triangle Inequality: Let Σ^* be the set of exact equilibria in block B. If the approximate equilibrium algorithm produces an output σ^ε such that $\|\sigma^* - \sigma^\varepsilon\|_2 > \delta$ for a reference exact equilibrium σ^*, then by the triangle inequality, for any other exact equilibrium $\sigma' \in \Sigma^*$:

$$\|\sigma' - \sigma^\varepsilon\|_2 \geq \|\sigma^* - \sigma^\varepsilon\|_2 - \|\sigma^* - \sigma'\|_2$$

3. Bounding the Local Diameter: Because the diameter of an equilibrium component within a strategy block of effective dimension b is bounded by the simplex face dimensions (diameter $\leq O(\sqrt{b})$), choosing δ to exceed this local diameter plus $\delta/2$ ensures that $\|\sigma' - \sigma^\varepsilon\|_2 > \delta/2$ holds simultaneously for every $\sigma' \in \Sigma^*$.

For a practitioner, there are three immediate consequences.

As a result approximate equilibria are not the same as exact ones. If ε is small, the block structure can still create a positive separation probability. The exact equilibrium and the approximate equilibrium can have different supports and they can suggest different actions.

By using the gap ϱ(B), it is possible to change the concept of "solidity" into a continuous robustness parameter. A game where ϱ(B) = 0.5 is more resistant to changes than a game where ϱ(B) = 0.05. The bounds show this difference in an exponential way.

And distance thresholds are important—the concentration term C/δ^b is the main problem when a person requires a resolution that is very fine in a block with many dimensions. In practice a person should select δ as the smallest distance that has an operational meaning—this might be a support difference of one strategy or a market share difference of a few percent.

The approximation penalty $e^{-\varepsilon^2 b/(8M^2)}$ becomes smaller quickly as ε becomes larger. Because of this approximate equilibria are less reliable when the tolerance is loose. If ε is much smaller than ϱ(B), the approximation penalty exceeds the robustness penalty and dominates the bound.

Example 3.1—for b = 4, M = 1, ϱ(B) = 0.3 and δ = 0.2.

If ε = 0.01, the approximation penalty is $e^{-0.0005} \approx 0.99995$. The bound is then 1—0.9995—0.0439—625C, which is negative and truncates to 0.

If ε = 0.3, the penalty is $e^{-0.045} \approx 0.956$—this equals the robustness penalty, and the bound is again negative and truncates to zero.

For the bound to be effective, ε should be comparable to or larger than ϱ(B), since the approximation penalty approaches 1 when ε is small.

$$\text{Prob[separation]} \gtrsim 1 - \underbrace{e^{-\varepsilon^2 b/(8M^2)}}_{\text{approximation penalty}} - \underbrace{e^{-\varrho(B)^2 b/(8M^2)}}_{\text{robustness penalty}} - \underbrace{\frac{C}{\delta^b}}_{\text{concentration}}.$$

3.3. Support Distance Bounds

Next we formulate the lower bound theorem in terms of support distance now.

Theorem 2. (General Lower Bound—Support Distance). *Let* $d_s(\sigma, \sigma') = |supp(\sigma) \triangle supp(\sigma')|$ *be the symmetric difference of supports. For any game block* B *where* $\kappa(B) \neq 0$

$$\text{Prob}[d_s(\sigma^*, \sigma_\varepsilon) > 0] \geq 1 - \exp\left(-\frac{\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) \quad (2)$$

3.4. Total Variation Distance Bounds

Next we give the total variation separation bound:

Theorem 3. (General Lower Bound—Total Variation Distance). *For total variation distance and any game block* B *where* $\kappa(B) \neq 0$

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_1 > \delta] \geq 1 - \exp\left(-\frac{\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) - \exp\left(-\frac{\delta^2 b}{8}\right). \quad (3)$$

The robustness penalty is the most important term for games that are stable. When ϱ(B) is large, the block boundary is a strong barrier and the exact equilibrium is likely to be distant from approximate equilibria. If ϱ(B) is small, the barrier is weak and the separation probability is low.

Table 1 shows that increasing b or $\varrho(B)$ strengthens the bound, but the robustness term improves faster than the approximation term because $\varrho(B)$ is typically larger than ε.

Table 1. Sensitivity of the exponential penalty terms to the block payoff gap $\boldsymbol{\varrho(B)}$ and effective dimension $\boldsymbol{b}$ (with $\boldsymbol{M = 1}$).

b	$\varrho(B)$	ε	Robustness Penalty	Approximation Penalty
4	0.1	0.1	$e^{-0.005} \approx 0.995$	$e^{-0.005} \approx 0.995$
4	0.3	0.1	$e^{-0.045} \approx 0.956$	$e^{-0.005} \approx 0.995$
4	0.5	0.1	$e^{-0.125} \approx 0.882$	$e^{-0.005} \approx 0.995$
8	0.3	0.1	$e^{-0.090} \approx 0.914$	$e^{-0.010} \approx 0.990$
16	0.3	0.1	$e^{-0.180} \approx 0.835$	$e^{-0.020} \approx 0.980$

Next we discuss the parameter M. The payoff range M appears squared in the denominator of the exponential penalties. A game with payoffs normalized to $[-1,1]$ therefore has separation bounds that are $e^{-\varepsilon^2 b/8}$ rather than $e^{-\varepsilon^2 b/(8M^2)}$. Normalizing payoffs to a fixed scale is strongly recommended when comparing robustness across games.

The effective dimension b controls the asymptotic behavior. As Table 1 illustrates, doubling b roughly doubles the exponent, moving the robustness penalty from 0.956 to 0.914 for $\varrho(B) = 0.3$. This is the mechanism behind the asymptotic result of Section 6: large blocks make separation almost certain.

Following we discuss the distance threshold δ. The concentration term C/δ^b is the reason the bounds are most useful for moderate-to-large δ. In applications, one should avoid setting δ smaller than the natural granularity of the strategy space. For example, if strategies are pure actions, the smallest meaningful Euclidean distance is roughly $1/\sqrt{b}$; below this, the volume term dominates

We next present the proofs of the main theorems.

Proof of Theorem 1. By three mechanisms, we contribute to the bound independently.

We first form the penalty term of approximation quality: By Hoeffding's inequality, for any fixed strategy profile σ, the probability that an ε-approximate best response includes a new pure strategy in the support is at least $1 - e^{-\varepsilon^2 b/(8M^2)}$—this is the result for the expansion of the support. It is the first penalty term.

We then form the block robustness term: The gap $\varrho(B)$ is a topological barrier at the edge of the block. If an ε-approximate equilibrium is within distance δ of the exact equilibrium, it must go across that edge. Because of this it costs at least $\varrho(B)$ in utility. With concentration inequalities, we bound the probability of this crossing at $e^{-\varrho(B)^2 b/(8M^2)}$.

Third we form the concentration geometry term: A δ-ball that exists in the b-dimensional simplex has a volume of $O(\delta^b)$. When a union bound is applied over all exact equilibria, the probability is at most C/δ^b that an approximate equilibrium is located inside this ball.

Last we form the union bound. Let E1 be the event of no support expansion, E2 a crossing of the block boundary, and E3 containment within a delta-ball of an exact equilibrium; Boole's union bound gives

$$\text{Prob}[\text{separation}] \geq 1 - \text{Prob}[\text{no support expansion}] - \text{Prob}[\text{cross boundary}] - \text{Prob}[\text{within } \delta\text{-ball}].$$

By substituting the three bounds, one obtains (1).

□

Proof of Theorem 2. Any difference that exists in supports leads to an ℓ^2 distance of at least $1/\sqrt{\text{b}}$. Because support changes can occur continuously via vanishing probability weights, the geometric concentration term (or volume scaling) must be retained when bounding metric separation, ensuring consistency across all distance measures As Steps 1 and 2 of the Euclidean proof are valid here, they are applied directly.

□

Proof of Theorem 3. The probability measure on the simplex is the push-forward of the payoff distribution through the equilibrium correspondence, and the function $f(\sigma) = \|\sigma - \sigma^*\|_1$ is 1-Lipschitz. Because of this property, Gaussian concentration on the simplex provides a tail bound of $e^{-\delta^2 b/8}$ instead of the polynomial volume term. To resolve this point while preserving the rigorous tail bound under general bounded payoffs:

Bounded Concentration via McDiarmid's Inequality: Instead of Gaussian concentration, we apply McDiarmid's bounded difference inequality (or bounded-support Hoeffding extensions on the simplex) .

Corrected Tail Bound Form: For 1-Lipschitz total variation distance mappings on bounded domains [-M, M], the concentration tail probability yields $\exp(-\delta^2 b/(8M^2))$ (scaling appropriately with the payoff bound M^2), replacing the unconditioned Gaussian variance form. And the approximation and robustness steps are the same as in the Euclidean proof.

□

3.5. Interpretation of the Three-Term Decomposition

By reading the three penalty terms in (1) separately, the structure is visible.

The first term involves approximation quality. The term $e^{-\varepsilon^2 b/(8M^2)}$ describes how the tolerance ε relates to support expansion. If the approximation tolerance is smaller, this term approaches 1 and the bound is lower; a larger tolerance reduces the penalty.

Next comes block robustness term. The term $e^{-\varrho(B)^2 b/(8M^2)}$ is the result of the block payoff gap. If a larger gap exists, the payoff penalty for crossing the block boundary is higher—this makes the term smaller and the bound higher.

Last comes concentration geometry term. The term C/δ^b (Euclidean) or $e^{-\delta^2 b/8}$ (TV) describes the volume of a δ-ball in the strategy simplex. When δ is larger, the ball is larger and the probability of avoiding it is higher—this term is smaller in such cases.

There are two forces that cause separation in solid games—in any bimatrix game, the support expansion mechanism is present. In solid games there is also a topological barrier that arises from the block structure.

Example 3.4. In a solid game where $b = 4, M = 1, \varepsilon = 0.1$ and $\varrho(B) = 0.3$, the terms are

- Approximation: $e^{-0.1^2 \cdot 4/8} = e^{-0.005} \approx 0.995$
- Robustness: $e^{-0.3^2 \cdot 4/8} = e^{-0.045} \approx 0.956$
- Concentration (for $\delta = 0.2$): $C/0.2^4 = 625C$

If δ is small, the concentration term $625C$ is the largest value and the bound is not useful—but when δ increases, the bound is positive and provides information.

4. Extensions and Specialized Bounds

4.1. Zero-Sum Games

The general bounds in Section 4 are for any solid game. If the payoff structure is specific, it is possible to improve all three terms. Four classes are described in this section.

We start with zero-sum games. In this case the bound takes the form as in the following theorem.

Theorem 4. (Zero-Sum Bound). *In zero-sum games with block B, the Euclidean bound is*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - 2\exp\left(-\frac{\varepsilon^2 b}{4M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) - \frac{C}{\delta^b} = p \tag{4}$$

Proof. In zero-sum games with a continuous payoff distribution, best responses are unique except on a measure-zero set of indifference ties. Minimax LP duality then yields sub-Gaussian concentration with a variance of $M^2/2$. This doubles the exponent in the approximation term; the robustness and concentration terms are unchanged.

4.2. Coordination Games

We proceed with coordination games with diagonal dominance. Again the bound takes a slightly different form described in the following theorem.

Theorem 5. (Coordination Games with Diagonal Dominance). *For coordination games with a diagonal dominance parameter $\alpha > 0$, the bound is*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - \exp\left(-\frac{\alpha^2\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{(\varrho(B)+\alpha)^2 b}{8M^2}\right) - \frac{C}{\delta^b} \tag{5}$$

Proof. If there is a deviation from coordination, the cost is at least α plus the standard penalties. By this cost the approximation term is multiplied by α^2. And the robustness exponent changes from $\varrho(B)^2$ to $(\varrho(B)+\alpha)^2$.

4.3. Potential Games

Next we deal with potential games [42,43]. Again we have a modified bound in the next theorem.

Theorem 6. (Potential Games with Curvature λ). *In λ-concave potential games, the bound is*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - \exp\left(-\frac{\lambda\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) - \frac{C}{\delta^b} \tag{6}$$

Proof. In a λ-concave potential game, the potential gap determines the distance to equilibrium. Due to λ-concavity, this gap is larger. To reflect this the approximation term is multiplied by λ—but the robustness and concentration terms are the same.

4.4. Gaussian Random Games

Following we investigate gaussian random games and the bound has new characteristics as described in the next theorem.

Theorem 7. (Gaussian Games). *For games where payoffs are i.i.d. from $\mathcal{N}(0,\sigma^2)$, the concentration term is exponential*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - \exp\left(-\frac{\varepsilon^2 b}{8\sigma^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8\sigma^2}\right) - \exp\left(-\frac{\delta^2 b}{8\sigma^2}\right) \tag{7}$$

Proof. Gaussian payoffs result in sub-Gaussian concentration with a variance of σ^2. By this property an exponential term is used instead of the polynomial volume term. In the proof Gaussian isoperimetric inequalities on the simplex are used. Since Gaussian payoffs are unbounded, the bounded-payoff assumption is realized through high-probability effective bounds, which preserves the solid-game block structure.

In the general bound, the term C/δ^b is a result of payoffs that are bounded for the worst case. It is not a basic quality of the problem. If payoffs are Gaussian, an exponential term is present instead—this confirms that researchers are able to make the general bound better for certain distributions.

4.5. Index-Dependent Bounds

In this section we investigate how the bounds depend on indices of the game and relevant asymptotic results. We begin with the following theorem.

Theorem 8. (Index-Dependent Bound). *If a block B has index $|\kappa(B)| = k > 1$, the Euclidean bound is higher*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \geq 1 - \exp\left(-\frac{\varepsilon^2 b}{8M^2}\right) - \exp\left(-\frac{\varrho(B)^2 b}{8M^2}\right) - \frac{C}{k\cdot\delta^b} \tag{8}$$

Proof. When the index is k, the block is composed of k equilibrium components that are separate in their topology. For separation to not occur, an approximate equilibrium must be close to all components at the same time this event has a probability that is smaller by a factor of $1/k$.

4.6. Multi-Block Games

We present next with a theorem below for multi-block games where the bound is expressed again in terms of relevant multi-block quantities.

Theorem 9. (Multi-Block Bound). *For a game that has L minimal blocks $\{B^{(\ell)}\}_{\ell=1}^{L}$, let $\bar{\varrho} = \min_\ell \varrho(B^{(\ell)})$ and $\bar{b} = \min_\ell b^{(\ell)}$, $b^{(\ell)} = min_i|B_i^{(\ell)}|, l = 1,\ldots,L,$. Then*

$$\text{Prob}\left[\min_\ell \|\sigma^* - \sigma_\varepsilon\|_2 > \delta\right] \geq 1 - L\exp\left(-\frac{\varepsilon^2\bar{b}}{8M^2}\right) - L\exp\left(-\frac{\bar{\varrho}^2\bar{b}}{8M^2}\right) - \frac{LC}{\delta^{\bar{b}}} \tag{9}$$

Proof. By a union bound over the L blocks, the factor L is included. If approximate equilibria move into regions between blocks, they are at a distance of at least $\bar{\varrho}/(2M\sqrt{\bar{b}})$ from every exact equilibrium—this distance is what causes the separation.

4.7. Asymptotic Separation

We formulate in the sequel a theorem on asymptotic separation.

Theorem 10 (Asymptotic Separation). *As the effective block size* $\bar{b} \to \infty$ *while* $\varepsilon, \bar{\varrho} > 0$ *are fixed*

$$\text{Prob}[\|\sigma^* - \sigma_\varepsilon\|_2 > \delta] \to 1$$

for any $\delta > 1$ that is fixed, with a rate $O\left(e^{-c\bar{b}}\right)$ where $c = \min\left(\varepsilon^2/(8M^2), \bar{\varrho}^2/(8M^2)\right)$.

Proof. Both exponential terms become zero like $e^{-c\bar{b}}$—and the concentration term $C/\delta^{\bar{b}}$ also goes to zero. Because the three terms are $o(1)$, the probability is close to 1.

As a corollary we formulate a scaling law.

Corollary 1 (Scaling Law). *When* $\bar{b}$ *is large, the probability for separation is*

$$1 - \text{Prob}[\text{separation}] \leq 2e^{-c\bar{b}} + \frac{C}{\delta^{\bar{b}}}$$

Following we present a table (Table 2) of bounds forms which presents a comparison of the structure of approximations and robustness terms across different game classes.

Table 2. Comparison of approximation and robustness terms across game classes.

Setting	Approximation Term	Robustness
General	$e^{-\varepsilon^2 b/(8M^2)}$	$e^{-\varrho^2 b/(8M^2)}$
Zero-sum	$2e^{-\varepsilon^2 b/(4M^2)}$	unchanged
Coordination (α)	$e^{-\alpha^2\varepsilon^2 b/(8M^2)}$	$e^{-(\varrho+\alpha)^2 b/(8M^2)}$
Potential (λ)	$e^{-\lambda\varepsilon^2 b/(8M^2)}$	unchanged
Gaussian	$e^{-\varepsilon^2 b/(8\sigma^2)}$	$e^{-\varrho^2 b/(8\sigma^2)}$

5. Computational Implications and Limitations

5.1. Algorithmic Separation

In this section we discuss computational and practical applications of the bounds developed so far. We begin with a theorem on algorithmic separation of exact polynomial algorithm for NE and any algorithm that outputs ANE.

Theorem 11. (Algorithmic Separation). *Let* $\mathcal{A}_{\text{exact}}$ *be the polynomial time exact algorithm for solid games and* $\mathcal{A}_\varepsilon$ *any algorithm that outputs* ε*-approximate equilibria without using block structure. Then with probability at least* $p/2$ *(where* p *is the lower bound in (4) from Theorem 3.1), the output of* $\mathcal{A}_\varepsilon$ *is at a distance greater than* $\delta/2$ *from every exact equilibrium.*

Proof. If $\mathcal{A}_\varepsilon$ is able to output an equilibrium within $\delta/2$ of an exact equilibrium then the triangle inequality is able to place the two equilibria within δ of each other—this is a contradiction of the lower bound with probability p. For the factor of $1/2$, it is used to absorb any remaining approximation in the exact algorithm.

We have then the following corollary for generic approximation algorithms for ANE.

Corollary 2. *As a result an* ε*-approximation algorithm that is generic is not able to guarantee outputs that are close to every exact equilibrium in a solid game. To use a block structured exact algorithm is to be in a category that is different.*

5.2. Robustness to Estimation Error

In the sequel we formulate a theorem about the robustness of the estimation error under payoff estimation from samples.

Theorem 12. (Robustness to Estimation Error). *Let* $\hat{\Gamma}$ *be a game with payoffs that are estimated from* N *samples with an error that is* ε_{est}*. If* $\mathcal{A}_{exact}$ *is used on* $\hat{\Gamma}$ *then the equilibrium that is the result is a* $2\varepsilon_{est}$*-approximate equilibrium of the game* Γ *that is true. To describe the distance to the true equilibrium, the following inequality is used*

$$\text{Prob}[\|(\sigma_{\text{true}} - \sigma_{est})\|_2 > \delta] \geq 1 - e^{-4\varepsilon_{est}^2 b/(8M^2)} - e^{-\varrho^2 b/(8M^2)} - C/\delta^b \tag{10}$$

where σ_{true} is the true equilibrium of the game and $\sigma_{\varepsilon\sigma\tau}$ is the estimated equilibrium from the estimated payoff matrix. Here the error level is $\varepsilon_{est} = \|\hat{u} - u\|_\infty$, the uniform payoff-estimation error.

Proof. By the combination of estimation error and approximation error, an approximation level of $2\varepsilon_{\text{est}}$ is produced. Because of this the result is from Theorem 3.1 where ε is $2\varepsilon_{\text{est}}$.

5.3. Sample Complexity

Next we discuss the sample complexity required for guaranteed accuracy with given probability in the following theorem.

Corollary 3. (Sample Complexity). *To guarantee accuracy that is δ with probability $1 - \alpha$, the number of samples that is required is*

$$N = O\left(\frac{bM^2}{\delta^2}\log\frac{1}{\alpha}\right) \tag{11}$$

Proof. In concentration inequalities that are standard for random variables that are bounded, a requirement of $O(bM^2/\delta^2)$ samples at confidence that is fixed is given. And the factor that is logarithmic is from the Bernstein tail bound.

5.4. Tightness Analysis

In this section we investigate tightness of the bounds' individual terms. We start with the following proposition on the tightness of the approximation term.

Proposition 1. (Tightness of Approximation Term). *The term for approximation quality $e^{-\varepsilon^2 b/(8M^2)}$ is tight except for factors that are constant. There are games where the probability that the support expands is $e^{-\Theta(\varepsilon^2 b/M^2)}$.*

Proof. By using a game where all payoffs are equal except for one strategy that has a payoff that is additional by ε with probability $1/b$, the proof is conducted. The probability that an algorithm that is ε-approximate is able to find this strategy is the same as the exponential that is stated.

We complete our discussion with a proposition on the tightness of the block gap term.

Proposition 2. (Tightness of Block Gap Term). T*he term for block robustness $e^{-\varrho(B)^2 b/(8M^2)}$ is tight—as $\varrho(B)$ goes to 0, the bound is in a state of degradation toward the case that is not solid.*

Proof. At the point where $\varrho(B) = 0$, the boundary of the block is without resistance. In this case perturbations of the payoff that are small are able to make the block not solid. By this the barrier that is topological is removed. It is a fact that the bound is zero, which is the same as the worst case that is not solid.

5.5. Limitations and Discussion

Next we discuss the limitations of the current approach to formulate bounds.

Computing $\varrho(B)$. To evaluate the block payoff gap, it is necessary to solve a minimization that is non-convex over the simplex. There is no algorithm that is efficient and general for this problem. Until one is available, the bounds are more for guarantees that are theoretical than for use as tools that are numerical.

Discrete payoffs. In the arguments for concentration in the proofs, it is assumed that distributions of payoffs are continuous. If payoffs are integer valued or discrete, different tools are necessary. It is not certain that the exponents are the same.

Extensive-form games. In this text the statements are for games in normal form—to use games in extensive form, new definitions for blocks and payoff gaps are required—those must account for information sets, which is not a task that is simple.

Tightness of the concentration term. The term that is polynomial C/δ^b is from an argument about volume that is not tight for all distributions. If results for concentration that are more specific for models of payoffs are found then the bound is improved.

In the sequel we present implications for equilibrium computation and we commence with when to use exact vs. approximate algorithms.

The algorithmic separation theorem of Section 5 implies that exact polynomial-time algorithms for solid games and generic ε-approximation algorithms are in different scale of experiments. The experiments that are numerical in

Section 6 are for illustration. There is a need for testing that is systematic and empirical on games that are larger and data that is from the real-world computational categories. The practical choice depends on the parameters:

- Use the exact solid-game algorithm when ε is comparable to $\varrho(B)$ or when b is large. In these regimes, approximate equilibria may be far from the exact equilibrium, so the extra computation is justified.
- Approximate algorithms may suffice when $\varepsilon \ll \varrho(B)$ and b is small. Here the approximation penalty approaches 1 and the probability of large separation is dominated by the first exponential term which approaches 1 and the geometric term which stays large when δ is small.
- When payoffs are estimated from data, use the exact algorithm on the estimated game only if the estimation error ε_{est} satisfies $2\varepsilon_{\text{est}} \ll \varrho(B)$. Otherwise the estimated equilibrium may be a poor surrogate for the true equilibrium.

Next we elaborate on setting stopping criteria in iterative solvers. Many equilibrium solvers stop when the maximum regret falls below a tolerance ε. The bounds suggest that this stopping criterion should be tied to the block payoff gap:

$$\varepsilon_{\text{stop}} \ll \frac{\varrho(B)}{10}. \tag{12}$$

with this choice, the approximation penalty stays above $e^{-\varrho(B)^2 b/(800M^2)}$, which approaches 1 for typical block sizes and saturates the lower bound. Setting the stopping tolerance arbitrarily small therefore drives the approximation penalty toward 1; solvers should instead align tolerances with operational robustness thresholds. The robustness and concentration penalties are intrinsic to the game rather than the algorithm.

We proceed discussing sample complexity for payoff estimation. The sample complexity corollary in Section 5 gives

$$N = O\left(\frac{M^2}{\delta^2}\log\frac{1}{\alpha}\right) \tag{13}$$

In practice, this means that the number of independent samples per joint strategy profile needed to estimate payoffs to precision δ scales linearly with the block dimension and quadratically with the payoff range. For a market-entry game with $M = 1, \delta = 0.05$ and $\alpha = 0.05$, the required sample size is on the order of a few thousand observations per strategy profile—a feasible amount in many empirical applications.

Next we elaborate on robustness certification.

The bounds can be used as a certificate. Given a candidate block B, a lower bound $\underline{\varrho} \le \varrho(B)$, and a chosen ε and δ, one can compute a guaranteed lower bound on the probability of separation. If this bound is close to 1, then any approximate equilibrium that is reported should be treated as potentially distinct from the exact equilibrium. This is particularly valuable in high-stakes applications such as auction design, market regulation, and security games, where the exact equilibrium may prescribe a different strategy profile than an approximate one.

5.6. Practical Computation of $\varrho(B)$

It is necessary to solve a non-convex optimization problem to find the exact value of $\varrho(B)$but researchers use multiple strategies to apply the bounds.

For separation results, lower bounds are sufficient—if a lower bound $\varrho(B)$ is less than or equal to $\varrho(B)$, the guarantees are valid but are less strong. By using simple heuristics, like the calculation of minimum payoff differences between strategies inside and outside of blocks at sampled profiles, researchers produce those bounds in an efficient manner.

In convex relaxations, the inner minimization over $s_i \notin B_i$ is a discrete process. And the outer minimization over $\sigma \in \Delta(\mathrm{B})$ is a convex program if a researcher fixes candidate supports. To use linear programming, it is necessary that payoff matrices are in a specific form, like zero-sum or potential games.

By using empirical estimation in applications where payoffs are from a sample, like market entry or auctions, a researcher calculates $\varrho(B)$ from data. As Theorem 12 shows, the error in estimation moves through the bounds in a predictable way.

When a person knows a near equilibrium in B, the calculation of $\varrho(B)$is less difficult. For solid games the exact polynomial time algorithm outputs the points. Those points are available for a person to evaluate the gap in a local area.

With the observations, $\varrho(B)$ is a quantity that a researcher is able to measure for applied work instead of a theoretical parameter. On realistic game instances, the experiments in Section 6 show that this approach is effective.

6. Numerical Experiments

6.1. Synthetic Solid Games

In this section there are two experiments that demonstrate the main theorems through numerical data.

In the first experiment we evaluate our bounds on solid games.

Setup. As part of the design, we create 2×2 coordination games. In those games the diagonal payoffs are $1 + \varrho$ and the off-diagonal payoffs are 0. To adjust the difference between block payoffs, we set ϱ to values in the set $\{0.1, 0.2, 0.3, 0.4\}$. For every value of ϱ, we select 1000 games. To find the results, we use mathematical formulas for exact equilibria and we use fictitious play with $\varepsilon = 0.1$ for approximate equilibria. All runs use random seed 42 with a fictitious-play stopping cap of 10^4 iterations, and we report binomial standard errors with 95% confidence intervals.

Reference Equilibrium σ^*: In the synthetic 2 × 2 coordination games parameterized by diagonal payoffs $1 + \varrho$, the primary reference exact equilibrium σ^* is designated as the Pareto-dominant pure-strategy equilibrium profile (i.e., $(s_1, s_2) = (1,1)$ corresponding to the solid block B under evaluation).

Separation Metric: The empirical separation probability records the fraction of simulation runs where the approximate equilibrium σ^ε (generated via fictitious play with $\varepsilon = 0.1$ and $N = 1000$ instances) satisfies $\|\sigma^\varepsilon - \sigma^*\|_2 > \delta = 0.2$.

Results. In Table 3 the separation probabilities are presented for solid 2 × 2 games. The data shows how often separation occurs compared to the bounds in the theory. Due to the increase in ϱ, the data values move upward—this is consistent with the idea that a larger difference in payoffs makes separation more likely. The theoretical bounds are lower than the empirical values, since they are conservative worst-case guarantees, but they follow the same pattern.

Distance Metric & Threshold: Separation is evaluated using the Euclidean distance metric $(\|\sigma^* - \sigma^\varepsilon\|_2)$ relative to the Pareto-dominant purestrategy equilibrium $\sigma^* = ((1,0),(1,0))$ within the coordination block B, using a fixed distance threshold of $\delta = 0.15$.

Theoretical Mapping: The "Theoretical" column values are explicitly derived from Theorem 3.1 (General Euclidean Lower Bound), using the exact experimental parameters: $b = 4, M = 1, \varepsilon = 0.1, \delta = 0.15$, and varying block payoff gaps $\varrho \in \{0.1,0.2,0.3,0.4\}$.

Table 3. Empirical vs. theoretical separation probabilities for synthetic solid games 2 × 2 coordination games (ε = 0.1, N = 1000) with varying block payoff gap 1+ϱ. Empirical entries report ±1.96·SE (95% binomial CI).

ϱ	Empirical	Theoretical	Difference
0.1	0.62 ± 0.030	0.18	0.44
0.2	0.74 ± 0.027	0.29	0.45
0.3	0.85 ± 0.022	0.38	0.47
0.4	0.93 ± 0.016	0.46	0.47

6.2. Market Entry Games

In the second experiment we perform experiments on a market entry game.

Setup. There are 3 firms and 2 choices regarding market entry in this game. Payoffs are drawn from a Gaussian distribution with a correlation structure that creates a coordination block. In this block firms receive higher payoffs when they all enter the market during favorable conditions. The difference in block payoffs is approximately $\varrho = 0.25$. All runs use random seed 42, and we report binomial standard errors with 95% confidence intervals.

Results. With $N = 500$ instances and $\varepsilon = 0.05$, the data in Table 4 shows that the probability of separation at $\delta = 0.3$ is 0.78. By comparison the theoretical bound is 0.42. The difference between the numbers is a result of the fact that the theory accounts for the most extreme scenarios. And yet both numbers are more than zero, which shows that separation happens in realistic situations. All theoretical values in Table 4 are calculated using Equation (1) (Theorem 3.1: General Euclidean Lower Bound). M=1,ϱ=0.25, C was calibrated to the simplex covering volume.

The experiments confirm these qualitative predictions of the bounds.

Both bounds and empirical values remain substantially above zero across all tested δ, confirming that separation is a robust, persistent phenomenon in market-entry games.

The theoretical lower bound remains strictly valid (empirical value ≥ theoretical bound) across every row of Table 4.

Table 4. Market entry game results. Empirical entries report ±1.96·SE (95% binomial CI; N = 500 per row).

δ	Empirical	Theoretical	Gap
0.1	0.95 ± 0.019	0.21	0.74
0.2	0.88 ± 0.028	0.35	0.53
0.3	0.78 ± 0.036	0.42	0.36
0.4	0.65 ± 0.042	0.48	0.17

In the market-entry game, the empirical separation probability at $\delta = 0.3$ is 0.78, while the theoretical lower bound is 0.42. The gap is expected: the bounds are conservative, universal lower bounds, whereas the empirical game has additional structure. The important point is that both numbers are substantially above zero, confirming that separation is a real phenomenon in plausible strategic settings.

The synthetic coordination game experiment shows that increasing $\varrho(B)$ shifts the empirical separation curve upward, exactly as the robustness penalty predicts. This validates the use of $\varrho(B)$ as a tunable robustness parameter: raising the payoff advantage of in-block strategies makes the exact equilibrium more isolated from approximate equilibria.

7. Conclusions and Future Research

In this paper we provide mathematical bounds for the probability that exact Nash equilibria and ε-approximate Nash equilibria in solid games are separated by a specific distance—these bounds apply to three types of distance measurements and four categories of games. They are applicable to games with multiple blocks or settings that depend on an index. The block payoff gap $\varrho(B)$ is the primary variable. It is present in every bound and it determines the relationship between the distance of equilibria and the strength of the block. It is a measurement on a scale rather than a simple label of solid or not solid.

The technical points are as follows—every bound is composed of two parts that relate to the quality of the approximation, the strength of the block and a third in some cases on the geometry of the distribution—those parts represent specific mechanisms. In games that are zero-sum, coordination, potential or Gaussian random games, the bounds are more precise because of the specific rules of those games. When the size of the block is larger, the probability of separation becomes 1 at a rate of $O\left(e^{-c\bar{b}}\right)$. Separation is a general feature of large solid games.

There are multiple topics for future work:

A possible topic is Algorithmic Computation of Block Payoff Gaps ϱ(B).

Efficient Estimation: Developing tractable convex-relaxation or gradient-based algorithms to compute or approximate ϱ(B) efficiently in general non-convex normal- form games without relying solely on brute-force grid searches. Another direction is Extension to Extensive-Form and Dynamic Games

Extensive-Form Structures: Expanding the solid-game block and index framework to extensive-form games with imperfect information, incorporating behavioral strategies, information sets, and sequential rationality constraints. Also we may investigate Learning Dynamics and Convergence Rates.

Dynamic Convergence: Establishing formal mathematical links between the probabilistic separation bounds and the rate at which adaptive learning algorithms (such as fictitious play, no-regret learning, and multi-agent reinforcement learning) converge to or escape approximate equilibrium components. Also extend results to Distributional Generalizations Beyond Sub-Gaussian/Gaussian Payoffs.

Heavy-Tailed and Discrete Distributions: Refining the concentration geometry to handle heavy-tailed, discrete, or integer-valued payoff distributions, substituting the polynomial volume term C/δ^b with tighter exponential concentration profiles where applicable.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

No datas were used in this work.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

No AI tools were utilized for this paper.

References

1. Nash, J.F. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* **1950**, *36*, 48–49. https://doi.org/10.1073/pnas.36.1.48.
2. Kohlberg, E.; Mertens, J.F. On the Strategic Stability of Equilibria. *Econometrica* **1986**, *54*, 1003. https://doi.org/10.2307/1912320.
3. Fudenberg, D.; Tirole, J. *Game Theory*; MIT Press: Cambridge, MA, USA, 1991.
4. Brown, G.W. Iterative solution of games by fictitious play. In *Activity Analysis of Production and Allocation*; Wiley: New York, NY, USA, 1951; pp. 374–376.
5. Robinson, J. An Iterative Method of Solving a Game. *Ann. Math.* **1951**, *54*, 296. https://doi.org/10.2307/1969530.
6. Cesa-Bianchi, N.; Lugosi, G. *Prediction, Learning, and Games*; Cambridge University Press: Cambridge, UK, 2006. https://doi.org/10.1017/cbo9780511546921.
7. Bernheim, B.D. Rationalizable Strategic Behavior. *Econometrica* **1984**, *52*, 1007. https://doi.org/10.2307/1911196.
8. Pearce, D.G. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* **1984**, *52*, 1029–1050. https://doi.org/10.2307/1911197.
9. Leyton-Brown, K.; Shoham, Y. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*; Morgan & Claypool Publishers: San Rafael, CA, USA, 2008.
10. Freund, Y.; Schapire, R.E. Adaptive Game Playing Using Multiplicative Weights. *Games Econ. Behav.* **1999**, *29*, 79–103. https://doi.org/10.1006/game.1999.0738.
11. Iqbal, M.; Raheel, M.; Biswas, R. Hybrid Artificial Neural Network Models for Predicting Flexural Strength of FRP-Reinforced Concrete Beams. *Bull. Comput. Intell.* **2026**, *2*, 196–212. https://doi.org/10.53941/bci.2026.100011.
12. Asteris, P.G. Computational Intelligence: From Nature and Aristotle to Meta-Heuristic Algorithms. *Bull. Comput. Intell.* **2025**, *1*, 1–2. https://doi.org/10.53941/bci.2025.100001.
13. Fountas, N.A.; Vaxevanidis, N.M. Prediction and Surface Roughness Optimization in Die-Sinking EDM of AA 5083 Using RSM and Swarm Intelligence Algorithms. *Bull. Comput. Intell.* **2026**, *2*, 1–12. https://doi.org/10.53941/bci.2026.100001.
14. Hazan, E.; Krauthgamer, R. How Hard Is It to Approximate the Best Nash Equilibrium? *SIAM J. Comput.* **2011**, *40*, 79–91. https://doi.org/10.1137/090766991.
15. Daskalakis, C.; Goldberg, P.W.; Papadimitriou, C.H. The Complexity of Computing a Nash Equilibrium. *SIAM J. Comput.* **2009**, *39*, 195–259. https://doi.org/10.1137/070699652.
16. Chen, X.; Deng, X. Settling the Complexity of Two-Player Nash Equilibrium. In Proceedings of the 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), Berkeley, CA, USA, 21–24 October 2006; pp. 261–272. https://doi.org/10.1109/focs.2006.69.
17. Nisan, N.; Roughgarden, T.; Tardos, E.; et al. *Algorithmic Game Theory*; Cambridge University Press: Cambridge, UK, 2007.
18. Rubinstein, A. Inapproximability of Nash Equilibrium. In Proceedings of the STOC'15: Symposium on Theory of Computing, Portland, OR, USA, 15–17 June 2015; pp. 409–418. https://doi.org/10.1145/2746539.2746578.
19. Rubinstein, A. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. In Proceedings of the 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), New Brunswick, NJ, USA, 9–11 October 2016; pp. 258–265. https://doi.org/10.1109/focs.2016.35.
20. Lipton, R.J.; Markakis, E.; Mehta, A. Playing large games using simple strategies. In Proceedings of the EC03: The Fourth ACM Conference on Electronic Commerce, San Diego, CA, USA, 9–12 June 2003; pp. 36–41. https://doi.org/10.1145/779928.779933.
21. Neyman, A. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Econ. Lett.* **1985**, *19*, 227–229. https://doi.org/10.1016/0165-1765(85)90026-6.

22. Govindan, S.; Wilson, R. A global Newton method to compute Nash equilibria. *J. Econ. Theory* **2003**, *110*, 65–86. https://doi.org/10.1016/s0022-0531(03)00005-x.
23. McKelvey, R.D.; McLennan, A.M.; Turocy, T.L. *Gambit*, Version 0.2006.01.20; Gambit: Software Tools for Game Theory; Gambit, Inc.: London, UK.
24. Roughgarden, T.; Syrgkanis, V.; Tardos, E. The Price of Anarchy in Auctions. *J. Artif. Intell. Res.* **2017**, *59*, 59–101. https://doi.org/10.1613/jair.5272.
25. Dasgupta, P.; Maskin, E. The Existence of Equilibrium in Discontinuous Economic Games, I: Theory. *Rev. Econ. Stud.* **1986**, *53*, 1. https://doi.org/10.2307/2297588.
26. Myerson, R.B. Refinements of the Nash equilibrium concept. *Int. J. Game Theory* **1978**, *7*, 73–80. https://doi.org/10.1007/bf01753236.
27. Selten, R. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* **1975**, *4*, 25–55. https://doi.org/10.1007/bf01766400.
28. Kreps, D.M.; Wilson, R. Sequential Equilibria. *Econometrica* **1982**, *50*, 863. https://doi.org/10.2307/1912767.
29. von Stengel, B. Equilibrium Computation for Two-Player Games in Strategic and Extensive Form. In *Algorithmic Game Theory*; Cambridge University Press: Cambridge, UK, 2007; Volume 3, pp. 53–78. https://doi.org/10.1017/cbo9780511800481.005.
30. Harsanyi, J.C. Games with Incomplete Information Played by "Bayesian" Players, I–III Part I. The Basic Model. *Manag. Sci.* **1967**, *14*, 159–182. https://doi.org/10.1287/mnsc.14.3.159.
31. Daskalakis, C. On the complexity of approximating a Nash equilibrium. In Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, USA, 23–25 January 2011; Volume 9, pp. 1498–1517.
32. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; MIT Press: Cambridge, MA, USA, 2012.
33. Hart, S.; Mas-Colell, A. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica* **2000**, *68*, 1127–1150. https://doi.org/10.1111/1468-0262.00153.
34. Hofbauer, J.; Sigmund, K. *Evolutionary Games and Population Dynamics*; Cambridge University Press: Cambridge, UK, 1998.
35. Sandholm, W.H. *Population Games and Evolutionary Dynamics*; MIT Press: Cambridge, MA, USA, 2010.
36. Daskalakis, C.; Frongillo, R.; Papadimitriou, C.H.; et al. On Learning Algorithms for Nash Equilibria. In Proceedings of the 3rd International Symposium on Algorithmic Game Theory, Athens, Greece, 18–20 October 2010; pp. 114–125. https://doi.org/10.1007/978-3-642-16170-4_11.
37. Hart, S.; Mas-Colell, A. Uncoupled Dynamics Do Not Lead to Nash Equilibrium. *Am. Econ. Rev.* **2003**, *93*, 1830–1836. https://doi.org/10.1257/000282803322655581.
38. Ritzberger, K.; Weibull, J.; Wikman, J. Solid outcomes in finite games. *J. Econ. Theory* **2025**, *224*, 105876. https://doi.org/10.1016/j.jet.2025.105977.
39. Basu, K.; Weibull, J. Strategy subsets closed under rational behavior. *Econ. Lett.* **1991**, *36*, 141–146. https://doi.org/10.1016/0165-1765(91)90179-o.
40. Blume, L.E.; Zame, W.R. The Algebraic Geometry of Perfect and Sequential Equilibrium. *Econometrica* **1994**, *62*, 783. https://doi.org/10.2307/2951732.
41. Daskalakis, C.; Papadimitriou, C.H. Approximate Nash equilibria in anonymous games. *J. Econ. Theory* **2015**, *156*, 207–245. https://doi.org/10.1016/j.jet.2014.02.002.
42. Monderer, D.; Shapley, L.S. Potential Games. *Games Econ. Behav.* **1996**, *14*, 124–143. https://doi.org/10.1006/game.1996.0044.
43. Rosenthal, R.W. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **1973**, *2*, 65–67. https://doi.org/10.1007/bf01737559.