

The Latent of a Game: Dimension-Free Representations for N-Player Strategic Interactions

Why Most Games Are Simpler Than They Look

Grade decomposition of payoff functions yields exponential decay of interaction orders for smooth games, with ρ as a new complexity measure.

Tamás Nagy, Ph.D.

tamas.nagy@thel latent.space

Draft • April 2026

Executive Summary (Non-Technical)

Game theory is the mathematics of strategic interaction. Nash proved in 1950 that every finite game has an equilibrium — a set of strategies where no player can improve by changing their own choice alone. He won the Nobel Prize for this. But Nash’s theorem only proves equilibria *exist*. It says nothing about how to *find* them.

Computing Nash equilibria is essentially impossible for large games. A game with 100 players, each choosing from 10 strategies, has a payoff table with 10^{100} entries — more than the atoms in the universe. Even approximate computation has been proved intractable in general (PPAD-complete). Every existing method hits this exponential wall as the number of players grows. This is the curse of dimensionality applied to strategic interaction.

This paper shows that **most games encountered in economics, finance, and engineering are dramatically simpler than their raw size suggests.** The reason is smoothness: when payoff functions are smooth — which virtually all economic models assume — the interactions between players decompose into *grades*. Grade 1 is individual behavior. Grade 2 is pairwise interaction. Grade 3 is three-player coalition effects. And so on. The key result: **for smooth games, higher-order interactions decay exponentially.** A 100-player oligopoly is dominated by pairwise effects. The three-way, four-way, and higher interactions are exponentially small.

The mathematical object that captures this structure is the **Latent of the game** — a finite, basis-free description of ALL strategic interactions, organized by interaction order. The Latent Theorem for games says the effective size of this description depends only on the game’s *regularity* (measured by a single number ρ , the Latent Number) and the desired accuracy — **not on the number of players**. One hundred players or one million: same Latent size, if the interaction structure is smooth.

This has immediate consequences. Nash equilibria of the low-grade approximation (typically grade 2, pairwise interactions) can be computed in polynomial time and approximate the true equilibria to accuracy controlled by ρ . The number ρ itself becomes a **new complexity measure for games**: $\rho \gg 1$ means the game is easy; $\rho = 1$ is a phase transition; $\rho < 1$ means the game is genuinely intractable. This is a finer classification than PPAD-completeness, which is a worst-case measure that cannot distinguish between easy and hard instances.

The framework also reveals a deep connection to cooperative game theory: the **Shapley value** — the fair division of collective payoff — decomposes naturally by interaction grade, and the exponential decay theorem implies that Shapley value computation becomes tractable for smooth games.

This paper does not claim to solve PPAD-complete problems in general. The result applies to games with smooth payoff functions, which covers most economic models (Cournot competition, congestion games, auction mechanisms with smooth valuations) but excludes games with discontinuous payoffs or combinatorial structure.

Abstract

We define the **Latent of an N-player game** as the element of a graded Hilbert tensor algebra that encodes the complete interaction structure of the game’s payoff functions, organized by interaction order. The grade- r component captures irreducible r -player interactions. We prove the **Interaction Decay Theorem**: if the payoff functions extend holomorphically to a Bernstein poly-ellipse with parameter $\rho > 1$, then the grade- r interaction component has norm bounded by $C\rho^{-r}$. The effective interaction order is $R^* = \lceil \log(1/\varepsilon)/\log \rho \rceil$, independent of the number of players N . We prove the **Truncated Equilibrium Theorem**: Nash equilibria of the grade- R^* truncated game approximate equilibria of the full game within ε in the strategy metric, with convergence rate $O(\rho^{-R^*})$. We introduce ρ as a **game complexity measure** that classifies individual game instances on a continuous spectrum from trivially compressible ($\rho \gg 1$) to fundamentally intractable ($\rho \leq 1$), providing finer resolution than worst-case PPAD-completeness. We decompose the **Shapley value** by interaction grade and show that the exponential decay theorem makes Shapley computation tractable for smooth cooperative games. We characterize **mean-field games** as the grade-1 truncation of the Latent, with explicit error bounds for the mean-field approximation. We apply the framework to Cournot oligopolies (exact grade- d representation for degree- d inverse demand), congestion games (grade-2 exactness for affine latency), and smooth auction mechanisms. All results specialize to the core Latent Theorem (Nagy, 2026) when the game has a single player.

1. Introduction

1.1 The Computational Wall in Game Theory

John Nash’s 1950 existence theorem is one of the cornerstones of mathematical economics. Every finite game has at least one Nash equilibrium: a strategy profile where no player benefits from unilateral deviation. The theorem is non-constructive — it follows from Brouwer’s fixed point theorem — and the question of *computing* equilibria has driven six decades of research.

The answer is sobering. Daskalakis, Goldberg, and Papadimitriou (2009) proved that computing a Nash equilibrium in a two-player game is PPAD-complete. For N -player games, the situation is worse: the payoff tensor has $\prod_{i=1}^N |S_i|$ entries, and no polynomial-time algorithm exists unless unexpected complexity-theoretic collapses occur. The curse of dimensionality is absolute for general games.

But general games include pathological constructions designed to be hard. The games that arise

in economics — oligopolies, auctions, congestion games, trading mechanisms — are not arbitrary. Their payoff functions are typically smooth, often polynomial, sometimes analytic. The question is not whether general games are hard (they are) but whether *the games we actually care about* are hard.

1.2 The Key Observation

Consider an N -firm Cournot oligopoly where firm i chooses quantity q_i and the market price is $P(Q)$ for $Q = \sum_j q_j$. Firm i 's profit is:

$$u_i(q) = q_i \cdot P(Q) - c_i(q_i)$$

When P is linear, u_i depends on q_i and $Q = \sum_j q_j$ but NOT on any individual q_j separately. The interaction is mediated entirely through the aggregate Q . In the language of this paper: grade 1 captures the individual cost $c_i(q_i)$, and grade 2 captures all pairwise interactions via Q . There are no higher-order interactions. The game is *exactly grade 2*, regardless of whether there are 3 firms or 3 million.

This is not an accident. Most economic models have low-order interaction structure because the mechanisms that create strategic interdependence — prices, congestion, externalities, information — typically operate through aggregates or bilateral channels, not through irreducible k -body effects for large k .

The Latent framework makes this observation precise and general.

1.3 Contribution

This paper:

1. **Defines the Latent of a game** — a graded tensor element where grade r captures exactly the r -player interaction component of the payoff structure (Section 3).
2. **Proves the Interaction Decay Theorem** — for games with analytic payoff functions, the grade- r component decays as ρ^{-r} where ρ is the Bernstein analyticity parameter (Section 4).
3. **Proves the Truncated Equilibrium Theorem** — Nash equilibria of the grade- R^* game approximate true equilibria within ε , where $R^* = O(\log(1/\varepsilon)/\log \rho)$ is independent of N (Section 5).
4. **Introduces the Latent Number ρ as a game complexity measure** — providing instance-specific complexity classification on a continuous spectrum (Section 6).
5. **Decomposes the Shapley value by interaction grade** and proves tractability for smooth cooperative games (Section 7).
6. **Characterizes mean-field games as grade-1 Latent truncation** with explicit approximation bounds (Section 8).

1.4 Related Work

ANOVA decomposition. The functional ANOVA (Hoeffding, 1948; Efron & Stein, 1981) decomposes a multivariate function into main effects and interactions. The exponential decay of ANOVA

components for analytic functions is known in the approximation theory literature (Griebel & Kuo, 2009; Potts & Schmischke, 2020). Our contribution is the game-theoretic interpretation and the equilibrium approximation theorem.

Mean-field games. Lasry & Lions (2006–2007) and Huang, Malhamé & Caines (2006) introduced the mean-field game framework for games with many symmetric players. The Latent framework subsumes mean-field as the grade-1 special case and provides explicit error bounds for the approximation.

Tensor methods in games. Tensor decomposition has been applied to symmetric games (Papadimitriou & Roughgarden, 2008) and to neural network-based equilibrium solvers (IJCAI, 2024). These operate on the payoff tensor directly. The Latent works on the payoff *function*, which is the continuous object the tensor discretizes.

Spectral game theory. Candogan, Menache, Ozdaglar & Parrilo (2011) decomposed games into potential, harmonic, and non-strategic components using graph spectral methods. This is a structural decomposition of the *game*, not of the *payoff function*. The Latent decomposition is orthogonal and complementary.

Algorithmic game theory. The PPAD-completeness results (Daskalakis et al., 2009; Chen, Deng & Teng, 2009) establish worst-case hardness. Our ρ -based classification provides instance-specific tractability results that live below the worst-case barrier.

2. Preliminaries

2.1 Strategic Games

An N -player strategic game is a triple $G = ([N], \{S_i\}, \{u_i\})$ where:

- $[N] = \{1, \dots, N\}$ is the player set.
- $S_i \subseteq \mathbb{R}^{d_i}$ is the strategy space of player i , assumed compact.
- $u_i : S \rightarrow \mathbb{R}$ is the payoff function of player i , where $S = \prod_{i=1}^N S_i$ is the joint strategy space.

The **ambient dimension** of the game is $D = \sum_{i=1}^N d_i$. A **Nash equilibrium** is a profile $s^* \in S$ such that $u_i(s^*) \geq u_i(s'_i, s^*_{-i})$ for all $i \in [N]$ and all $s'_i \in S_i$.

2.2 The Graded Hilbert Tensor Algebra

Let $H_i = L_0^2(S_i, \mu_i)$ be the space of square-integrable, zero-mean functions on S_i with respect to a reference measure μ_i (e.g., uniform on S_i). The **graded Hilbert tensor algebra** generated by H_1, \dots, H_N is:

$$\mathcal{L} = \bigoplus_{r=0}^N \mathcal{L}_r, \quad \mathcal{L}_r = \bigoplus_{|T|=r} \bigotimes_{i \in T} H_i$$

where $T \subseteq [N]$ ranges over all subsets of size r . The grade- r component \mathcal{L}_r collects all r -player interaction spaces.

For $r = 0$: $\mathcal{L}_0 = \mathbb{R}$ (constants). For $r = 1$: $\mathcal{L}_1 = H_1 \oplus \dots \oplus H_N$ (individual effects). For $r = 2$: $\mathcal{L}_2 = \bigoplus_{i < j} H_i \otimes H_j$ (pairwise interactions).

2.3 The Functional ANOVA Decomposition

Every $f \in L^2(S, \mu)$ has a unique **ANOVA decomposition**:

$$f(s) = f_\emptyset + \sum_{i=1}^N f_{\{i\}}(s_i) + \sum_{i < j} f_{\{i,j\}}(s_i, s_j) + \dots + f_{[N]}(s)$$

where each component f_T depends only on $(s_i)_{i \in T}$ and satisfies the orthogonality condition:

$$\int_{S_k} f_T(s_T) d\mu_k(s_k) = 0 \quad \text{for all } k \in T$$

This decomposition is an isometry: $\|f\|^2 = \sum_{T \subseteq [N]} \|f_T\|^2$.

2.4 Bernstein Analyticity

A function $f : [-1, 1]^D \rightarrow \mathbb{R}$ is ρ -analytic if it extends to a holomorphic function on the **Bernstein poly-ellipse** $E_\rho = \prod_{j=1}^D E_\rho^{(j)}$, where $E_\rho^{(j)} = \{z \in \mathbb{C} : |z - 1| + |z + 1| \leq \rho + \rho^{-1}\}$ is the Bernstein ellipse of parameter ρ in coordinate j .

The **Latent Number** of f is:

$$\rho(f) = \sup\{\rho > 0 : f \text{ extends holomorphically to } E_\rho\}$$

This is the distance to the nearest singularity in the complexified domain. It measures how “smooth” the function is: $\rho = \infty$ for entire functions (polynomials, exponentials), $\rho > 1$ for analytic functions, $\rho = 1$ for functions that are smooth but not analytic at the boundary.

3. The Latent of a Game

3.1 Definition

Definition 1 (Latent of a Game). Let $G = ([N], \{S_i\}, \{u_i\})$ be an N -player strategic game with square-integrable payoff functions. The **Latent of G** is the element $\Lambda(G) \in \mathcal{L}^N$ defined by:

$$\Lambda(G) = (\Lambda(u_1), \dots, \Lambda(u_N))$$

where $\Lambda(u_i) \in \mathcal{L}$ is the ANOVA decomposition of u_i viewed as an element of the graded tensor algebra:

$$\Lambda(u_i) = \sum_{T \subseteq [N]} (u_i)_T, \quad (u_i)_T \in \bigotimes_{j \in T} H_j$$

Definition 2 (Interaction Grade). The **grade- r Latent** of G is:

$$\Lambda_r(G) = (\Lambda_r(u_1), \dots, \Lambda_r(u_N)), \quad \Lambda_r(u_i) = \sum_{|T|=r} (u_i)_T$$

This captures the totality of r -player interactions in the game.

3.2 Game-Theoretic Interpretation

The grade decomposition has a natural strategic reading:

Grade	Component	Strategic Meaning
$r = 0$	$\Lambda_0(u_i) = \mathbb{E}[u_i]$	Average payoff (no strategy matters)
$r = 1$	$\Lambda_1(u_i) = \sum_j (u_i)_{\{j\}}(s_j)$	Individual effects: how each player's strategy affects i 's payoff <i>in isolation</i>
$r = 2$	$\Lambda_2(u_i) = \sum_{j < k} (u_i)_{\{j,k\}}(s_j, s_k)$	Pairwise interaction: how player <i>pairs</i> jointly affect i 's payoff, beyond individual effects
r	$\Lambda_r(u_i)$	Irreducible r -player coalitional effects

Grade 1 is pure individual optimization — no strategic interdependence. Grade 2 introduces the essential game-theoretic element: one player's strategy affects another's payoff. Grade $r \geq 3$ captures irreducible multi-body effects that no pairwise model can represent.

3.3 The Truncated Game

Definition 3 (Grade- R Truncated Game). The **grade- R truncated game** $G^{(R)}$ has the same player set and strategy spaces as G , with payoff functions:

$$u_i^{(R)}(s) = \sum_{r=0}^R \Lambda_r(u_i)(s)$$

The truncated game keeps only interactions of order $\leq R$, discarding all coalition effects involving more than R players simultaneously.

Key special cases:

- $G^{(0)}$: the **trivial game** — each player receives their average payoff regardless of strategy. Every profile is an equilibrium.
- $G^{(1)}$: the **independent game** — each player optimizes independently. No strategic interaction. Equilibrium is the profile of individual optima.
- $G^{(2)}$: the **pairwise game** — each player's payoff depends on pairwise interactions only. This is a graphical game on the complete graph.

4. Exponential Decay of Interaction Orders

4.1 Main Theorem

Theorem 1 (Interaction Decay). Let $G = ([N], \{S_i\}, \{u_i\})$ be an N -player game where each $S_i = [-1, 1]^{d_i}$ (after affine rescaling) and each u_i is ρ -analytic with $\rho > 1$. Then for all $i \in [N]$ and all $r \geq 0$:

$$\|\Lambda_r(u_i)\|_{L^2} \leq C_i \cdot \binom{N}{r} \cdot \rho^{-r}$$

where $C_i = \|u_i\|_{L^\infty(E_\rho)}$ is the sup-norm on the Bernstein poly-ellipse.

Moreover, the **relative grade- r energy** satisfies:

$$\frac{\|\Lambda_r(u_i)\|_{L^2}^2}{\|u_i\|_{L^2}^2} \leq \binom{N}{r} \cdot \rho^{-2r}$$

Proof sketch. Each ANOVA component $(u_i)_T$ for $|T| = r$ is the integral projection of u_i onto the r -player interaction subspace. For ρ -analytic functions, Chebyshev coefficient decay gives $|a_{\mathbf{k}}| \leq C\rho^{-|\mathbf{k}|}$ where $|\mathbf{k}| = \sum_j k_j$. The ANOVA component of order r involves only multi-indices supported on exactly r coordinate groups. The minimum total degree of such a multi-index is r (at least one nonzero index per group), giving the ρ^{-r} bound. The $\binom{N}{r}$ factor counts the number of subsets of size r . \square

4.2 The Effective Interaction Order

Corollary 1. For accuracy $\varepsilon > 0$, the payoff function u_i is approximated within ε (in L^2 norm) by its grade- R^* truncation, where:

$$R^* = \left\lceil \frac{\log(C_i/\varepsilon)}{\log \rho - \log(eN/R^*)} \right\rceil$$

For $\rho \gg N$ (strongly analytic games), this simplifies to:

$$R^* \approx \left\lceil \frac{\log(1/\varepsilon)}{\log \rho} \right\rceil$$

which is **independent of N** — the same formula as the core Latent Theorem.

For $\rho \sim N$ (moderately analytic games), R^* grows logarithmically with N . For $\rho < 1 + c/N$, the decay rate fails to overcome the combinatorial $\binom{N}{r}$ growth, and the game is not compressible. This identifies the **critical regime**: the interaction of smoothness and combinatorial complexity.

4.3 Sharpness

The bound is tight. Consider the N -player symmetric game with payoff $u_i(s) = \prod_{j=1}^N g(s_j)$ for an analytic function g with Bernstein parameter ρ . The ANOVA decomposition of a product has nonzero components at every grade, with the grade- r component decaying exactly as ρ^{-r} (from the

Chebyshev expansion of g). At the phase transition $\rho = 1$, the components do not decay, and the full product structure is needed.

5. Equilibrium in the Latent Algebra

5.1 Nash Equilibrium of the Truncated Game

Theorem 2 (Truncated Equilibrium Approximation). Let G be a game with ρ -analytic payoff functions ($\rho > 1$) and compact, convex strategy spaces. Let $s^{*(R)}$ be a Nash equilibrium of the grade- R truncated game $G^{(R)}$, and let s^* be a Nash equilibrium of the full game G . Assume the payoff functions satisfy a **stability condition**: the Jacobian of the best-response map has spectral radius < 1 at equilibrium. Then:

$$\|s^{*(R)} - s^*\| \leq K \cdot \rho^{-R}$$

where K depends on the Lipschitz constant of the best-response map and the stability margin but not on N .

Proof sketch. The best-response map $\text{BR}_i(s_{-i}) = \arg \max_{s_i} u_i(s_i, s_{-i})$ depends on u_i through its first-order condition $\nabla_{s_i} u_i = 0$. The gradient of the grade- R truncated payoff differs from the gradient of the full payoff by the tail:

$$\|\nabla_{s_i} u_i - \nabla_{s_i} u_i^{(R)}\| \leq \sum_{r>R} \|\nabla_{s_i} \Lambda_r(u_i)\| \leq C' \rho^{-R}$$

By the implicit function theorem (applied to the equilibrium fixed-point equation under the stability condition), this perturbation in the first-order condition propagates to a perturbation of the same order in the equilibrium. \square

5.2 Computational Consequence

Corollary 2 (Polynomial-Time Equilibrium Approximation). For $R = O(1)$ (constant effective interaction order), the grade- R truncated game $G^{(R)}$ is a **polymatrix game** (when $R = 2$) or a **low-degree polynomial game** (when $R > 2$). For polymatrix games, Nash equilibria can be computed in polynomial time via Lemke-Howson-type algorithms or linear complementarity methods. Therefore:

For any $\varepsilon > 0$ and any game with $\rho > 1$, an ε -approximate Nash equilibrium can be computed in time polynomial in N and ρ^{R^} , where R^* depends only on ε and ρ .*

This does not contradict PPAD-completeness, which is a worst-case result over ALL games (including games with $\rho \leq 1$). Our result applies to the subclass of smooth games, which is the economically relevant class.

5.3 Grade-2 Games Are Tractable

When ρ is large enough that $R^* = 2$, the truncated game is a polymatrix game (also called a separable network game): each player’s payoff is a sum of pairwise interactions. Computing Nash equilibria of polymatrix games has been extensively studied:

- For zero-sum polymatrix games: solvable in polynomial time via LP (Cai & Daskalakis, 2011).
- For general polymatrix games: approximate equilibria via PTAS methods.
- For potential polymatrix games: global optimization of the potential function.

The Latent framework identifies WHEN the polymatrix approximation is valid (when ρ is large enough that $R^* = 2$) and provides explicit error bounds.

6. The Latent Number as a Game Complexity Measure

6.1 Definition for Games

Definition 4 (Latent Number of a Game). The **Latent Number** $\rho(G)$ of a game G is:

$$\rho(G) = \min_{i \in [N]} \rho(u_i)$$

where $\rho(u_i)$ is the Bernstein analyticity parameter of the i -th payoff function.

The Latent Number measures the *least smooth* payoff — the bottleneck for compressibility.

6.2 Complexity Classification

The Latent Number induces a continuous complexity classification:

Regime	ρ	Structure	Example
Polynomial	∞	Payoffs are polynomial; game is exactly grade- d for degree d	Cournot with linear demand
Strongly analytic	$\gg N$	$R^* = O(1)$, independent of N	Most smooth economic models
Weakly analytic	$\sim N$	R^* grows logarithmically with N	Games near criticality
Critical	1	Phase transition; all grades contribute	Threshold mechanisms
Sub-critical	< 1	Game is not compressible	Genuinely combinatorial games

6.3 ρ vs. PPAD-Completeness

PPAD-completeness says: *there exist games where computing a Nash equilibrium is hard*. The Latent Number says: *this specific game G has complexity $\rho(G)$, and if $\rho(G) > 1$, we can compute its equilibrium efficiently*.

The relationship is:

- PPAD-completeness is a property of the *problem class* (all games).
- The Latent Number is a property of an *individual game instance*.
- PPAD-hard instances must have $\rho \leq 1$ (otherwise our polynomial algorithm would contradict the hardness result).
- The converse is open: does $\rho \leq 1$ imply hardness? We conjecture yes.

Conjecture 1 (Latent Hardness Conjecture). The constructions used in PPAD-completeness reductions (Daskalakis et al., 2009) produce games with $\rho \leq 1$.

If true, $\rho = 1$ is not just a phase transition in approximation quality but the precise boundary between tractable and intractable game instances.

6.4 The ρ Landscape of Classical Games

Game Class	ρ	Effective Grade
Cournot oligopoly (linear demand)	∞	2 (exact)
Cournot oligopoly (polynomial demand, degree d)	∞	d (exact)
Cournot oligopoly (analytic demand, singularity at Q_c)	Q_c/Q_{\max}	$O(\log(1/\varepsilon)/\log \rho)$
Congestion game (affine latency)	∞	2 (exact)
Congestion game (BPR latency, degree 4)	∞	4 (exact)
First-price sealed-bid auction (smooth distributions)	$\rho > 1$	$O(\log(1/\varepsilon))$
All-pay auction (smooth distributions)	$\rho > 1$	$O(\log(1/\varepsilon))$
Voting game (threshold rule)	1	N (full)
Matching pennies (N -player)	1	N (full)

The pattern is clear: economically natural games have $\rho \gg 1$; games constructed to be hard (matching pennies, threshold mechanisms) have $\rho = 1$.

7. Connection to Cooperative Game Theory

7.1 The Shapley Value and Interaction Grades

A **cooperative game** (N, v) has a characteristic function $v : 2^N \rightarrow \mathbb{R}$ with $v(\emptyset) = 0$. The **Shapley value** $\phi_i(v)$ gives each player their fair contribution:

$$\phi_i(v) = \sum_{T \subseteq [N] \setminus \{i\}} \frac{|T|!(N - |T| - 1)!}{N!} [v(T \cup \{i\}) - v(T)]$$

The characteristic function has its own interaction decomposition (the **Möbius inversion** or **Harsanyi dividends**):

$$v(T) = \sum_{S \subseteq T} \Delta_S(v), \quad \Delta_S(v) = \sum_{T \subseteq S} (-1)^{|S| - |T|} v(T)$$

where $\Delta_S(v)$ is the **Harsanyi dividend** of coalition S — the irreducible value created by the members of S acting together.

Theorem 3 (Shapley–Latent Decomposition). The Shapley value decomposes by interaction grade:

$$\phi_i(v) = \sum_{r=1}^N \phi_i^{(r)}(v), \quad \phi_i^{(r)}(v) = \sum_{\substack{S \ni i \\ |S|=r}} \frac{\Delta_S(v)}{|S|}$$

where $\phi_i^{(r)}$ is player i 's Shapley value contribution from grade- r interactions.

Proof. This follows from the well-known identity $\phi_i(v) = \sum_{S \ni i} \Delta_S(v)/|S|$, which is a rearrangement of Shapley's formula using Harsanyi dividends. Grouping by $|S| = r$ gives the grade decomposition. \square

7.2 Tractable Shapley via Exponential Decay

Corollary 3 (Shapley Truncation). If the cooperative game arises from a strategic game with ρ -analytic payoffs via $v(T) = \text{val}(G|_T)$ (the equilibrium value of the subgame restricted to coalition T), and if the subgame values inherit analyticity, then:

$$\left| \phi_i(v) - \sum_{r=1}^{R^*} \phi_i^{(r)}(v) \right| \leq C \rho^{-R^*}$$

The grade- R^* Shapley approximation requires computing Harsanyi dividends only for coalitions of size $\leq R^*$. There are $\sum_{r=1}^{R^*} \binom{N}{r} = O(N^{R^*})$ such coalitions, polynomial in N for fixed R^* .

Remark. Computing the exact Shapley value is #P-hard in general (Deng & Papadimitriou, 1994). The Latent framework does not overcome this for arbitrary games but provides polynomial-time approximation for smooth games — the same pattern as for Nash equilibria.

8. Mean-Field Games as Grade-1 Truncation

8.1 The Mean-Field Limit

In the mean-field game framework (Lasry & Lions, 2006–2007; Huang, Malhamé & Caines, 2006), N symmetric players interact through the empirical distribution $m_N = \frac{1}{N} \sum_{j=1}^N \delta_{s_j}$. As $N \rightarrow \infty$, the game converges to a coupled PDE system:

$$\begin{aligned} \partial_t v + H(x, \nabla v, m) &= 0 && \text{(Hamilton-Jacobi-Bellman)} \\ \partial_t m - \operatorname{div}(m \cdot D_p H(x, \nabla v, m)) &= 0 && \text{(Fokker-Planck)} \end{aligned}$$

This is a limit where individual identity vanishes and only the aggregate distribution matters.

8.2 The Latent Interpretation

In the Latent framework, the mean-field limit is precisely the **grade-1 truncation** of the symmetric game.

Theorem 4 (Mean-Field as Grade-1 Latent). For a symmetric N -player game with payoff $u_i(s) = \tilde{u}(s_i, m_N^{-i})$ where $m_N^{-i} = \frac{1}{N-1} \sum_{j \neq i} \delta_{s_j}$, the grade-1 component of u_i satisfies:

$$\Lambda_1(u_i)(s_i) = \tilde{u}(s_i, \bar{m}) + O(1/N)$$

where \bar{m} is the population mean strategy. The mean-field limit keeps only the grade-1 effect (interaction with the aggregate) and discards all finite-player corrections.

8.3 Beyond Mean-Field: Latent Corrections

The Latent framework provides explicit corrections to mean-field:

- **Grade-2 correction:** captures pairwise deviations from the mean-field approximation. These are the $O(1/N)$ finite-size effects that mean-field theory drops.
- **Grade-3 correction:** three-body effects. These decay as ρ^{-3} and are typically negligible for smooth games.

Proposition 1 (Mean-Field Error via Latent). For a symmetric game with ρ -analytic payoffs, the mean-field approximation error satisfies:

$$\|u_i - u_i^{(1)}\|_{L^2} \leq C \cdot N \cdot \rho^{-2}$$

This is a STRUCTURAL error bound (independent of N up to the ρ^{-2} factor), in contrast to the standard $O(1/\sqrt{N})$ probabilistic bound from the law of large numbers. The two bounds address different sources of error:

- The $O(1/\sqrt{N})$ bound: statistical fluctuation of the empirical measure around the mean.
- The $O(\rho^{-2})$ bound: structural interaction beyond the mean-field (higher-order coalitions).

For strongly analytic games ($\rho \gg 1$), the structural error is small even for moderate N . For weakly analytic games (ρ near 1), the structural error dominates.

9. Applications

9.1 Cournot Oligopoly

Consider N firms with quantities $q_i \in [0, \bar{q}]$ and inverse demand $P(Q)$ for $Q = \sum_j q_j$.

Linear demand: $P(Q) = a - bQ$. The payoff is $u_i = q_i(a - bQ) - c_i q_i$. Expanding:

$$u_i = (a - c_i)q_i - bq_i^2 - b \sum_{j \neq i} q_i q_j$$

The first two terms are grade 1 (depend only on q_i). The sum is grade 2 (pairwise terms $q_i q_j$). There are no higher grades. $\rho = \infty$, **effective grade = 2, exact**.

The Nash equilibrium of the pairwise game IS the Nash equilibrium of the full game. For $N = 1000$ firms, the pairwise game has $O(N^2) = O(10^6)$ interaction terms, trivially solvable.

Isoelastic demand: $P(Q) = Q^{-1/\eta}$. Now ρ is determined by the singularity at $Q = 0$. If the total quantity is bounded away from zero ($Q \geq Q_{\min} > 0$), then $\rho = Q_{\min}/\bar{q}$ is finite and > 1 for sensible parameter ranges. The game is still compressible, with $R^* = O(\log(1/\varepsilon)/\log(Q_{\min}/\bar{q}))$.

9.2 Congestion Games

In a congestion game, N players choose routes in a network. The cost of edge e depends on the number of players using it: $c_e(\ell_e)$ where $\ell_e = \sum_i \mathbb{1}[e \in r_i]$ is the load.

Affine latency: $c_e(\ell) = a_e + b_e \ell$. The cost for player i is:

$$u_i(r) = \sum_{e \in r_i} (a_e + b_e \ell_e) = \sum_{e \in r_i} a_e + b_e + b_e \sum_{j \neq i} \mathbb{1}[e \in r_j]$$

Grade 1: the free-flow cost $\sum_{e \in r_i} (a_e + b_e)$. Grade 2: the pairwise congestion $b_e \mathbb{1}[e \in r_i \cap r_j]$. No higher grades. $\rho = \infty$, **effective grade = 2, exact**. This is consistent with the known result that congestion games with affine latencies are potential games (Rosenthal, 1973).

BPR latency: $c_e(\ell) = t_e(1 + \alpha(\ell/C_e)^\beta)$ **with** $\beta = 4$ (**US Bureau of Public Roads**). Grade ≤ 4 , exact. Still polynomial for any N .

9.3 Smooth Auction Mechanisms

Consider a sealed-bid auction with N bidders. Bidder i has private value v_i drawn from a smooth distribution F_i and submits bid b_i . The payoff depends on ALL bids through the allocation and payment rules.

First-price auction: $u_i = (v_i - b_i) \cdot \mathbb{1}[b_i > \max_{j \neq i} b_j]$. The indicator function is discontinuous, so $\rho = 1$ at the decision boundary. However, the EXPECTED payoff (over opponents' strategies) $\mathbb{E}_{b_{-i}}[u_i] = (v_i - b_i) \prod_{j \neq i} F_j(b_i)$ is smooth if the F_j are smooth, giving $\rho > 1$ for the Bayesian game.

Revenue-optimal mechanism (Myerson, 1981): The optimal mechanism for a single item with N bidders allocates to the highest “virtual value” $\psi_i(v_i) = v_i - (1 - F_i(v_i))/f_i(v_i)$. If F_i are analytic, ψ_i are analytic, and the allocation function is a composition of analytic functions away from ties. $\rho > 1$ **generically**; $\rho = 1$ **at ties**.

For multi-item auctions (the frontier of mechanism design), the allocation function $x : V_1 \times \dots \times V_N \rightarrow [0, 1]^M$ is high-dimensional, and current methods cannot handle $N, M > 3$. The Latent framework compresses x by interaction grade, making tractable the regime where valuations are smooth and complementarities are bounded.

10. The Phase Transition: When Games Become Hard

10.1 The $\rho = 1$ Boundary

At $\rho = 1$, the exponential decay ρ^{-r} becomes $1^{-r} = 1$ for all r . Every interaction grade contributes equally. The game cannot be compressed.

What makes $\rho = 1$?

1. **Discontinuous payoffs.** A binary threshold ($u_i = \mathbb{1}[\text{condition}]$) has a singularity on the decision boundary. This is not analytic, so $\rho = 1$.
2. **Combinatorial structure.** Games where the payoff depends on the *identity* of which specific players choose which strategies (rather than on aggregates) often have $\rho = 1$. Example: the coloring game — player i receives payoff 1 if their color differs from all neighbors, 0 otherwise.
3. **Pathological constructions.** The reductions in PPAD-completeness proofs construct games with payoff functions that encode Boolean circuits. These have $\rho = 1$ by construction.

10.2 Characterization of the Transition

Theorem 5 (Phase Transition). For an N -player game with ρ -analytic payoffs:

- If $\rho > eN$: the game is **supercritical** — the truncation error decreases geometrically with grade, and $R^* = O(1)$.
- If $1 < \rho < eN$: the game is **subcritical** — the truncation error initially increases (due to the $\binom{N}{r}$ combinatorial factor) before the exponential decay takes over. $R^* = O(\log N)$.
- If $\rho = 1$: **critical** — no finite truncation suffices. The game is incompressible.

The critical exponent is $\rho_c = 1$, but the effective transition between “easy” and “hard” occurs at $\rho \sim N$ due to the interplay between smoothness and combinatorial structure.

11. Discussion and Open Problems

11.1 What This Framework Provides

The Latent of a game is not a computational algorithm but a **structural theorem** about strategic interaction. It says:

1. Every game has a natural hierarchy of interaction orders.
2. For smooth games, this hierarchy decays exponentially.
3. The effective depth of this hierarchy (the Latent Number ρ) is a new game-theoretic primitive.
4. Existing concepts (mean-field, polymatrix, Shapley) are special cases of the grade truncation.

11.2 Open Problems

Problem 1 (Latent Hardness Conjecture). Prove that PPAD-hard games have $\rho \leq 1$. This would establish $\rho = 1$ as the exact phase boundary for computational tractability.

Problem 2 (Dynamic Games). Extend the framework to extensive-form and repeated games. The grade decomposition applies at each stage, but the temporal coupling creates new technical challenges. Connection to differential games (Hamilton-Jacobi-Isaacs equations) is developed in the companion paper (Nagy, 2026).

Problem 3 (Bayesian Games). For games of incomplete information, the payoff function also depends on types (private information). The Latent decomposition extends naturally to the type-strategy product space, with type grades and strategy grades. The interaction between type complexity and strategic complexity is unexplored.

Problem 4 (Quantum Games). In quantum game theory, strategies are quantum operations and payoffs involve entanglement. The Hilbert space structure of the Latent algebra maps naturally onto the quantum state space. The grade decomposition corresponds to entanglement order. Characterizing ρ for quantum games is open.

Problem 5 (Algorithmic Implications). Develop practical algorithms that exploit the grade structure. Can the grade- R^* truncation be computed efficiently without first computing the full payoff tensor? For specific game classes (Cournot, congestion, auction), the answer is yes (Section 9). For general smooth games, this requires efficient ANOVA computation in high dimensions.

Problem 6 (Market Design). The Latent Number ρ depends on the mechanism designer's choices (allocation rule, payment rule). Can the designer choose a mechanism that maximizes ρ , making the resulting game as compressible as possible? This is an inverse problem: design the game to be easy, not just optimal.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Candogan, O., Menache, I., Ozdaglar, A., & Parrilo, P. A. (2011). Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3), 474–503.
- Chen, X., Deng, X., & Teng, S.-H. (2009). Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 1–57.
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 195–259.
- Deng, X., & Papadimitriou, C. H. (1994). On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2), 257–266.
- Efron, B., & Stein, C. (1981). The jackknife estimate of variance. *The Annals of Statistics*, 9(3), 586–596.

- Griebel, M., & Kuo, F. Y. (2009). Smoothing and decay estimates for functions of the form $f(x) = f(x_1 + \dots + x_s)$. *Proceedings of the International Congress of Mathematicians*.
- Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3), 293–325.
- Huang, M., Malhamé, R. P., & Caines, P. E. (2006). Large population stochastic dynamic games. *Communications in Information and Systems*, 6(3), 221–252.
- Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1), 229–260.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58–73.
- Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48–49.
- Papadimitriou, C. H., & Roughgarden, T. (2008). Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), 1–29.
- Potts, D., & Schmischke, M. (2020). Approximation of high-dimensional periodic functions with Fourier-based methods. *SIAM Journal on Numerical Analysis*, 59(5), 2393–2429.
- Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1), 65–67.
- Shapley, L. S. (1953). A value for n-person games. In *Contributions to the Theory of Games, Vol. II*, Annals of Mathematics Studies 28 (pp. 307–317). Princeton University Press.
- Cai, Y., & Daskalakis, C. (2011). On minmax theorems for multiplayer games. *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 217–234.
- Nagy, T. (2026). The Latent: Finite Sufficient Representations of Smooth Systems. Working Paper.
- Nagy, T. (2026). The Latent Algebra: A Universal Representational Language for Smooth Systems. Working Paper.