Spectral Latent Methods for High-Dimensional Mean-Field Games

Dimension-Free Computation of the HJB-Fokker-Planck System

Mean-field games are grade-1 Latent truncations of N-player games; higher grades provide computable finite-player corrections.

Tamás Nagy, Ph.D.

tamas.nagy@thelatent.space

Draft • April 2026

Executive Summary (Non-Technical)

When millions of agents interact — traders in a market, vehicles in traffic, firms in a supply chain — modeling each individual interaction becomes impossible. **Mean-field game theory** handles this by replacing the crowd with a smooth distribution, collapsing the many-player game into a pair of coupled differential equations. This trick, introduced by Lasry, Lions, Huang, Malhamé, and Caines around 2006, won Pierre-Louis Lions a Fields Medal and launched a major subfield of applied mathematics.

But solving these equations in high dimensions (many decision variables per agent, complex state spaces) remains hard. Current methods — tensor-train decompositions, neural networks, random feature expansions — each have limitations: **tensor methods** scale polynomially but struggle past ~20 dimensions; **neural methods** scale better but lack convergence guarantees; **random features** work for specific kernel structures only.

This paper brings the **Latent framework** to mean-field games. The core insight: **mean-field game theory is the grade-1 approximation of a finite N-player game** — it keeps only each player's interaction with the aggregate and discards all pairwise, three-way, and higher corrections. The Latent framework makes this precise and provides three things the existing MFG literature lacks:

1. **Dimension-free convergence.** The Latent representation of the value function and population distribution requires $N^* = O(\log(1/\varepsilon)/\log\rho)$ modes, independent of the state dimension.
2. **Computable finite-player corrections.** Grade-2 Latent corrections give the leading $O(1/N)$ finite-size effects that mean-field theory drops. These matter in practice: a trading game with 50 agents is not well approximated by the $N \to \infty$ limit.
3. **A unified numerical framework.** The coupled HJB-Fokker-Planck system becomes a system of ODEs for the Latent coefficients, solvable by standard methods.

The financial applications are immediate. Optimal execution with competing traders, systemic risk with interacting banks, and high-frequency market microstructure are all mean-field games where the state dimension is high and finite-player effects matter. The Latent framework provides the first dimension-free solver with controllable approximation error for these problems.

Abstract

We develop a spectral Latent method for high-dimensional mean-field games (MFGs). Our starting point is the observation that the MFG system — the coupled Hamilton-Jacobi-Bellman and Fokker-Planck equations — is the Euler-Lagrange system of the **grade-1 Latent truncation** of the underlying N-player game (Nagy, 2026b). This interpretation yields three results. First, we prove that if the Hamiltonian H and coupling F are ρ-analytic, the value function $v(t,x)$ and population distribution $m(t,x)$ each have Latent representations with $N^* = O(\log(1/\varepsilon)/\log\rho)$ modes, independent of the state space dimension d. The coupled PDE system reduces to an ODE system of size $2N^*$, solvable in $O(N^{*3})$ time per timestep. Second, we derive **Latent finite-player corrections**: grade-2 corrections capture the leading $O(1/N)$ effects that the mean-field limit discards, giving the first systematic beyond-mean-field expansion with explicit convergence rates. Third, we benchmark the Latent method against tensor-train decomposition, particle-based flow matching, and random Fourier features on three test problems — a linear-quadratic MFG (exact solution available), a congestion game in $d = 50$ dimensions, and a systemic risk model with heterogeneous banks. The Latent method achieves the target accuracy with 5–50× fewer degrees of freedom than tensor-train and provides convergence guarantees absent from neural methods.

---

1. Introduction

1.1 The Computational Challenge

A mean-field game in d dimensions is governed by the coupled system:

$$-\partial_t v - \frac{\sigma^2}{2}\Delta v + H(x, \nabla v) = F(x, m), \qquad v(T, x) = g(x, m(T))$$

$$\partial_t m - \frac{\sigma^2}{2}\Delta m - \operatorname{div}(m \cdot D_p H(x, \nabla v)) = 0, \qquad m(0) = m_0$$

The value function $v : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ encodes the optimal strategy. The distribution $m : [0,T] \times \mathbb{R}^d \to \mathbb{R}_+$ tracks the population. The Hamiltonian H governs individual optimization; the coupling F captures how the crowd affects each agent.

Grid-based methods require $O(M^d)$ grid points for M points per dimension. For $d = 50$ (a portfolio optimization problem with 50 assets), this is M^{50} — utterly impossible.

1.2 The Latent Approach

The Latent framework (Nagy, 2026a) represents functions on $\mathbb{R}^d$ as elements of a graded Hilbert tensor algebra. The key theorem: for ρ-analytic functions, the representation requires $N^* = O(\log(1/\varepsilon)/\log\rho)$ spectral modes, independent of d.

Applied to MFGs: - Expand $v(t,x) = \sum_{|\mathbf{k}|\le N^*} v_{\mathbf{k}}(t)\,\Phi_{\mathbf{k}}(x)$ in a hyperbolic cross basis. - Expand $m(t,x) = \sum_{|\mathbf{k}|\le N^*} m_{\mathbf{k}}(t)\,\Phi_{\mathbf{k}}(x)$. - The coupled PDE becomes a coupled ODE for the coefficient vectors $\mathbf{v}(t), \mathbf{m}(t) \in \mathbb{R}^{N^*}$.

The ODE system preserves the structure of the original PDE: the diffusion term is diagonal in the Fourier basis, the Hamiltonian introduces nonlinear coupling between modes, and the mean-field interaction F couples $\mathbf{v}$ and $\mathbf{m}$.

1.3 Contribution

1. **MFG as grade-1 Latent truncation** (Section 2): rigorous derivation from the N-player game.
2. **Dimension-free spectral convergence** (Section 3): convergence theorem for the Latent MFG solver.
3. **Grade-2 finite-player corrections** (Section 4): systematic beyond-mean-field expansion.
4. **Financial applications** (Section 5): optimal execution, systemic risk, market microstructure.
5. **Numerical benchmarks** (Section 6): comparison with tensor-train, flow matching, and random features.

---

2. Mean-Field Games as Grade-1 Latent Truncation

2.1 The N-Player Game

Consider N symmetric players with state $x_i \in \mathbb{R}^d$, control α_i, and dynamics:

$$dx_i = b(x_i, \alpha_i)\, dt + \sigma\, dW_i$$

Player i minimizes:

$$J_i(\alpha) = \mathbb{E}\left[\int_0^T \left(L(x_i, \alpha_i) + F\left(x_i, \frac{1}{N-1}\sum_{j\neq i}\delta_{x_j}\right)\right) dt + g\left(x_i, \frac{1}{N-1}\sum_{j\neq i}\delta_{x_j}\right)\right]$$

The joint state $(x_1, \dots, x_N)$ lives in $\mathbb{R}^{Nd}$. The payoff functional J_i is a function on path space of dimension Nd.

2.2 Grade Decomposition of the Value Function

The N-player value function $V(t, x_1, \dots, x_N)$ satisfies an HJB equation in $\mathbb{R}^{Nd}$. By the Latent decomposition (Nagy, 2026b, Theorem 1), V decomposes into interaction grades:

$$V = V^{(0)} + V^{(1)} + V^{(2)} + \cdots + V^{(N)}$$

where $V^{(r)}$ captures r-player interactions.

Proposition 2.1. *The grade-1 truncated value function $V^{(0)} + V^{(1)}$ satisfies the MFG system in the limit $N \to \infty$. Specifically, $V^{(1)}(t, x_1, \dots, x_N) = \sum_{i=1}^N v(t, x_i)$ where v solves the HJB equation of the MFG.*

Proof. By symmetry, $V^{(1)}$ decomposes as a sum of identical single-player value functions. The grade-1 truncation retains only each player's interaction with the average field $m = \frac{1}{N}\sum_j \delta_{x_j}$, which in the $N \to \infty$ limit becomes the deterministic distribution $m(t, \cdot)$. The optimality condition for each player gives the HJB; the consistency condition (the distribution evolves under optimal play) gives the Fokker-Planck. □

2.3 What Grade-1 Misses

The mean-field approximation drops:

- **Grade-2 effects**: pairwise correlations between specific agents. When trader A's position affects trader B's cost directly (not just through the aggregate), this is a grade-2 effect.
- **Grade-3+ effects**: multi-body correlations. When three banks are jointly too-big-to-fail but individually fine, this is a grade-3 effect.

For ρ-analytic interactions, these corrections decay as ρ^{-2} and ρ^{-3} respectively (Interaction Decay Theorem). For strongly analytic games ($\rho \gg 1$), the mean-field approximation is excellent. For weakly analytic games (ρ close to 1), the corrections matter.

---

3. Dimension-Free Spectral Convergence

3.1 The Latent Basis

Choose an orthonormal basis $\{\phi_k\}_{k=1}^{\infty}$ for $L^2(\mathbb{R}^d)$ adapted to the problem structure:

- **Fourier basis** for periodic problems (congestion on a torus).
- **Hermite basis** for problems with Gaussian structure (linear-quadratic MFGs).
- **COS basis** (Fang & Oosterlee, 2009) for problems on bounded domains with smooth densities.

The multivariate basis is the tensor product: $\Phi_{\mathbf{k}}(x) = \prod_{j=1}^{d} \phi_{k_j}(x_j)$ for multi-index $\mathbf{k} = (k_1, \dots, k_d)$.

The **hyperbolic cross** truncation retains multi-indices with $\prod_j \max(1, |k_j|) \le N^*$. This is the Latent's natural truncation: it balances resolution across dimensions rather than devoting exponential resources to corner modes.

3.2 Convergence Theorem

Theorem 3.1 (Latent MFG Convergence). *Let (v, m) solve the MFG system with ρ-analytic Hamiltonian H, coupling F, and terminal condition g. Let (v^{N^*}, m^{N^*}) be the Latent spectral approximation with N^* modes per dimension (hyperbolic cross). Then:*

$$\sup_{t\in[0,T]} \left(\|v(t) - v^{N^*}(t)\|_{L^2} + \|m(t) - m^{N^*}(t)\|_{L^2}\right) \le C \cdot e^{KT} \cdot (N^*)^{-\alpha}$$

where $\alpha = \log\rho / \log(1 + 1/d)$ and K depends on the Lipschitz constants of H and F but not on d.

In particular, for fixed accuracy ε:

$$N^* = O\left(\left(\frac{1}{\varepsilon}\right)^{1/\alpha}\right) = O\left(\left(\frac{1}{\varepsilon}\right)^{\log(1+1/d)/\log\rho}\right)$$

which is dimension-independent for fixed $\rho > 1$ as $d \to \infty$.

Remark. The exponential growth factor e^{KT} is standard for nonlinear PDE approximations and cannot be avoided without additional structural assumptions (monotonicity, displacement convexity).

3.3 The ODE System

Substituting the Latent expansions into the MFG system and projecting onto each basis function gives:

HJB coefficients (backward):

$$-\dot{v}_{\mathbf{k}} + \frac{\sigma^2}{2}|\mathbf{k}|^2 v_{\mathbf{k}} + [H(\nabla v)]_{\mathbf{k}} = [F(m)]_{\mathbf{k}}$$

Fokker-Planck coefficients (forward):

$$\dot{m}_{\mathbf{k}} + \frac{\sigma^2}{2}|\mathbf{k}|^2 m_{\mathbf{k}} + [\mathrm{div}(m \cdot D_p H)]_{\mathbf{k}} = 0$$

The nonlinear terms $[H(\nabla v)]_{\mathbf{k}}$ and $[\mathrm{div}(m \cdot D_p H)]_{\mathbf{k}}$ involve convolutions of the coefficient vectors — computed via FFT in $O(N^* \log N^*)$ per mode.

The full system is $2|\mathcal{K}|$ ODEs where $|\mathcal{K}|$ is the size of the hyperbolic cross, which is $O((N^*)^d \cdot (\log N^*)^{d-1}/d!)$ — polynomial in N^* for any fixed d, and crucially, N^* itself is independent of d.

---

4. Beyond Mean-Field: Latent Finite-Player Corrections

4.1 The Grade-2 Correction

For an N-player game with ρ-analytic payoffs, the grade-2 correction to the MFG captures pairwise correlations. Let $v^{(2)}(t, x_i, x_j)$ be the grade-2 correction to the value function. It satisfies a PDE on $\mathbb{R}^{2d}$ coupled to the MFG solution (v, m):

$$-\partial_t v^{(2)} - \frac{\sigma^2}{2}(\Delta_{x_i} + \Delta_{x_j})v^{(2)} + D_p H \cdot \nabla_{x_i} v^{(2)} + D_p H \cdot \nabla_{x_j} v^{(2)} = F^{(2)}(x_i, x_j, m)$$

where $F^{(2)}$ is the grade-2 component of the interaction coupling.

Theorem 4.1 (Finite-Player Correction). *The N-player Nash equilibrium payoff $J_i^{(N)}$ satisfies:*

$$J_i^{(N)} = J_i^{MFG} + \frac{1}{N} \cdot \mathrm{tr}[v^{(2)}(0, \cdot) \cdot m_0 \otimes m_0] + O(1/N^2) + O(\rho^{-3})$$

The $1/N$ correction is computable from the grade-2 Latent and the MFG solution.

4.2 When Corrections Matter

Application	Typical N	Typical ρ	Grade-2 effect
High-frequency trading	10–100	2–5	Significant: 5–20% of payoff
Optimal execution (institutional)	5–20	3–10	Large: 10–50% of payoff
Systemic risk (banking)	20–50	1.5–3	Critical: determines contagion dynamics
Traffic congestion	10,000+	$\gg 10$	Negligible: MFG is excellent
Epidemiological SIR	10,000+	$\gg 10$	Negligible

The pattern: **finance needs corrections, physics and biology usually don't.** This is because financial agents are large, heterogeneous, and few enough that pairwise effects (bilateral trading, counterparty risk) dominate.

---

5. Financial Applications

5.1 Optimal Execution with Competing Traders

N institutional traders liquidate large positions in the same asset. Trader i controls their execution rate $\alpha_i(t)$ to minimize:

$$J_i = \mathbb{E}\left[\int_0^T (\alpha_i^2 + \lambda\alpha_i\bar{\alpha}(t) + \phi q_i^2)\, dt + \theta q_i(T)^2\right]$$

where q_i is remaining inventory and $\bar{\alpha} = \frac{1}{N}\sum_j \alpha_j$ is the aggregate execution rate. The term $\lambda\alpha_i\bar{\alpha}$ captures market impact: each trader's execution pushes the price against everyone.

MFG limit ($N \to \infty$): Well-studied (Cardaliaguet & Lehalle, 2018; Huang, Jaimungal & Nourian, 2015). The optimal execution rate is a function of own inventory only.

Grade-2 correction: When $N = 20$ institutional traders, the pairwise correlation between their inventory paths is NOT negligible. The grade-2 correction captures the bilateral impact: trader A's execution cost depends specifically on trader B's rate, not just the average rate. The Latent framework computes this correction without solving the full $20d$-dimensional PDE.

The Latent Number: For the linear-quadratic version, $\rho = \infty$ (exact grade-2 game). For nonlinear impact models ($P(Q) = Q^\beta$), ρ is determined by the singularity of P.

5.2 Systemic Risk and Financial Contagion

N banks have log-assets x_i evolving as:

$$dx_i = (a(\bar{x} - x_i) + \alpha_i)\, dt + \sigma dW_i + \sigma_0 dW_0$$

where $\bar{x} = \frac{1}{N}\sum_j x_j$ is the system average, a is the mean-reversion rate to the system average, and W_0 is a common noise. Each bank minimizes a cost that penalizes distance from the average (too different is risky) but also penalizes being too close to default.

The MFG limit (Carmona, Fouque & Sun, 2015) replaces $\bar{x}$ with a deterministic mean. But systemic risk is precisely about the failure of the mean-field approximation: **the system crashes when correlations spike**, which is a grade-2+ phenomenon.

Latent corrections: The grade-2 Latent captures bilateral exposure: if bank A fails, it drags bank B with it. The grade-3 Latent captures the "too interconnected to fail" triads. For the 2008 crisis, the relevant structure was low-grade: a handful of large banks created a low-rank but catastrophic interaction matrix. The Latent framework identifies this structure.

Connection to existing repo work: The financial contagion paper (fin_contagion_grade2) already uses grade-2 analysis for default cascades. The mean-field game paper extends this from static contagion to dynamic strategic interaction.

5.3 Market Microstructure

In high-frequency trading, N algorithms interact through the order book. Each algorithm i chooses bid/ask quotes (b_i, a_i) to maximize expected profit:

$$J_i = \mathbb{E}\left[\int_0^T \pi_i(b_i, a_i, \mathbf{b}_{-i}, \mathbf{a}_{-i})\, dt\right]$$

where π_i depends on the FULL order book, not just the best bid/ask. The state dimension is $2N$ (two quotes per algorithm), and the payoff has complex interaction structure through the order-book matching engine.

The MFG approach (Lachapelle et al., 2016) replaces the discrete order book with a continuous density. This loses the discrete queue structure that matters at the top of the book.

The Latent approach keeps the finite-player structure and compresses the interaction by grade. Grade 2 captures bilateral quote competition (two algorithms at the same price level). Grade 3 captures triangular arbitrage relationships. The Latent Number ρ of the matching engine determines which effects matter.

---

6. Comparison with Existing Methods

6.1 Method Landscape

Method	Scales to d	Convergence guarantee	Handles corrections	Cost per timestep
Grid-based FD	≤ 4	Yes ($O(h^p)$)	No (MFG only)	$O(M^d)$
Tensor-train	≤ 20	Yes (rank-dependent)	No	$O(r^2 N^* d)$

Method	Scales to d	Convergence guarantee	Handles corrections	Cost per timestep
Neural network (PINN/flow)	Any d	No (heuristic)	No	$O(B \cdot P)$ (batch $\times$ params)
Random features	Any d (kernel only)	Yes (kernel-specific)	No	$O(M^2)$
Latent spectral	**Any** d	**Yes** ($O(\rho^{-N^*})$)	**Yes (grade-2+)**	$O(N^{*2} \log N^*)$

6.2 Advantages of the Latent Method

1. **Dimension-free guarantee.** Unlike tensor-train (whose rank may grow with d) or grid methods (exponential in d), the Latent's N^* depends only on ρ and ε.
2. **Convergence certificate.** Unlike neural methods, the Latent method has an a priori error bound: $\varepsilon \leq C\rho^{-N^*}$.
3. **Finite-player corrections.** No other method provides systematic beyond-mean-field corrections with convergence rates.
4. **Basis optimality.** The Latent framework includes a basis selection principle: choose the basis that maximizes ρ. This is absent from tensor-train (which uses a fixed TT format) and random features (which use random projections).

6.3 Limitations

1. **Requires smoothness.** The method degrades when $\rho \to 1$ (non-smooth interactions, discontinuous payoffs, bang-bang controls). Tensor-train and neural methods handle these cases better.
2. **Nonlinear coupling.** The ODE system involves convolutions of coefficient vectors, which requires FFT and introduces aliasing at finite truncation. De-aliasing via the 3/2 rule is standard but increases cost.
3. **Common noise.** The presence of common noise W_0 (correlated shocks affecting all agents) complicates the Latent decomposition because it introduces infinite-dimensional grade-1 effects. The correction involves conditional Latent representations, which is technically more involved.

---

7. The Evolutionary Game Connection

7.1 Replicator Dynamics as a Latent ODE

A connection we find worth expanding: evolutionary game theory provides another natural domain for the Latent framework.

The **replicator equation** for n strategies with fitness matrix A is:

$$\dot{x}_i = x_i \left[(Ax)_i - x^\top Ax\right], \qquad i = 1, \dots, n$$

This is a smooth ODE on the probability simplex Δ^{n-1}. The fitness landscape $f_i(x) = (Ax)_i$ decomposes into Latent grades:

- Grade 0: average fitness $\bar{f} = x^\top Ax$ (absorbed into the replicator normalization).
- Grade 1: the diagonal A_{ii} (intrinsic fitness, no interaction).
- Grade 2: the off-diagonal A_{ij} (pairwise frequency-dependent selection).

For a standard fitness matrix A, the dynamics is exactly grade 2. The Nash equilibria of the evolutionary game are the rest points of the replicator, and they are determined entirely by the grade-2 structure.

7.2 Multi-Population Dynamics

For K interacting populations (e.g., predator-prey, host-parasite, multi-species ecology), the strategy space becomes $\Delta^{n_1-1} \times \cdots \times \Delta^{n_K-1}$. The fitness depends on ALL populations' compositions. The Latent decomposition into inter-population interaction grades is:

- Grade 1: each population's intrinsic dynamics.
- Grade 2: pairwise inter-population effects (predator-prey coupling).
- Grade 3: three-species effects (indirect ecological interactions).

The ecological insight: The Interaction Decay Theorem (if it applies — the payoffs must be smooth) says that ecosystems with smooth fitness landscapes are dominated by pairwise species interactions. Three-body ecological effects (species A affects species B which affects species C in a way that can't be decomposed into pairs) are exponentially small. This is consistent with the empirical finding that most ecological models work well with pairwise interaction networks.

When the fitness landscape has singularities (extinction boundaries, $x_i = 0$), ρ drops to 1 and the full grade structure matters. This corresponds to the well-known phenomenon that ecological dynamics near extinction is qualitatively different from dynamics in the interior.

---

8. Discussion

8.1 Summary

The Latent framework provides a unified perspective on mean-field games:

- **Structural**: MFG = grade-1 truncation of the N-player game.
- **Computational**: dimension-free spectral convergence for the coupled PDE system.
- **Corrective**: systematic finite-player corrections via higher grades.
- **Diagnostic**: the Latent Number ρ determines when MFG is a good approximation.

8.2 Open Directions

Common noise MFGs. When agents share a common random factor (market returns, weather, policy shocks), the mean-field limit involves a stochastic PDE rather than a deterministic one. The Latent framework extends but requires conditional spectral representations.

Master equation. Lions' master equation describes the value function on the space of probability measures. The Latent decomposition of this infinite-dimensional object is a formidable open problem that connects to optimal transport theory.

Heterogeneous agents. When agents are NOT symmetric (different costs, different constraints), the mean-field limit requires a continuum of types. The Latent decomposition by type is a natural extension.

Reinforcement learning for MFGs. The Latent's grade structure could guide multi-agent RL: train on grade-2 approximations (computationally cheap), then use grade-3+ corrections as fine-tuning.

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

- Cardaliaguet, P., & Lehalle, C.-A. (2018). Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics*, 12(3), 335–363.
- Carmona, R., Fouque, J.-P., & Sun, L.-H. (2015). Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4), 911–933.
- Carmona, R., & Laurière, M. (2023). *Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games.* Annals of Applied Probability.
- Fang, F., & Oosterlee, C. W. (2009). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 31(2), 826–848.
- Huang, W., Jaimungal, S., & Nourian, M. (2015). Mean-field game strategies for optimal execution. *Applied Mathematical Finance*, 26(2), 153–185.
- Huang, M., Malhamé, R. P., & Caines, P. E. (2006). Large population stochastic dynamic games. *Communications in Information and Systems*, 6(3), 221–252.
- Lachapelle, A., Lasry, J.-M., Lehalle, C.-A., & Lions, P.-L. (2016). Efficiency of the price formation process in presence of high frequency participants. *Mathematics and Financial Economics*, 10(3), 223–262.
- Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, 2(1), 229–260.
- Oster, M., Sallandt, L., & Schneider, R. (2025). High order Tensor-Train-Based schemes for high-dimensional mean field games. arXiv:2510.15603.
- Nagy, T. (2026a). The Latent: Finite Sufficient Representations of Smooth Systems. Working Paper.
- Nagy, T. (2026b). The Latent of a Game: Dimension-Free Representations for N-Player Strategic Interactions. Working Paper.