The α-Potential Game Paradigm: Theory, Algorithms, and Applications

YUFEI ZHANG

Imperial College London

Advances in Financial Mathematics 2026, Paris

27-30 January 2025

Based on the following works

- Xin Guo, Xinyu Li, and Y. Zhang. "An α-potential game framework for N-player games." *SIAM J. Control Optim.*, 2025.
- Xin Guo, Xinyu Li, and Y. Zhang. "Distributed games with jumps: an α-potential game approach." Revision at *Math. Finance*, arXiv:2508.01929. 2025.
- Xuan Di, Anran Hu, Zhexin Wang, Y. Zhang. "α-potential games for decentralized control of connected and automated vehicles." arXiv:2512.05712. 2025.

N-player games exist ubiquitously in finance and engineering.

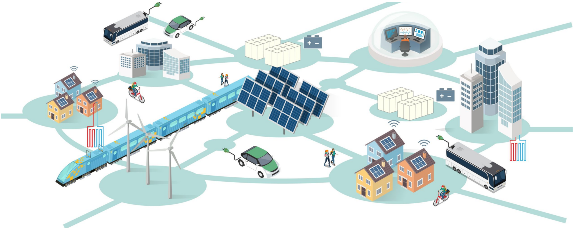

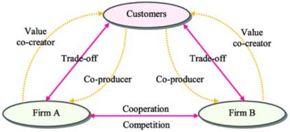

N-player games in modern life

Automated trading

N-player games exist ubiquitously in finance and engineering.

N-player games in modern life

Autonomous driving

N-player games exist ubiquitously in finance and engineering.

N-player games

- Players interact within stochastic environment and optimize their individual objectives.
- Coorperative games can be analyzed, for example, under Pareto/social optimality criteria, as high dimensional control problems
- Non-cooperative games mostly relies on the concept of Nash Equilibrium (NE): a set of strategies where given the strategies of other players, no one has incentive to deviate from the chosen strategy.

Challenges in N-player non-cooperative games

- NEs are difficult to analyze because of interdependence among players.
  - Computational complexity increasing exponentially in N.
  - Iterative algorithms may fail to converge, even in two-player games (Shapley, 1964; Mazumdar et al., 2020)

Challenges in N-player non-cooperative games

- NEs are difficult to analyze because of interdependence among players.
  - Computational complexity increasing exponentially in N.
  - Iterative algorithms may fail to converge, even in two-player games (Shapley, 1964; Mazumdar et al., 2020)
- Homogeneity and weak interaction assumptions for mean-field theory appears restrictive.
- Mean-field approximation focuses only on asymptotic behavior as $N \to \infty$.
- Our framework captures richer game structures beyond N

Static potential game

- Introduced in Monderer and Shapley (*Games and Economic Behavior*, 1996).
- $[N] = \{1, \ldots, N\}$ is a set of N players.
- Player i's action set $\mathcal{A}_i$.
- Player i's utility function $u_i : \mathcal{A}^{(N)} \to \mathbb{R}$, where $\mathcal{A}^{(N)} = \Pi_{i \in [N]} \mathcal{A}_i$ is the set of joint action profiles.

Static potential game

- Introduced in Monderer and Shapley (*Games and Economic Behavior*, 1996).
- $[N] = \{1, \ldots, N\}$ is a set of N players.
- Player i's action set $\mathcal{A}_i$.
- Player i's utility function $u_i : \mathcal{A}^{(N)} \to \mathbb{R}$, where $\mathcal{A}^{(N)} = \Pi_{i \in [N]} \mathcal{A}_i$ is the set of joint action profiles.
- There exists a potential function $\phi : \mathcal{A}^{(N)} \to \mathbb{R}$ such that

$$u_i(a_i', a_{-i}) - u_i(a_i, a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i, a_{-i}), \ \forall i \in [N].$$

  That is, any change of utility function u_i due to player i's policy change equals to the change of a potential function ϕ.

Advantage of static potential game

Popular framework for static games with heterogeneous players:

- Any minimizer of the potential function is an NE.
- Finding NE is reduced to an optimization problem
- Convergent algorithms for NE, including best response, fictitious play, and policy gradient methods.

Advantage of static potential game

Popular framework for static games with heterogeneous players:

- Any minimizer of the potential function is an NE.
- Finding NE is reduced to an optimization problem
- Convergent algorithms for NE, including best response, fictitious play, and policy gradient methods.

Question: How about dynamic games?

Earlier attempt: Markov potential game (MPG)

- Introduced by Leonardos et al. ('21) for discrete-time state dynamics with Markovian transitions and policies.
- Direct generalization of the static potential game:

$$J_i(a_i', a_{-i}) - J_i(a_i, a_{-i}) = \Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i}).$$

That is, the change of utility function J_i due to player i's policy change equals to the change of a potential function Φ.

Earlier attempt: Markov potential game (MPG)

- Introduced by Leonardos et al. ('21) for discrete-time state dynamics with Markovian transitions and policies.
- Direct generalization of the static potential game:

$$J_i(a_i', a_{-i}) - J_i(a_i, a_{-i}) = \Phi(a_i', a_{-i}) - \Phi(a_i, a_{-i}).$$

  That is, the change of utility function J_i due to player i's policy change equals to the change of a potential function Φ.
- Any minimizer of Φ is an NE.
- Various algorithms converge to an NE for an MPG.

Issues with MPGs

- Most dynamic games with Markovian transition and policies are not MPGs.
- Markov games are not MPGs even when each step is static potential game.
- Certifying a game to be an MPG is challenging.

Issues with MPGs

- Most dynamic games with Markovian transition and policies are not MPGs.
- Markov games are not MPGs even when each step is static potential game.
- Certifying a game to be an MPG is challenging.

A general form of potential game is necessary for dynamic games.

Solution: α-potential game (Guo, Li, Z. '25)

Consider a dynamic game $\mathcal{G}$:

- state space $\mathcal{S}$,
- player i's admissible strategy set $\mathcal{A}_i$,
- $J_i(\boldsymbol{a})$ is player i 's expected cost if players take the strategy profile $\boldsymbol{a} \in \mathcal{A}^{(N)}$, and the state starts with a fixed state $s_0 \in \mathcal{S}$.

For any $i \in [N]$, player i minimizes the function J_i over all admissible strategies in $\mathcal{A}_i$.

Solution: α-potential game

- A game $\mathcal{G}$ is called an α-potential game, if there exists $\alpha \geq 0$ and $\Phi : \mathcal{A}^{(N)} \to \mathbb{R}$ such that

$$|J_i((a_i', a_{-i})) - J_i((a_i, a_{-i})) - (\Phi((a_i', a_{-i})) - \Phi((a_i, a_{-i})))| \leq \alpha$$

  for any $i \in [N]$, $a_i, a_i' \in \mathcal{A}_i$, $a_{-i} \in \mathcal{A}_{-i}^{(N)}$.
- Φ is called an α-potential function for $\mathcal{G}$.

α-NE for α-potential game

Theorem

If $\boldsymbol{a}^* \in \mathcal{A}^{(N)}$ minimizes Φ, then $\boldsymbol{a}^*$ is an α-Nash equilibrium of $\mathcal{G}$:

$$J_i\left(\boldsymbol{a}^*\right) \leq J_i\left(\left(a_i, \boldsymbol{a}^*_{-i}\right)\right) + \alpha, \quad \forall a_i \in \mathcal{A}_i, i \in [N].$$

- Finding approximate NE of N-player game reduces to a optimization problem of Φ.

α-NE for α-potential game

Theorem

If $\boldsymbol{a}^* \in \mathcal{A}^{(N)}$ minimizes Φ, then $\boldsymbol{a}^*$ is an α-Nash equilibrium of $\mathcal{G}$:

$$J_i\left(\boldsymbol{a}^*\right) \leq J_i\left(\left(a_i, \boldsymbol{a}^*_{-i}\right)\right) + \alpha, \quad \forall a_i \in \mathcal{A}_i, i \in [N].$$

- Finding approximate NE of N-player game reduces to a optimization problem of Φ.

Analytical challenges:

- Construct α-potential function Φ;
- Characterize α;
- Optimize α-potential function Φ for α-NE

Fix $\boldsymbol{z} \in \mathcal{A}^{(N)}$ and define $\Phi : \mathcal{A}^{(N)} \to \mathbb{R}$ by

$$\Phi(\boldsymbol{a}) = \int_0^1 \sum_{i=1}^{N} \frac{\delta J_i}{\delta a_i} (\boldsymbol{z} + r(\boldsymbol{a} - \boldsymbol{z}); a_i - z_i) \, \mathrm{d}r,$$

where $\frac{\delta J_i}{\delta a_i}$ is the derivative of player i's objective with respect to their control. Then Φ is an α-potential function of $\mathcal{G}$ with

$$\alpha \leq 2 \sup_{i \in [N], a_i' \in \mathcal{A}_i, \boldsymbol{a}, \boldsymbol{a}'' \in \mathcal{A}^{(N)}} \sum_{j=1}^{N} \left| \frac{\delta^2 J_i}{\delta a_i \delta a_j} \left(\boldsymbol{a}; a_i', a_j'' \right) - \frac{\delta^2 V_j}{\delta a_j \delta a_i} \left(\boldsymbol{a}; a_j'', a_i' \right) \right| .$$

Stochastic differential games

For $i \in [N]$, player i considers the state dynamics

$$\mathrm{d}X_t^i = u_t^i \mathrm{d}t + \sigma_i \mathrm{d}W_t, \quad X_0^i = \xi_i,$$

and minimizes the objective function:

$$J_i(\mathbf{u}) = \mathbb{E}\left[\int_0^T f_i(t, \mathbf{X}_t^{\mathbf{u}}, \mathbf{u}_t)\,\mathrm{d}t + g_i(\mathbf{X}_T^{\mathbf{u}})\right],$$

over all adapted controls u^i, where f_i and g_i are twice differentiable functions.

Derive α-potential function

An α-potential function is given by

$$\Phi(\mathbf{u}) = \mathbb{E}\left[\int_0^T F(t, \mathbf{X}_t^{\mathbf{u}}, \mathbf{u}_t)\mathrm{d}t + G(\mathbf{X}_T^{\mathbf{u}})\right],$$

where $\mathrm{d}\mathbf{X}_t = \mathbf{u}_t \mathrm{d}t + \sigma \mathrm{d}W_t$, and

$$F(t,x,a) \coloneqq \sum_{i=1}^N \int_0^1 \begin{pmatrix} x_i \\ a_i \end{pmatrix}^\top \begin{pmatrix} \partial_{x_i} f_i \\ \partial_{a_i} f_i \end{pmatrix} (t, rx, ra)\,\mathrm{d}r,$$

$$G(x) \coloneqq \sum_{i=1}^N \int_0^1 x_i^\top \left(\partial_{x_i} g_i\right)(rx)\,\mathrm{d}r.$$

- Finding α-NE of the N-player game reduces to a finite-dimensional control problem!

Quantifying α

α captures the intensity and level of heterogeneous interactions

Theorem

$$\alpha \le C_{T,U,\xi} \sup_{i \in [N]} \sum_{j \in [N] \setminus \{i\}} \Big(|\partial^2_{x_i x_j} \Delta^f_{i,j}|_{L^\infty} + (|\partial^2_{x_i a_j} \Delta^f_{i,j}|_{L^\infty} + |\partial^2_{a_i x_j} \Delta^f_{i,j}|_{L^\infty}) + |\partial^2_{a_i a_j} \Delta^f_{i,j}|_{L^\infty} + |\partial^2_{x_i x_j} \Delta^g_{i,j}|_{L^\infty} \Big),$$

where $\Delta^f_{i,j} := f_i - f_j$, $\Delta^g_{i,j} := g_i - g_j$.

Quantifying α: games with strong interactions

α captures the intensity and level of heterogeneous interactions

Consider the following cost functional (Oelschläger, 1985):

$$J_i(\mathbf{u}) = \mathbb{E}\left[\int_0^T \left(\ell_i(u_t^i) + \sum_{j=1}^N K_N(X_t^i - X_t^j)\right) \mathrm{d}t + g_i(X_T^i)\right],$$

where

$$K_N(z) = N^{\beta-1}\rho(N^\beta |z|), \quad \beta \in [0,1],$$

with $\rho : [0,\infty) \to [0,\infty)$ satisfies $\lim_{|z|\to\infty} \rho(z) = 0$.

- Weak interaction $\beta = 0$: interaction with players of distance $O(1)$, with strength $O(1/N)$.
- Strong interaction $\beta = 1$: interaction with players of distance $O(1/N)$, with strength $O(1)$.

Theorem

$\alpha = 0$ *For all* $\beta \in [0,1]$, *i.e., the N-player game is a potential game.*

Consider the following cost functional:

$$J_i(\mathbf{u}) = \mathbb{E}\left[\int_0^T \left(\ell_i(u_t^i) + \frac{1}{N}\sum_{j=1}^N q_{ij} K(|X_t^i - X_t^j|)\right) \mathrm{d}t + g_i(X_T^i)\right],$$

where

- the function K models the interaction between players i and j,
- the weight $q_{ij} \geq 0$ reflects interaction strength.

Theorem

$$\alpha \leq C_{T,U,\xi} |\partial_{xx}^2 K|_{L^\infty} \frac{1}{N} \max_{i\in[N]} \sum_{j\neq i} |q_{ji} - q_{ij}|.$$

Quantifying α

α captures the intensity and level of heterogeneous interactions

- $\alpha = 0$ if the interactions are pairwise symmetric:

$$q_{ij} = q_{ji}, \quad \forall i \neq j.$$

This includes mean field or symmetric network interaction.

Quantifying α

α captures the intensity and level of heterogeneous interactions

- $\alpha = 0$ if the interactions are pairwise symmetric:

$$q_{ij} = q_{ji}, \quad \forall i \neq j.$$

  This includes mean field or symmetric network interaction.
- For asymmetric interactions, consider a network game specified by a graph G, where
  - the vertices are the set of players,
  - the edge indicates the connectivity between the players.

  Assume the interaction asymmetry decays in the graph distance between players, and G has a bounded degree $d_G \geq 2$.

Quantifying α

α captures the intensity and level of heterogeneous interactions

- If the degree of asymmetry decays exponentially:

$$|q_{ij} - q_{ji}| \leq w_{i,j}\rho^{c(i,j)}, \quad \text{for some } \rho \in (0,1),$$

$$\alpha_N = \begin{cases} \mathcal{O}\left(N^{\frac{\ln \rho}{\ln d_G}}\right), & \text{if } \rho \in (1/d_G, 1), \\ \mathcal{O}\left(\frac{\ln N}{N}\right), & \text{if } \rho = 1/d_G, \\ \mathcal{O}\left(N^{-1}\right), & \text{if } \rho \in (0, 1/d_G). \end{cases}$$

Quantifying α

α captures the intensity and level of heterogeneous interactions

- If the degree of asymmetry decays exponentially:

$$|q_{ij} - q_{ji}| \le w_{i,j}\rho^{c(i,j)}, \quad \text{for some } \rho \in (0,1),$$

$$\alpha_N = \begin{cases} \mathcal{O}\left(N^{\frac{\ln \rho}{\ln d_G}}\right), & \text{if } \rho \in (1/d_G, 1), \\ \mathcal{O}\left(\frac{\ln N}{N}\right), & \text{if } \rho = 1/d_G, \\ \mathcal{O}\left(N^{-1}\right), & \text{if } \rho \in (0, 1/d_G). \end{cases}$$

- If the degree of asymmetry decays polynomially:

$$|q_{ij} - q_{ji}| \le w_{i,j}c(i,j)^{-\beta}, \quad \text{for some } \beta > 0,$$

$$\alpha_N = \mathcal{O}\left(\frac{\ln \ln N}{(\ln N)^\beta}\right).$$

Quantifying α

α captures the intensity and level of heterogeneous interactions

- If the degree of asymmetry decays exponentially:

$$|q_{ij} - q_{ji}| \leq w_{i,j} \rho^{c(i,j)}, \quad \text{for some } \rho \in (0,1),$$

$$\alpha_N = \begin{cases} \mathcal{O}\left(N^{\frac{\ln \rho}{\ln d_G}}\right), & \text{if } \rho \in (1/d_G, 1), \\ \mathcal{O}\left(\frac{\ln N}{N}\right), & \text{if } \rho = 1/d_G, \\ \mathcal{O}\left(N^{-1}\right), & \text{if } \rho \in (0, 1/d_G). \end{cases}$$

- If the degree of asymmetry decays polynomially:

$$|q_{ij} - q_{ji}| \leq w_{i,j} c(i,j)^{-\beta}, \quad \text{for some } \beta > 0,$$

$$\alpha_N = \mathcal{O}\left(\frac{\ln \ln N}{(\ln N)^{\beta}}\right).$$

- $\alpha_N \to 0$ for random networks such as Erdős-Rényi network (both sparse and dense regimes), Small world network (Rokade et al. '25)

Minimize Φ over adapted controls

$$\Phi(\mathbf{u}) = \mathbb{E}\left[\int_0^T F(t, \mathbf{X}_t^{\mathbf{u}}, \mathbf{u}_t)\mathrm{d}t + G(\mathbf{X}_T^{\mathbf{u}})\right].$$

with

$$\mathrm{d}\mathbf{X}_t = \mathbf{u}_t \mathrm{d}t + \sigma \mathrm{d}W_t, \quad \mathbf{X}_0 = \boldsymbol{\xi}.$$

- DPP and verification theorem via HJB equations.
- The α-NE is given by $\mathbf{u}^* = \phi(t, \mathbf{X}_t)$.
- The α-NE can be numerically computed by optimizing over ϕ using policy gradient methods.

Numerical illustration

Aversion in weakly interacting vehicles: Weak Interaction

Aversion in strongly interacting vehicles: Strong Interaction

Aversion and obstacle avoidance in strongly interacting vehicles: Obstacle avoidance

More generally

For each control profile $\mathbf{u} = (u_i)_{i \in [N]}$, the state process $\mathbf{X}^{\mathbf{u}} = (X_i^{\mathbf{u}})_{i=1}^N$ satisfies for all $i \in [N]$,

$$\mathrm{d}X_{t,i} = b_i(t, \mathbf{X}_t, \mathbf{u}_t)\mathrm{d}t + \sigma_i(t, \mathbf{X}_t, \mathbf{u}_t)\mathrm{d}W_t, \quad X_{0,i} = x_i.$$

Player i's objective function is given by

$$J_i(\mathbf{u}) = \mathbb{E}\left[\int_0^T f_i(t, \mathbf{X}_t^{\mathbf{u}}, \mathbf{u}_t)\,\mathrm{d}t + g_i(\mathbf{X}_T^{\mathbf{u}})\right].$$

Here b_i, σ_i, f_i and g_i are twice differentiable functions.

Derive α-potential function

The α-potential function $\Phi(\mathbf{u}) = \int_0^1 \sum_{i=1}^N \frac{\delta J_i}{\delta u_i}(r\mathbf{u}; u_i)\mathrm{d}r$ can be written as

$$\Phi(\mathbf{u}) = \int_0^1 \sum_{i=1}^N \mathbb{E}\left[\int_0^T \begin{pmatrix} \mathbf{Y}_t^{r\mathbf{u},u_i} \\ u_{t,i} \end{pmatrix}^\top \begin{pmatrix} \partial_x f_i \\ \partial_{u_i} f_i \end{pmatrix} (t, \mathbf{X}_t^{r\mathbf{u}}, \mathbf{u}_t)\,\mathrm{d}t + (\mathbf{Y}_T^{r\mathbf{u},u_i})^\top (\partial_x g_i)(\mathbf{X}_T^{r\mathbf{u}})\right]\mathrm{d}r,$$

where

- $\mathbf{Y}^{r\mathbf{u},u_i}$ is the derivative of $\mathbf{X}_i^{r\mathbf{u}}$ w.r.t. player i's control u_i. It satisfies a (controlled) linear SDE.

Derive α-Nash equilibrium (Guo, Li, Z. '25)

Minimize Φ over $\mathbb{F}^W$-adapted controls:

$$\Phi(\mathbf{u}) = \int_0^1 \sum_{i=1}^N \mathbb{E}\left[\int_0^T \begin{pmatrix} \mathbf{Y}_t^{r\mathbf{u},u_i} \\ u_{i,t} \end{pmatrix}^\top \begin{pmatrix} \partial_x f_i \\ \partial_{u_i} f_i \end{pmatrix} (t, \mathbf{X}_t^{r\mathbf{u}}, \mathbf{u}_t)\, \mathrm{d}t + (\mathbf{Y}_T^{r\mathbf{u},u_i})^\top (\partial_x g_i)(\mathbf{X}_T^{r\mathbf{u}}) \right] \mathrm{d}r.$$

- Φ involves aggregating players' control over $r \in [0,1]$.
- Verification theorem via lifting into a McKean-Vlasov control problem.
- For LQ games, the HJB can be solved analytically using a system of ODEs.
- The corresponding α can be bounded using the number of player N, the state coefficients, and the derivatives of $f_i - f_j$ and $g_i - g_j$.

Summary

- Propose an α-potential game framework for N-player dynamic games.
- Reduce finding α-NE to minimizing the α-potential function.
- Construct α-potential function analytically, with α in terms of the magnitude of the asymmetry of objective functions' second-order derivatives.
- For a large class of games, α-NE can be analyzed by solving a finite-dimensional control problem.
- α is explicitly identified in terms of the number of players, and the strength and degree of heterogeneous interactions.
- In general, α-NE can be analyzed under the MKV control framework, with explicit solutions to LQ system.

A host of open problems

- Sharp estimate of α
- Properties of α-NE
- Multiple NE through different constructions of potential functions
- Efficient computational and learning algorithms
- Applications in finance problems

General α-potential function

- If J_i is continuously Fréchet differentiable in $\mathcal{A}_i$. Fix $\boldsymbol{z} \in \mathcal{A}^{(N)}$, let $p_i : [0,1] \times \mathcal{A}_i \to \mathcal{A}_i$ be a continuously differentiable reparameterization of the strategy class $\mathcal{A}_i$ such that for all $a_i \in \mathcal{A}_i$, $p_i(0, a_i) = z_i$ and $p_i(1, a_i) = a_i$. Then

$$\Phi(\boldsymbol{a}) = \int_0^1 \sum_{i=1}^N (\partial_{a_i} J_i)(p(r, \boldsymbol{a})) \partial_t p_i(r, a_i) \mathrm{d}r,$$

  where $\partial_{a_i} J_i$ is the Fréchet derivative of J_i.
- This generalizes Monderer and Shapley's criterion to dynamic games with infinite-dimensional strategy sets and arbitrary $\alpha \geq 0$.

Consider parameterized polices $(\phi_\theta)_{\theta\in\mathbb{R}^L}$:

$$\Phi(\theta) := \mathbb{E}\left[\int_0^T F(t, \mathbf{X}_t^\theta, \phi_\theta(t, \mathbf{X}_t^\theta))\mathrm{d}t + G(\mathbf{X}_T^\theta)\right],$$

with

$$\mathrm{d}\mathbf{X}_t = \phi_\theta(t, \mathbf{X}_t)\mathrm{d}t + \sigma \mathrm{d}W_t, \quad \mathbf{X}_0 = \boldsymbol{\xi}.$$

- Minimizer $\Phi(\theta)$ using policy gradient methods.