A Glimpse into Mean Field Particle Systems and Games

Lane Chun Yeung (Illinois Institute of Technology)

Research Collaboratorium

September 24, 2025

A Beautiful Mind

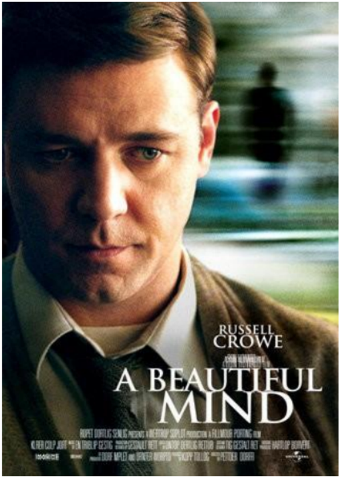

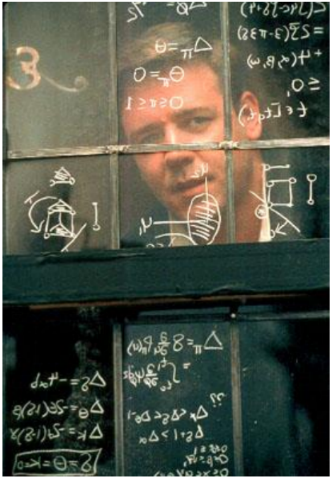

A Beautiful Mind

- John Nash (1928–2015) wrote a 28-page thesis on non-cooperative games
- Awarded the Nobel Prize in Economics in 1994

- Nash's theory in the movie (which is *not* an example of Nash equilibrium)
- `https://www.youtube.com/watch?v=bbNMTbcuitA`

A variant of Prisoner's Dilemma

Figure: Golden Balls: Split or Steal

- An amount $\$J$ is on the table.
- If both Split: fair share.
- If one Steals: stealer gets all, splitter gets nothing.
- If both Steal: both walk away with nothing.

Golden Balls: Split or Steal (A vs B)

	B: Split	**B: Steal**
A: Split	$\left(\frac{J}{2},\ \frac{J}{2}\right)$	$\left(0,\ J\right)$
A: Steal	$\left(J,\ 0\right)$	$\left(0,\ 0\right)$

- **If B chooses Split:** A gets J by **Steal** vs. $J/2$ by **Split** $\Rightarrow$ A prefers **Steal**.
- **If B chooses Steal:** A gets 0 whether **Split** or **Steal** $\Rightarrow$ indifferent.
- Thus B should steal whatever A does.
- By symmetry, A should also steal.
- Conclusion: Both of them should steal.
- The strategy profile $(\text{Steal}, \text{Steal})$ is a **Nash equilibrium**
- Note that it is not the "social optimum": $(\text{Split}, \text{Split})$ yields $(\frac{J}{2}, \frac{J}{2})$.

Nash equilibrium: Two-player, single-period game

- Two players with action sets A_1 and A_2.
- A strategy profile is $\alpha = (\alpha_1, \alpha_2) \in A_1 \times A_2$.
- Let $J_1, J_2 : A_1 \times A_2 \to \mathbb{R}$ be the players' *costs* (to be minimized).
- A strategy profile α^* is a **Nash equilibrium** iff

$$\begin{aligned} J_1(\alpha_1^*, \alpha_2^*) &\leq J_1(\alpha_1, \alpha_2^*) \quad \forall\, \alpha_1 \in A_1, \\ J_2(\alpha_1^*, \alpha_2^*) &\leq J_2(\alpha_1^*, \alpha_2) \quad \forall\, \alpha_2 \in A_2. \end{aligned}$$

- Interpretation: at α^*, neither player can improve their outcome by unilaterally deviating

n-player, single-period game

- Players $i \in \{1, \dots, n\}$ with action sets A_i
- **Notation.** For any player i, write

$$\alpha_{-i} = (\alpha_1, \dots, \alpha_{i-1}, \alpha_{i+1}, \dots, \alpha_n) \in \prod_{j \neq i} A_j$$

- Cost functions: $J_i : \prod_{j=1}^n A_j \to \mathbb{R}$, for $i = 1, \dots, n$.
- A strategy profile $\alpha^* \in \prod_{j=1}^n A_j$ is a **Nash equilibrium** if for each $i = 1, \dots, n$,

$$J_i(\alpha_i^*, \alpha_{-i}^*) \ \leq \ J_i(\alpha_i, \, \alpha_{-i}^*) \qquad \forall \, \alpha_i \in A_i.$$

- Similar definitions apply to **multi-period/repeated** games.

Continuous-time stochastic games

- **Time horizon:** $[0, T]$, where $T \in (0, \infty)$.

- **Strategy:** Each player $i \in \{1, \dots, n\}$ chooses their strategy $\alpha^i = (\alpha^i_t)_{t \in [0,T]}$ among A^i, a set of feasible actions (a.k.a. admissible controls)

- **State dynamics**: Each player's state is given by

$$\mathrm{d}X^i_t \;=\; b\big(X^i_t, \mu^n_t, \alpha^i_t\big)\,\mathrm{d}t \;+\; \mathrm{d}W^i_t,$$

  where $\mu^n_t = \frac{1}{n}\sum_{j=1}^n \delta_{X^j_t}$ is the empirical measure of all players' states.

- **Cost functions:** Each player tries to minimize

$$J_i(\alpha^1, \dots, \alpha^n) \;=\; \mathbb{E}\Bigg[\underbrace{\int_0^T f\big(X^i_t, \mu^n_t, \alpha^i_t\big)\,\mathrm{d}t}_{\text{running cost}} + \underbrace{g\big(X^i_T, \mu^n_T\big)}_{\text{terminal cost}}\Bigg].$$

- A strategy profile $\alpha^* = (\alpha^1, \dots, \alpha^n)$ is a **Nash equilibrium** if for each i,

$$J_i\big(\alpha^{i,*}, \alpha^{-i,*}\big) \;\leq\; J_i\big(\alpha^i, \alpha^{-i,*}\big), \quad \forall \alpha^i \in A^i$$

- **Issue:** Very difficult to compute Nash equilibria even if n is reasonably large!

Mean field game (MFG) paradigm

- **Idea**: A "typical" or "representative" player interacts with a **continuum** of others only through the population state distribution μ_t.
- Fix a measure flow $\mu = (\mu_t)_{t \in [0,T]}$ representing a continuum of agents' state process
- Solve the optimal control problem faced by a "typical" player

$$\begin{cases} \alpha^* = \arg\min_{\alpha} \mathbb{E}\left[\int_0^T f\big(X_t, \mu_t, \alpha_t\big)\, \mathrm{d}t \; + \; g\big(X_T, \mu_T\big) \right], \\ \mathrm{d}X_t \; = \; b\big(X_t, \mu_t, \alpha_t\big)\, \mathrm{d}t \; + \; \mathrm{d}W_t. \end{cases}$$

- **Fixed point problem:** find μ such that $\mathrm{Law}(X_t^{\alpha^*}) = \mu_t$ for all $t \in [0,T]$.
- Any pair (μ, α^*) satisfying this is an **MFG equilibrium**.

Nash's existence proof of equilibrium in n-person games

- ◇ Uses Kakutani's fixed point theorem.

EQUILIBRIUM POINTS IN N-PERSON GAMES

By John F. Nash, Jr.*

Princeton University

Communicated by S. Lefschetz, November 16, 1949

One may define a concept of an n-person game in which each player ha
a finite set of pure strategies and in which a definite set of payments to th
n players corresponds to each n-tuple of pure strategies, one strateg
being taken for each player. For mixed strategies, which are probabilit

distributions over the pure strategies, the pay-off functions are the expectations of the players, thus becoming polylinear forms in the probabilities with which the various players play their various pure strategies.

Any n-tuple of strategies, one for each player, may be regarded as a point in the product space obtained by multiplying the n strategy spaces of the players. One such n-tuple counters another if the strategy of each player in the countering n-tuple yields the highest obtainable expectation for its player against the $n - 1$ strategies of the other players in the countered n-tuple. A self-countering n-tuple is called an equilibrium point.

The correspondence of each n-tuple with its set of countering n-tuples gives a one-to-many mapping of the product space into itself. From the definition of countering we see that the set of countering points of a point is convex. By using the continuity of the pay-off functions we see that the graph of the mapping is closed. The closedness is equivalent to saying: if $P_1, P_2, \ldots$ and $Q_1, Q_2, \ldots, Q_n, \ldots$ are sequences of points in the product space where $Q_n \to Q$, $P_n \to P$ and Q_n counters P_n then Q counters P.

Since the graph is closed and since the image of each point under the mapping is convex, we infer from Kakutani's theorem[1] that the mapping has a fixed point (i.e., point contained in its image). Hence there is an equilibrium point.

In the two-person zero-sum case the "main theorem"[2] and the existence of an equilibrium point are equivalent. In this case any two equilibrium points lead to the same expectations for the players, but this need not occur in general.

* The author is indebted to Dr. David Gale for suggesting the use of Kakutani's theorem to simplify the proof and to the A. E. C. for financial support.

[1] Kakutani, S., *Duke Math. J.*, 8, 457–459 (1941).

[2] Von Neumann, J., and Morgenstern, O., *The Theory of Games and Economic Behaviour*, Chap. 3, Princeton University Press, Princeton, 1947.

The analytic PDE approach to mean field games

- For a fixed measure flow $\mu = (\mu_t)$, the value function

$$V^\mu(t,x) = \inf_{\alpha \in \mathcal{A}} \mathbb{E}\left[\int_t^T f\big(X_s, \mu_s, \alpha_s\big)\,\mathrm{d}s \;+\; g\big(X_T, \mu_T\big) \;\Big|\; X_t = x \right]$$

  solves the (backward in time) **Hamilton–Jacobi–Bellman** equation:

$$\begin{cases} \partial_t V^\mu(t,x) \;+\; \inf_\alpha \Big\{ f(x,\mu_t,\alpha) + \nabla V^\mu(t,x) \cdot b(x,\mu_t,\alpha) + \dfrac{1}{2}\,\Delta V^\mu(t,x) \Big\} = 0, \\ V^\mu(T,x) \;=\; g(x,\mu_T). \end{cases}$$

- The fixed point step is implemented by requiring $\mu = (\mu_t)$ solves the (forward) **Fokker–Planck equation**:

$$\partial_t \mu_t \;=\; \tfrac{1}{2}\,\Delta \mu_t - \nabla \cdot \Big(b\big(x, \mu_t, \alpha^\mu\big)\,\mu_t \Big).$$

- This is a system of strongly coupled nonlinear PDEs!
- There is also a popular **probabilistic approach** using forward-backward stochastic differential equations (FBSDEs).

Key progresses over the years

Theory

- ◇ **Existence** of MFG equilibria.
- ◇ **Uniqueness** in monotone settings.
- ◇ **Approximate Nash equilibria** for N-player games from MFG limit.
- ◇ **Convergence** of n-player games to MFG limit
- ◇ **Computation** guarantees: convergence of iterative schemes.
- ◇ ...

Applications

- ◇ Systemic risk and interbank lending.
- ◇ Flocking and herding models in biology
- ◇ Crowd motion and congestion.
- ◇ Algorithmic trading and execution.
- ◇ Cybersecurity, bank runs
- ◇ ...

Many active areas/open problems!
Read some prerequisites on the next page and come talk to me.

Prerequisite knowledge

- Foundational
  - Probability (e.g., *Probability with Martingales* by Williams).
  - Stochastic processes: Brownian motion, SDEs, stochastic calculus (e.g. *Brownian Motion and Stochastic Calculus* by Karatzas and Shreve).
- More advanced
  - McKean–Vlasov equations, interacting particle systems, mean field games (e.g. *Mean Field Games and Interacting Particle Systems* by Daniel Lacker, *Probabilistic Theory of Mean Field Games with Applications* by René Carmona and François Delarue)
  - Stochastic control theory (e.g. *Continuous-Time Stochastic Control and Optimization with Financial Applications* by Huyên Pham)

Thank you very much for your attention!