International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

STOCHASTIC GAME THEORY AND MEAN-FIELD MODELS FOR AUTONOMOUS VEHICLE COORDINATION

[1]Dr. J. Prakash, [2]Dr. Sam Abraham, [3]G. Lakshmi, [4]A. Nishanandhini, [5]Dr. D. Karthikeswaran, [6]R Shri Keerthana

[1]Professor, Department of Information science and Engineering,

Moodlakatte Institute of Technology,

Kundapura, Udupi district, Karnataka, India

Orc id: 0000-0003-0343-2047

prakashj@mitkundapura.com

[2]Associate Professor, Department of Computer Science and Engineering (Data Science),

Moodlakatte Institute of Technology,

Kundapura, Udupi District, Karnataka, India

Orc id: 0000-0001-8355-863X

sam@mitkundapura.com

[3]Assistant Professor, Department of Information Science and Engineering,

Moodlakatte Institute of Technology,

Kundapura, Udupi district, Karnataka, India

Orc id: https://orcid.org/0009-0004-2238-2617

lakshmi@mitkundapura.com

[4]Assistant Professor, Department of Computer Science and engineering

Hindusthan Institute of Technology

Nishanandhini.a@hit.edu.in

[5]Professor Department of Computer Science and Engineering,

CMS College of Engineering and Technology,

Appachigoundenpathy, Kumittipathi Post Coimbatore, India

karthids@gmail.com

[6]Assistant professor Department of Computer Science and Design

Erode Sengunthar Engineering College

shri4keerthana@gmail.com

Received: August 04, 2025

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

1. Introduction

Coordinating large numbers of autonomous vehicles (AVs) on roads is a complex, uncertain, and decentralized control problem. Each vehicle must decide in real time how to accelerate, brake, change lanes, merge, or cross intersections while anticipating and influencing the actions of many others—some of which are human-driven and inherently unpredictable. The problem is made more difficult by partial observability (no vehicle sees everything), communication constraints, noisy sensors, delays, heterogeneous objectives (comfort, energy, time), safety requirements, and non-stationary environments.

Two mathematical frameworks have emerged as especially promising for principled coordination at scale: **stochastic game theory** and **mean-field models**. Stochastic (dynamic) games capture the strategic interdependence of a finite set of agents whose states and payoffs evolve in response to joint actions and exogenous randomness. **Mean-field games (MFGs)** approximate the influence of a very large population by an aggregate distribution rather than explicit agent-to-agent coupling, enabling scalable planning and learning with performance guarantees such as ε-Nash optimality as the population grows.

This article explains how these frameworks model AV coordination, what equilibrium notions and algorithmic tools they use, how learning is integrated, and how safety, robustness, and human factors are addressed. We cover canonical tasks—platooning, lane changes, merges, intersections, and speed harmonization—then discuss hybrid architectures that combine game-theoretic planning with model predictive control, deep reinforcement learning, incentive design, and formal verification. The goal is a practical, engineering-oriented view with enough theory to justify why these methods work and where they still fall short.

2. Why Games? Problem Structure in Traffic

Strategic Coupling

AV decisions are **strategic**: my acceleration affects your feasible actions (e.g., closing a gap prevents your lane change) and your actions affect my risk and progress. Classic optimal control treats other agents as disturbances; this is conservative and often inefficient. Game theory explicitly models **best responses** and **equilibria**, capturing mutual adaptation (Mo et al., 2024).

Uncertainty and Dynamics

Road dynamics are stochastic due to sensor noise, friction, driver variability, and random arrivals. **Stochastic dynamic games** extend Markov decision processes (MDPs) to multi-agent settings, with state transitions that depend on joint actions and noise. This matches road realities: the same maneuver has different outcomes in different traffic states.

Scale

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Urban highways may have thousands of vehicles. Modeling all pairwise interactions is intractable; **mean-field** models replace many-to-many coupling with **one-to-distribution** interactions, dramatically reducing complexity while preserving aggregate effects like density waves and shock fronts.

3. Stochastic Game Formulations for AVs

Consider a finite set of players $i = 1, \ldots, N$(vehicles), discrete time t, and global state $s_t$that includes positions, velocities, gaps, and possibly environment variables (signals, weather). Player $i$selects an action a_t^i(acceleration, lane change, merge accept/reject). The joint action $a_t = (a_t^1, \ldots, a_t^N)$induces a transition $s_{t+1} \sim P(\cdot \mid s_t, a_t)$. Each player has a reward $r^i(s_t, a_t)$capturing progress, comfort (jerk penalties), energy use, and safety (barrier penalties or hard constraints).

Equilibrium Notions

- **Markov Perfect Equilibrium (MPE):** Each player uses a stationary policy $\pi^i(a^i \mid s)$that is optimal given others' stationary policies, based only on the current state. MPE is natural for ongoing traffic but computing it is hard when $N$is moderate to large and state spaces are continuous.
- **Correlated Equilibrium (CE):** A coordinating signal (e.g., a roadside unit message) draws actions from a joint distribution; no agent benefits by deviating given its signal (Zheng et al., 2025). CE can exploit communication to reduce inefficiency without requiring binding agreements.
- **Stackelberg Equilibria:** Useful for hierarchical scenarios (e.g., an infrastructure controller or platoon leader acts first; followers best respond). This captures **cooperative-competitive** patterns such as ramp meters influencing on-ramp merges.

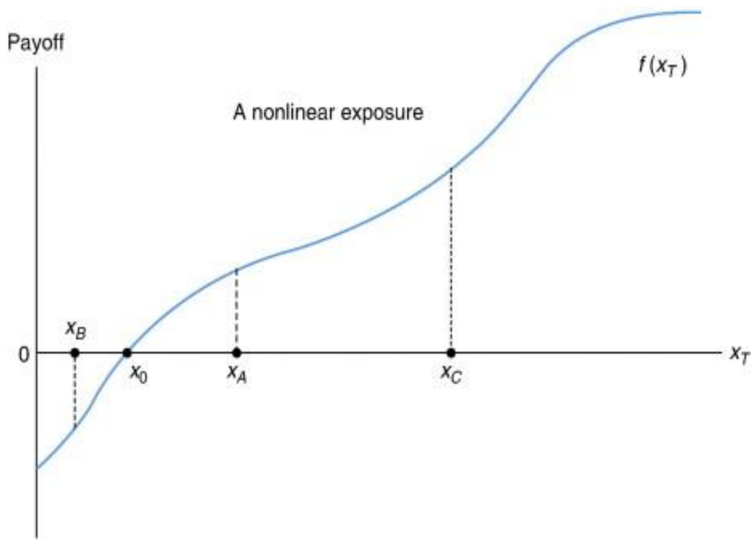

Figure 1: Payoff surface & best-response curves

Received: August 04, 2025

International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

(Source: sciencedirect.com, 2025)

Types of Games in AV Settings

- **Potential games** (e.g., lane allocation with congestion costs) admit Lyapunov-like potential functions; best-response dynamics converge.
- **Congestion games** model lane choice where cost depends on the number of users; equilibria reflect Wardrop principles (Elokda et al., 2024).
- **Stochastic pursuit-evasion or differential games** appear in collision avoidance at short time scales.

Constraints and Safety

Safety can be encoded as:

- **Hard constraints** $g(s_t, a_t) \leq 0$(no collision, keep headway).
- **Control barrier functions (CBFs)** producing safe action sets; games with **safe feasible sets** avoid unsafe equilibria.
- **Chance constraints** $\mathbb{P}(\text{collision}) \leq \delta$, yielding risk-aware policies.

4. Mean-Field Games (MFGs) for Scalable Coordination

Basic Idea

As $N \to \infty$, the influence of any single agent on the population is negligible. Each representative vehicle interacts with the **mean field** m_t, the distribution of states (positions, velocities, lanes). The agent solves an optimal control problem assuming $m_t$evolves according to a **Fokker–Planck** (continuous time) or **Kolmogorov forward** (discrete time) equation induced by the common policy.

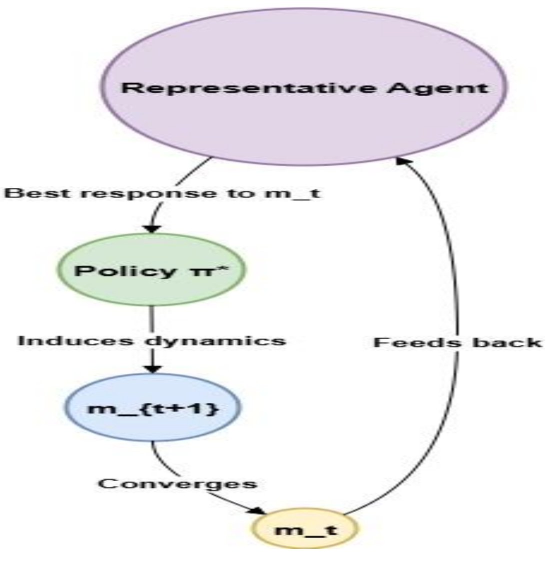

Figure 2: Mean-Field Game Fixed-Point Loop

Received: August 04, 2025

International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

(Source: draw io)

Fixed points where the policy is optimal given $m_t$and $m_t$is generated by the policy constitute an **MFG equilibrium**. In large but finite populations, such policies are **ε-Nash**: no agent can gain more than ε by unilateral deviation.

Why MFGs Suit Traffic

- **Local interactions, global effects:** Individual car-following creates macroscopic density waves; MFGs capture this micro-macro loop.
- **Heterogeneity via types:** Vehicle types (car, truck, AV level), risk preferences, and target speeds can be modeled as type-indexed distributions.
- **Computational tractability:** Solving for a distribution over states is cheaper than simulating thousands of explicit interactions in planning loops.

Common MFG Structures in AVs

- **LQG-MFG (Linear-Quadratic-Gaussian):** Linear dynamics (e.g., velocity, headway), quadratic costs (deviation from desired speed and spacing), Gaussian noise. Leads to Riccati equations and fixed-point coupling; analytically tractable for speed harmonization.
- **Discrete-state MFGs:** Lane choice, route choice, and gap-acceptance can be modeled on graphs with transition probabilities depending on local densities (Zhang et al., 2022).
- **Risk-sensitive and entropy-regularized MFGs:** Incorporate CVaR or exponential utility for safety/comfort; entropy (or KL) regularization encourages smooth, exploratory policies.

5. Canonical Coordination Problems

Platooning and String Stability

Goal: Maintain tight formations to reduce drag and improve capacity without amplifying disturbances (string stability).

- **Stochastic game view:** Each vehicle’s action influences the spacing errors of followers. Payoffs penalize spacing error, control effort, and safety violations. MPE or potential-game formulations support decentralized policies with convergence guarantees.
- **MFG view:** Treat the platoon as a dense line where each vehicle reacts to the mean spacing/velocity distribution; yields feedback gains derived from LQG-MFG equations.
- **Safety:** CBFs ensure no rear-end collisions; stochastic headway constraints or probabilistic time-to-collision limits produce risk-aware spacing.
- **Learning:** Multi-agent policy-gradient with opponent modeling (or fictitious play against population statistics) adapts to mixed human/AV driving.

Merging (On-ramp, zipper merge)

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Goal: Allow equitable, safe, and efficient merges when multiple streams compete for limited gaps.

- **Stochastic game:** Short horizon, severe coupling; equilibrium resembles bargaining over gaps. Stackelberg variants model a leader vehicle that yields or asserts priority (Hua and Fan, 2023).
- **MFG:** Approximate mainline and ramp flows by densities; compute optimal merging rates and gap creation policies that minimize delay and collision risk.
- **Incentives:** Infrastructure can issue **pricing** or **priority tokens** to shape equilibria toward socially efficient throughput.

Lane Changing and Overtaking

Goal: Decide when to change lanes given stochastic payoffs (speed gain vs. risk).

- **Game:** Discrete moves with externalities (cut-in reduces downstream headway). Correlated equilibrium emerges if V2V signals coordinate yields.
- **MFG:** Lane occupancy is the mean field; transitions follow softmax of estimated advantages penalized by lane-change friction. Equilibrium conditions balance marginal benefit with congestion externality.

Unsignalized Intersections

Goal: Resolve right-of-way without a traffic light.

- **Game:** Timing game with incomplete information about acceleration intents; equilibrium may be inefficient (both wait). Correlated devices (roadside beacons) improve efficiency.
- **MFG:** Arrival processes feed density at conflict points; solve for crossing rates and reservation policies (time slots) to minimize total delay subject to safety buffers. ε-Nash guarantees scale to many approaches.

Speed Harmonization and Shockwave Damping

Goal: Smooth speeds upstream of bottlenecks to avoid stop-and-go waves.

- **MFG-LQG:** Desired speed tracks a density-dependent target; optimal policies reduce variance.
- **Mechanism:** A small fraction of compliant AVs can dissipate waves by acting on the mean field; this is a **control-by-sparsity** effect.

International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6. Learning in Games and Mean-Field Systems

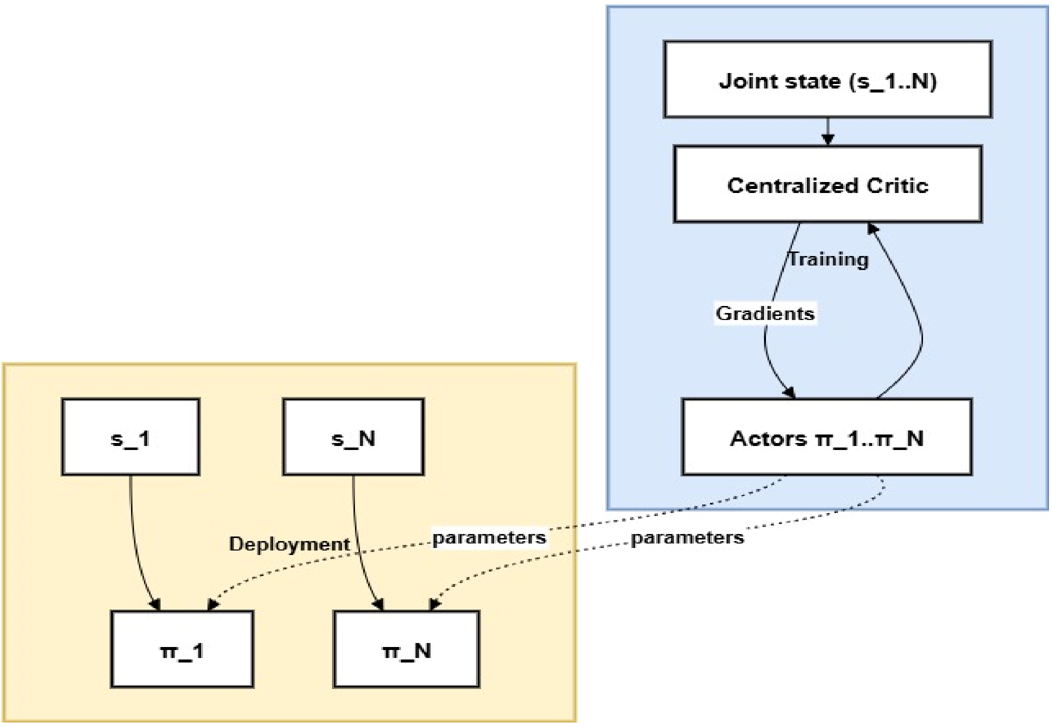

Figure 3: Centralized Training and Decentralized Execution (CTDE) Framework

(Source: draw io, 2025)

Model-based vs. Model-free

- **Model-based:** Estimate dynamics, then solve stochastic game/MFG by dynamic programming or Riccati-type equations; more sample-efficient and interpretable.
- **Model-free MARL:** Q-learning or actor–critic extended to multi-agent settings with centralized training and decentralized execution (CTDE). Useful when dynamics are complex or partially unknown.

Equilibrium-Seeking Algorithms

- **Fictitious play / mean-field fictitious play:** Iteratively best-respond to empirical distributions; converges in potential games and many MFGs.
- **No-regret learning:** Guarantees convergence to coarse correlated equilibria; compatible with bandit feedback.
- **Policy gradient with mean-field coupling:** Alternate between solving the representative agent's RL problem (given m_t) and updating $m_t$via population rollouts until fixed point.

Dealing with Partial Observability

Traffic is partially observable: occlusions, limited range, missing intentions. Tools include:

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- **POMDPs and belief states** per agent.
- **Centralized critics** with privileged information at training time; decentralized recurrent policies at test time (An et al., 2023).
- **Information-sharing protocols** (V2V/V2I) that implement cheap signals to approximate correlated equilibria.

Robustness and Risk

- **Distributional RL** provides estimates of return distributions; enables CVaR-aware decisions under rare but catastrophic events.
- **Robust MFGs** assume worst-case model perturbations (ambiguity sets), producing conservative yet stable policies—useful near safety boundaries or during weather changes.

7. Safety, Guarantees, and Verification

Safety Shields and Barrier Certificates

A common architecture learns a high-performance policy then filters unsafe actions through a **safety shield**:

- **Control Barrier Functions (CBFs)** define a safe set $\mathcal{S}$; a quadratic program minimally adjusts the action to keep trajectories within $\mathcal{S}$.
- **Game-aware shields:** The feasible set depends on predicted opponent actions; conservative approximations use reachable sets under opponent bounds.

Temporal Logic and Compositional Reasoning

Mission specs (e.g., “always avoid collision,” “eventually merge”) can be encoded in **Signal Temporal Logic (STL)** or LTL. Controllers are synthesized to satisfy specs with high probability. In multi-agent settings, **assume-guarantee** contracts decompose verification across agents or lanes.

Equilibrium Stability and Price of Anarchy

Game equilibria may be inefficient compared with a centralized optimum. **Price of anarchy** quantifies the ratio. Tools:

- **Potential games** with convex costs often have small inefficiency bounds.
- **Incentive mechanisms** (tolls, priority auctions, platoon credits) align private and social costs.

8. Communication, Incentives, and Mechanism Design

Limited V2V/V2I

Communication reduces uncertainty but is bandwidth-limited and unreliable. Design principles:

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- Send **minimal sufficient statistics** (estimated intent horizon, preferred gap).
- Use **common randomization** (a seed) to induce correlated equilibria without heavy messaging.
- Apply **reservation systems** at intersections, granting time-space slots; this is a Stackelberg mechanism with the intersection as leader (Ren et al., 2023).

Incentivizing Cooperative Behavior

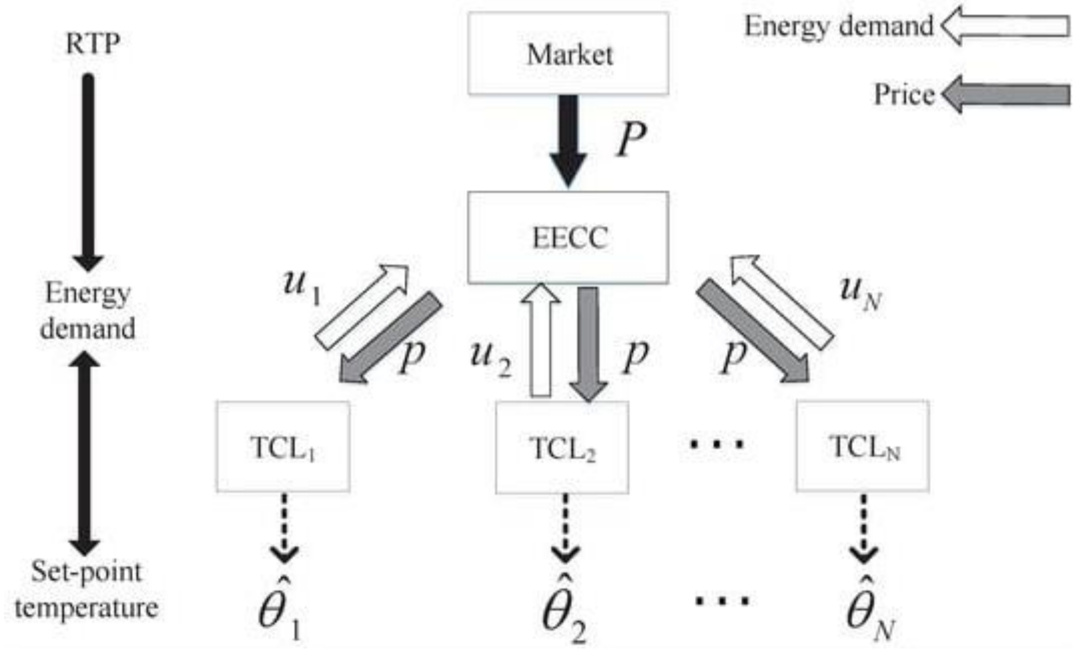

Figure 4: Stackelberg controller with pricing loop

(Source: mdpi.com, 2025)

If AVs are owned by different stakeholders, self-interest might create aggressive equilibria. Mechanisms to improve social welfare:

- **Dynamic tolls** calibrated to density (Pigouvian taxes).
- **Token-based priority** tradable for future advantages (e.g., faster merge next time).
- **Reputation systems** affecting right-of-way probability.

9. Mixed Traffic with Humans

AVs must coexist with human drivers exhibiting bounded rationality, diverse risk attitudes, and occasional rule-breaking.

- **Behavioral games:** Model humans as **level-k** or **quantal response** players (probabilistic best responses), then compute AV best responses.

International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- **Inverse reinforcement learning (IRL):** Learn human reward parameters; embed them in the game.
- **Mean-field with types:** Human types (aggressive/cautious) are proportions in the mean field; AV policies adapt to type distribution estimates.
- **Safety buffers:** Increase CBF margins near human-dense zones; robustify against unpredictable cut-ins.

10. Algorithmic Blueprints

Two-Loop MFG-RL for Speed Harmonization

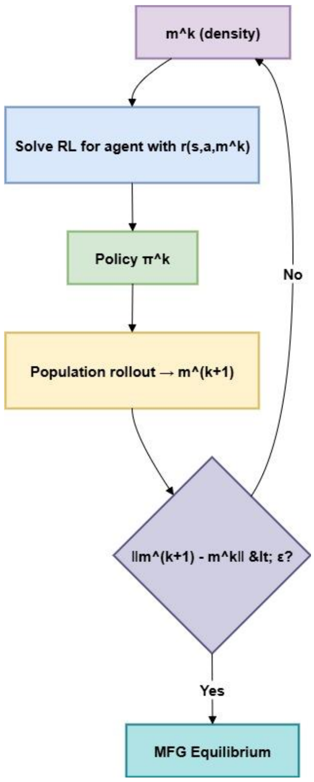

Figure 5: MFG Iterative Learning Process

(Source: draw io, 2025)

1. **Initialize** a mean-field density m^0over longitudinal states.
2. **Inner loop:** Solve a single-agent RL problem with reward $r(s, a, m^k)$(penalizing deviations from target speed/density, energy, discomfort). Train a policy π^k.
3. **Population rollout:** Simulate many agents using $\pi^k$to obtain m^{k+1}.
4. **Fixed-point step:** If $\| m^{k+1} - m^k \|$small, stop; else $k \leftarrow k + 1$.

Enhancements: entropy regularization for smoother controls; CVaR penalties to manage tail risks; CBF-QP safety layer during rollouts.

Game-Aware Lane Change with CTDE

- **State:** ego kinematics, gaps, relative speeds, neighbor intents (beliefs).
- **Action:** accelerate/decelerate, keep lane, change left/right.
- **Training:** centralized critic observes joint state; actors are decentralized.
- **Opponent modeling:** predict neighbors as quantal responders to incentives; train with randomized opponent types for robustness (Stella, Bauso and Colaneri, 2021).
- **Deployment:** shield with headway/lateral safety constraints; tie-break with correlation device when two AVs conflict.

Intersection Reservation Mechanism

- **Stackelberg stage:** intersection publishes pricing/slots.
- **Follower responses:** vehicles compute utility (delay vs. price) and bid.
- **Allocation:** solve a convex program (or greedy) to assign non-conflicting time-space trajectories; guarantee no collision by design.
- **Learning prices:** use online dual ascent so average occupancy targets are met and queues remain stable.

11. Data, Simulation, and Evaluation

Data Sources

High-resolution trajectory datasets (e.g., naturalistic highway and urban recordings) calibrate game payoffs and noise models. Simulation platforms (microscopic simulators, physics-based engines) allow **in-silico** evaluation of equilibria, price-of-anarchy, and robustness.

Metrics

- **Safety:** collision rate, minimum time-to-collision distribution, constraint violation probability.
- **Efficiency:** average speed, throughput, delay, merge success rate.
- **Comfort/energy:** jerk variance, fuel or kWh per km.
- **Stability:** damping of stop-and-go waves (string stability metrics).
- **Fairness:** delay dispersion across agents/types.
- **Scalability:** compute time vs. agent count; communication load.

Experimental Protocols

- **A/B equilibria:** compare selfish MPE vs. incentivized equilibria.
- **Population shifts:** stress-test policies under changes in human/AV mix and demand surges.

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- **Domain shifts:** weather, sensor noise, degraded comms.
- **Safety cases:** evaluate worst-case adversarial cut-ins and hard-brake events with risk-sensitive returns.

12. Theoretical Guarantees and Practical Limits

Existence and Uniqueness

- Many MFGs admit equilibria under monotonicity (Lasry–Lions conditions) or contractive couplings (Liu et al., 2023).
- Uniqueness matters: multiple equilibria can cause indecision or oscillations. Mechanisms and correlation devices help select the desirable equilibrium.

ε-Nash Optimality

In large-population settings, the MFG policy is an ε-Nash equilibrium for the finite game; ε decays as $1/\sqrt{N}$(typical order under regularity). This justifies deploying MFG policies in realistic (but finite) traffic.

Sample Complexity and Non-Stationarity

MARL sample complexity grows with dimensionality and non-stationarity. CTDE and mean-field approximations mitigate this, but real-world deployment requires curriculum learning, domain randomization, and continual adaptation.

Robustness Gaps

Model mismatch (e.g., unusual human maneuvers) can invalidate equilibrium assumptions. Robust MFGs and distributional RL reduce sensitivity, but certification remains challenging, especially under rare multi-vehicle edge cases.

13. Hybrid Architectures: From Theory to On-Road Control

Practical stacks often integrate multiple layers:

1. **Strategic layer (seconds to minutes):** Game/MFG planner sets intents: target speed, desired lane, merge time windows, platoon targets. Could run in infrastructure (V2I) for coordination at bottlenecks.
2. **Tactical layer (hundreds of ms to seconds):** Local game with a handful of neighbors refines maneuver specifics. Use MARL policies trained with population statistics or explicit best-response computation in small games.
3. **Servo layer (tens of ms):** Model predictive control with **CBF-QP** ensures feasibility and safety; handles actuation limits and tracking (Schwarting et al., 2021).
4. **Verification layer:** Runtime monitors enforce temporal logic constraints; conflict with learned policy triggers safe fallback (e.g., decelerate and yield).

Received: August 04, 2025

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

5. **Learning layer:** Check of category membership of human categorizations, density maps and value functional modification; both of which are checked by offline and online detecting distortion.

This game (MPC) is the medium in this case that safeguards the conjunction of the predictive balance logic and the potent safety protocols.

14. Case Studies (Conceptual)

Speed Harmonization with Sparse AV Penetration

A density-sensitive policy of 10-20 percent AV penetration may help lessen its effects through a density-dependent speed MFG-LQG policy. AVs are roadblocks that slow the pace in the traffic so that the entire traffic will move and some slight refund on the payment will work to compensate this bringing the case to what the society perceives as good.

Fair Merging via Tokens

Self-serving equilibria present aggressive alterations in the lanes in case there is vacuity in a congested on-ramp. On a Stackelberg token system, converging priorities are granted; vehicles are given tokens by engaging in cooperative actions (e.g. opening gaps upstream) and spent tokens when required. The result of such a stability, which follows, as that the rate of near-collections is less, as the postponements.

Lane Change Courtesy Protocol

Bit-single signal is transmitted with a certified device: ego is in such position. The induced CE is not in deadlocks as vehicles are responding to a probability, which also causes minimal changes in the lanes. The performance improvement results will follow after the endurance of the coordination opportunities, the spirit of 5-10 was not taken into account.

15. Engineering Recommendations

1. Start with large scale benchmarks using large size scales: calculate density sensitive target strategies: solve these problems, locally, using small-games.

2. Use CTDE MARL: protective barriers: the use of a centralized evaluator can be investigated in order of enhancing the speed of learning; the use of CBF-QP screens can be implemented to ensure its limitation.

3. Correlation leverage methods: V2V does not need a lot of bandwidth to make it as virtually performant as CE.

4. Failure The introduction of incentives on the part of institutions is needed, that is, prices, tokens, even preference can shake non-effective equilibria.

5. Model human heterogeneity: maintain online estimates of driver types; adapt courtesy and gap-creation policies.

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6. Verify with temporal logic monitors: test above the usually-case safety mission-level guarantee.

7. Embrace risk-sensitive objectives: tune CVaR to control tail risks in merges and intersections (Qu et al., 2022).

8. Plan for non-stationarity: domain randomization, continual learning, and anomaly detectors handle weather, events, and fleet updates.

9. Measure price of anarchy: quantify inefficiency and evaluate mechanisms that narrow the gap to social optimum.

10. Deploy incrementally: begin with infrastructure-supported sites (ramp meters, cooperative intersections) where coordination benefits are strongest.

16. Open Challenges

- **Heterogeneous perception and latency:** Equilibria derived under synchronized time may degrade with asynchronous sensing and actuation delays.
- **Multi-modal objectives:** Preference aggregation is required to make sure that there is some trade of comfort, energy and time with mis-specification that may lead to wrong-headed equilibria.
- **Ambiguity in human intent:** Predictive uncertainty is greater in the urban environment that is characterized by high clutter and which can be reduced due to the presence of the conservative safety margin.
- **Scaling verification:** The verification of enriched and learned policies in saturated traffic can be scaled with the help of compositional techniques, which are now being converted into a push-button (Zhu et al., 2023).
- **Security and adversarial behavior:** Malicious actors can use the coordination protocols and should be able to detect and offer high protection to such actors.
- **Regulatory and ethical concerns:** The incentives and priority systems should be based on the standards of fairness and the legal right-of-way.

17. Conclusion

Stochastic game theory as well as the mean-field models can provide scalable, secure and effective autonomous vehicle coordination. The concepts of the stochastic game and the equilibrium (MPE, CE, Stackelberg) can be used to summarize the strategic, uncertain and interactive features of the local behavior merger, lane changes and the crossing intersections and prediction, as well as to organize behavior. Mean-field games are akin to the case of corridors and networks in that they simulate the impact of the population in terms of distributions and solutions to LQG models are the analysis and fixed-point learning applications to complex models.

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Marrying these frameworks with modern tools—centralized-training MARL, safety shields via barrier functions, temporal logic monitors, and mechanism design—yields deployable architectures that both anticipate others and guarantee safety. While open challenges remain in robustness, verification at scale, and human variability, the combined game-theoretic and mean-field approach has demonstrated that a small fraction of well-coordinated AVs can stabilize flows, increase throughput, and reduce risk. The strategic, population-level perspective is not a theoretical luxury but a practical necessity for the next generation of intelligent transportation systems.

References

Journals

1. An, L., Yang, X., Kim, D.K., Hu, J., Wang, G. and Xu, Z., 2023. Speed harmonization for partially connected and automated traffic. *IEEE Transactions on Intelligent Transportation Systems*, *24*(10), pp.10491-10500.
2. Elokda, E., Bolognani, S., Censi, A., Dörfler, F. and Frazzoli, E., 2024. Dynamic population games: A tractable intersection of mean-field games and population games. *IEEE Control Systems Letters*, *8*, pp.1072-1077.
3. Hua, C. and Fan, W., 2023. Dynamic speed harmonization for mixed traffic flow on the freeway using deep reinforcement learning. *IET Intelligent Transport Systems*, *17*(12), pp.2519-2530.
4. Liu, M., Kolmanovsky, I., Tseng, H.E., Huang, S., Filev, D. and Girard, A., 2023. Potential game-based decision-making for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, *24*(8), pp.8014-8027.
5. Mo, Z., Chen, X., Di, X., Iacomini, E., Segala, C., Herty, M. and Lauriere, M., 2024. A game-theoretic framework for generic second-order traffic flow models using mean field games and adversarial inverse reinforcement learning. *Transportation Science*, *58*(6), pp.1403-1426.
6. Qu, D., Zhang, K., Song, H., Jia, Y. and Dai, S., 2022. Analysis and modeling of lane-changing game strategy for autonomous driving vehicles. *IEEE Access*, *10*, pp.69531-69542.
7. Ren, Z., Tan, X., Touzi, N. and Yang, J., 2023. Entropic optimal planning for path-dependent mean field games. *SIAM Journal on Control and Optimization*, *61*(3), pp.1415-1437.
8. Schwarting, W., Pierson, A., Karaman, S. and Rus, D., 2021. Stochastic dynamic games in belief space. *IEEE Transactions on Robotics*, *37*(6), pp.2157-2172.
9. Stella, L., Bauso, D. and Colaneri, P., 2021. Mean-field game for collective decision-making in honeybees via switched systems. *IEEE Transactions on Automatic Control*, *67*(8), pp.3863-3878.

International Journal of Applied Mathematics

Volume 38 No. 6s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

10. Zhang, H., Lu, C., Tang, H., Wei, X., Liang, L., Cheng, L., Ding, W. and Han, Z., 2022. Mean-field-aided multiagent reinforcement learning for resource allocation in vehicular networks. *IEEE Internet of Things Journal*, *10*(3), pp.2667-2679.
11. Zheng, L., Wang, X., Li, F., Mao, Z., Tian, Z., Peng, Y., Yuan, F. and Yuan, C., 2025. A Mean-Field-Game-Integrated MPC-QP Framework for Collision-Free Multi-Vehicle Control. *Drones*, *9*(5), p.375.
12. Zhu, L., Yang, D., Cheng, Z., Yu, X. and Zheng, B., 2023. A model to manage the lane-changing conflict for automated vehicles based on game theory. *Sustainability*, *15*(4), p.3063.

Websites

13. mdpi.com, (2025). Stackelberg controller with pricing loop. [Online] Available at: https://www.mdpi.com/2076-3417/8/8/1370 [Accessed on: 08/10/2025]
14. sciencedirect.com, (2025). *Payoff surface & best-response curves*. [Online] Available at: https://www.sciencedirect.com/topics/mathematics/payoff-function [Accessed on: 08/10/2025]