Entangled Traffic: A Quantum Game-Theoretic Framework for Lane-Changing Cooperation in Mixed Traffic

Sungyong Chung*[1], Tina Radvand[1], and Alireza Talebpour[1]

[1]Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Short summary

As automated vehicles (AVs) enter mixed traffic, proactively anticipating human driving behavior during interactions like lane changes is essential. Classical Evolutionary Game Theory (EGT) assumes driver independence, predicting unrealistic full cooperation that contradicts the stable 42% cooperation rate observed in real-world Waymo Open Motion Dataset (WOMD) trajectories. To resolve this discrepancy, this study introduces a Quantum Game Theory (QGT) framework utilizing the Marinatto-Weber quantization scheme. By embedding latent correlations between drivers into a single interaction's payoff structure via an entanglement parameter ($|b|^2_{HDV} \approx 0.52$), our model accurately reproduces the observed mixed equilibrium. Furthermore, simulations of various AV deployment strategies reveal that human adaptation depends critically on the underlying AV algorithm. This framework enables stakeholders to simulate repeated interactions and anticipate how human drivers will evolve in response to specific AV software designs, ensuring safer mixed-traffic integration.

Keywords: Automated Vehicle, Cooperation, Evolutionary Game Theory, Lane-Changing, Quantum Game Theory

1 Introduction

The transition toward mixed traffic systems, characterized by the simultaneous operation of human-driven vehicles (HDVs) and automated vehicles (AVs), introduces highly complex interactive challenges. During maneuvers such as lane changing, drivers do not act in isolation. They must continuously anticipate the intentions of surrounding vehicles. While the behavior of AVs during lane changes has been extensively studied (An et al., 2023; Chung et al., 2024), the interactive evolution of behavior between opportunistic human drivers and programmed AVs remains an open question (Chung et al., 2025). It is vital to determine whether human drivers will exploit cautious AVs or if AV presence will inadvertently foster a more cooperative traffic environment.

Recent empirical analyses utilizing the Waymo Open Motion Dataset (WOMD) demonstrated that during lane-changing interactions, human drivers exhibit a stable cooperation rate of approximately 42% (Chung et al., 2025). However, when these empirically extracted behavioral payoffs are simulated using classical Evolutionary Game Theory (EGT), the traffic population rapidly diverges from reality, converging instead to an unrealistic state of near-full cooperation. This discrepancy exposes a foundational flaw in classical EGT applications for traffic: the assumption of statistical independence. Classical games presume that each interaction is an isolated event between independent agents. In reality, however, human drivers bring established social norms, regional driving cultures, and latent psychological correlations into every encounter.

To overcome this theoretical barrier, this research adapts concepts from Quantum Game Theory (QGT) to model traffic interactions. Specifically, we utilize quantum entanglement not as a physical phenomenon, but as a robust mathematical formalism to embed these latent behavioral correlations directly into the game's payoff matrix. By introducing specific AV behavioral profiles into a quantized simulation environment, this paper provides a novel platform for anticipating how human driving strategies will evolve across various stages of AV market penetration.

2 Methodology

This study constructs a data-driven simulation framework that maps observed lane-changing trajectories into a quantized game-theoretic environment.

Empirical Data and Utility Extraction

The baseline interactions are derived from 7,636 lane-changing events (Chung et al., 2025) extracted from the WOMD (Ettinger et al., 2021). Each event is treated as a two-player game involving an active vehicle (initiating the lane change) and a passive vehicle (the lag vehicle in the target lane). Using the k-means algorithm, we categorized the kinematic behaviors of the active and passive vehicles during a lane change into cooperative (C) and defective (D) strategies. Linear utility functions were then fitted to the pre-lane-change state variables (e.g., relative speeds, accelerations, and gap sizes) to quantify the empirical payoffs for mutual cooperation (R), mutual defection (P), temptation (T), and the sucker's payoff (S). In other words, we calibrated the payoffs for each behavior based on a given set of pre-lane-change state variables to understand which strategy (cooperative or defective) the active and passive vehicles would choose. Under classical rules, these utilities dictate the fitness of a driving strategy over time.

Game Quantization Schemes

To relax the classical assumption of independent decision-making, we look to quantization schemes. The two most prominent methods for quantizing classical games are the Eisert-Wilkens-Lewenstein (EWL) scheme (Eisert & Wilkens, 2000) and the Marinatto-Weber (MW) scheme (Marinatto & Weber, 2000). For a comprehensive overview of these two schemes, see Khan et al. (2018).

The EWL scheme utilizes a quantum circuit approach where an initial separable state is entangled via an operator $\hat{J}$ before players apply continuous unitary strategies. A disentangling gate $\hat{J}^\dagger$ is then applied prior to measurement. While EWL successfully demonstrates how players can escape the sub-optimal Nash equilibrium of the Prisoner's Dilemma to achieve Pareto-optimal mutual cooperation, its continuous strategy space is less suited to the binary (cooperate vs. defect) nature of our empirically observed driving behaviors.

Conversely, the MW scheme omits the disentangling gate and restricts the strategy space to standard classical operators, specifically the identity operator $\hat{I}$ (representing cooperation) and the Pauli-X operator $\hat{\sigma}_x$ (representing defection). In the MW framework, the game originates from an already entangled initial state defined in the computational basis:

$$|\psi_{in}\rangle = a|00\rangle + b|11\rangle \tag{1}$$

where $|a|^2 + |b|^2 = 1$. Here, the parameter $|b|^2 \in [0, 1]$ explicitly quantifies the degree of latent correlation, or "entanglement," between the two drivers prior to their strategic choice.

Evolutionary Stability in the Prisoner's Dilemma

The power of the MW scheme lies in its ability to alter the Evolutionarily Stable Strategy (ESS) of a population based solely on the correlation parameter b. In a classical Prisoner's Dilemma, defection strictly dominates, making mutual defection the unique ESS. However, when the MW scheme is applied to the Prisoner's Dilemma, the introduction of entanglement fundamentally rewrites the expected payoff structures.

Previous theoretical analyses (Iqbal & Toor, 2001) have proven that the evolutionary dynamics of this game transition through three distinct phases depending on the magnitude of $|b|^2$. When entanglement is low ($|b|^2 \leq 1/3$), the game mirrors classical mechanics, and pure defection remains the ESS. When entanglement is extraordinarily high ($|b|^2 \geq 2/3$), the strong latent correlation flips the dilemma entirely, stabilizing pure cooperation.

Crucially, in the intermediate regime ($1/3 < |b|^2 < 2/3$), the system undergoes a phase transition where neither pure strategy is stable. Instead, the unique ESS becomes a mixed strategy. This mathematical property is vital for traffic modeling, as it provides a mechanistic explanation for why human drivers naturally settle into a stable equilibrium of partial cooperation (42%) rather than absolute cooperation or total aggression.

Traffic Simulation Framework

Applying the MW logic to our empirically derived traffic payoffs (R, S, T, and P), the quantum expected payoffs (R^q, S^q, T^q, P^q) for a vehicle in a specific lane-changing state $\mathbf{S}_i$ are calculated as a weighted superposition of the classical payoffs:

$$R^q(\mathbf{S}_i) = R(\mathbf{S}_i)(1 - |b|^2) + P(\mathbf{S}_i)|b|^2 \tag{2}$$

$$S^q(\mathbf{S}_i) = S(\mathbf{S}_i)(1 - |b|^2) + T(\mathbf{S}_i)|b|^2 \tag{3}$$

$$T^q(\mathbf{S}_i) = T(\mathbf{S}_i)(1 - |b|^2) + S(\mathbf{S}_i)|b|^2 \tag{4}$$

$$P^q(\mathbf{S}_i) = P(\mathbf{S}_i)(1 - |b|^2) + R(\mathbf{S}_i)|b|^2 \tag{5}$$

Evolution is simulated on a 20×20 lattice grid. Agents iteratively interact with neighbors, accumulating payoffs based on equations (2-5). Using a standard pairwise comparison mechanism, vehicles periodically update their strategy by mimicking more successful neighbors, subject to a noise parameter. We assign specific entanglement parameters to represent distinct vehicle types, denoted as $|b|^2_{HDV}$ for human drivers and $|b|^2_{AV}$ for automated agents. For cross-type interactions, the effective entanglement is the arithmetic mean of the two parameters.

3 Results and discussion

Calibrating Human Entanglement

To establish a baseline, we first simulated a 100% HDV environment to calibrate the $|b|^2_{HDV}$ parameter against the empirical WOMD data. As illustrated in Figure 1, classical assumptions ($|b|^2 = 0$) erroneously pushed the grid toward complete cooperation. By sweeping the parameter space, we identified that setting $|b|^2_{HDV} \approx 0.52$ forces the evolutionary simulation to stabilize perfectly at the empirically observed 42% cooperation rate. This confirms that human driving decisions are heavily influenced by latent social correlations, acting effectively as a natural entanglement mechanism that prevents absolute optimization.

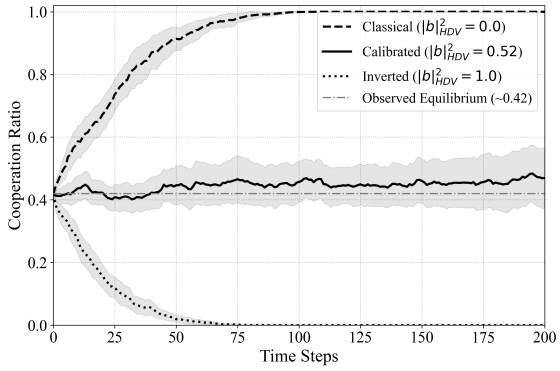


Figure 1: Evolution of cooperation in 100% HDV traffic under different entanglement regimes. The classical model (dashed line, $|b|^2_{HDV} = 0.0$) diverges from reality, while the calibrated quantum model (solid line) accurately reproduces the observed mixed equilibrium. The inverted model (dotted line) is also shown for reference.

Impact of AV Algorithms on Mixed Traffic

With the human baseline fixed at 0.52, we introduced AVs into the grid across varying Market Penetration Rates (MPRs). We tested three distinct AV programming profiles on a static spatial grid to observe how human drivers adapt:

- **Classical AVs ($|b|^2_{AV} = 0.0$):** Programmed as purely rational, independent agents.
- **Entangled AVs ($|b|^2_{AV} = 0.5$):** Utilizing V2V communication to act with high correlation.

- **Inverted AVs ($|b|^2_{AV} = 1.0$):** Operating under an inverted, highly defective reward structure.

Simulations revealed stark differences in human adaptation, as shown in Figure 2. When Classical AVs are deployed, they maintain high cooperation, but HDVs learn to exploit this predictability, actively reducing their own cooperation rates, especially near 50% MPR. Conversely, Inverted AVs, which inherently act defectively, compel the human population to become more cooperative to secure higher payoffs. Crucially, at low MPRs, deploying Inverted AVs paradoxically results in higher overall system cooperation than deploying Classical AVs. Finally, Entangled AVs provided the most stable integration; by operating near the human entanglement level, they maintained a consistent mixed equilibrium without triggering exploitative behaviors from the HDV population.

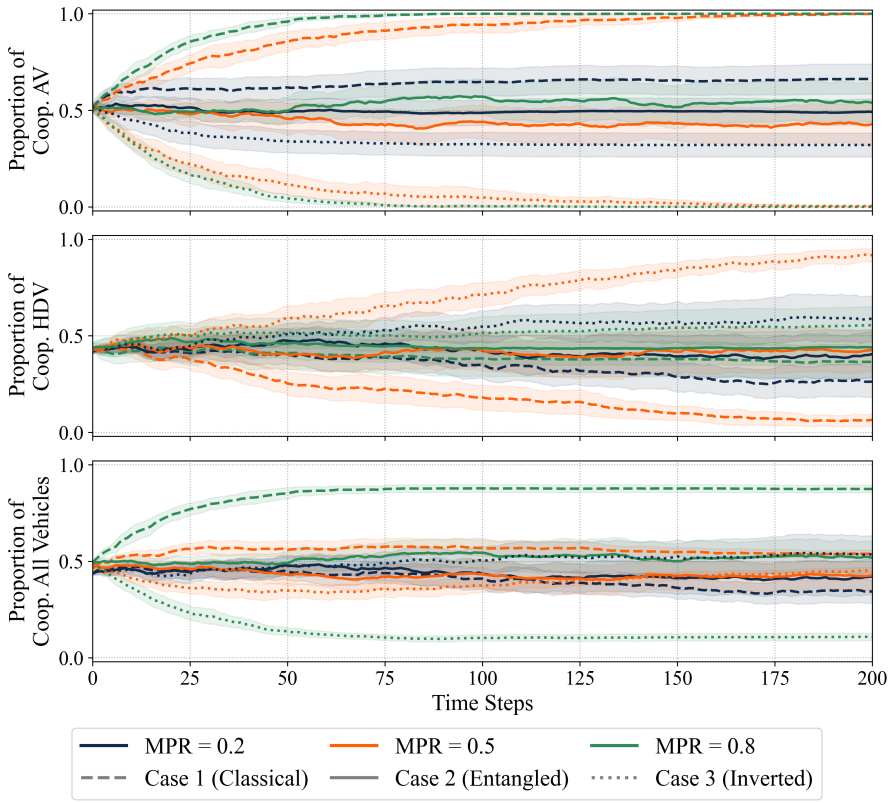


Figure 2: Comparison of evolutionary dynamics among classical AVs (dashed lines), entangled AVs (solid lines), and inverted AVs (dotted lines) on a static spatial grid. The rows represent the cooperation ratio of AVs, HDVs, and all vehicles, respectively. The shaded regions depict the 95% confidence intervals calculated from 20 simulation runs.

4 Conclusions

This study highlights the dangers of relying on classical independence assumptions when modeling human driving behavior. By integrating the Marinatto-Weber quantization scheme, we successfully embedded the latent social correlations of human drivers into an evolutionary traffic simulation. The findings clearly demonstrate that the evolution of human lane-changing behavior is not static; it will dynamically adapt based on the specific software architecture of the AVs sharing the road. Consequently, this quantized framework provides researchers and policymakers with a proactive tool to simulate and evaluate the long-term systemic impacts of various AV deployment strategies before they are implemented on public networks.

Acknowledgements

This work was supported by the National Science Foundation under Grant No. 2047937. The authors would also like to acknowledge the use of AI-assisted tools for language refinement during the preparation of this manuscript. All underlying research, data analysis, and methodological development remain original work.

References

An, G., Bae, J. H., & Talebpour, A. (2023). An optimized car-following behavior in response to a lane-changing vehicle: A bézier curve-based approach. *IEEE Open Journal of Intelligent Transportation Systems*, *4*, 682–689.

Chung, S., Ka, D., Kim, Y., & Lee, C. (2024). Gap setting control strategy for connected and automated vehicles in freeway lane-drop bottlenecks. *IET Intelligent Transport Systems*, *18*(12), 2641–2659.

Chung, S., Talebpour, A., & Hamdar, S. H. (2025). Characterizing lane-changing behavior in mixed traffic. *arXiv preprint arXiv:2512.07219*.

Eisert, J., & Wilkens, M. (2000). Quantum games. *Journal of Modern Optics*, *47*(14-15), 2543–2556.

Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., ... Anguelov, D. (2021, October). Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the ieee/cvf international conference on computer vision (iccv)* (pp. 9710–9719).

Iqbal, A., & Toor, A. H. (2001). Evolutionarily stable strategies in quantum games. *Physics Letters A*, *280*(5-6), 249–256.

Khan, F. S., Solmeyer, N., Balu, R., & Humble, T. S. (2018). Quantum games: a review of the history, current state, and interpretation. *arXiv preprint arXiv:1803.07919*.

Marinatto, L., & Weber, T. (2000). A quantum approach to static games of complete information. *Physics Letters A*, *272*(5-6), 291–303.