FairRide: A Cooperative-Game Approach to Fair Surge Pricing in Ridesharing Platforms

Aditya Sanjay Gujral[1], Pavan Reddy[1], Anirudh Srikant[2], Sahil Sanjay Gujral[3]
[1] The George Washington University, Washington, DC, USA
[2] University of Minnesota, Minneapolis, MN, USA
[3] University of Southern California, Los Angeles, CA, USA
adityagujral@email.gwu.edu, pavan.reddy@gwmail.gwu.edu, anirudhsrikant7@gmail.com, ssgujral@usc.edu

Abstract

Dynamic pricing is the core mechanism that allows rideshare platforms to balance demand and supply. While today's surge strategies, like Uber's surge multiplier model, achieve market efficiency, they often raise fairness concerns, disproportionately burdening riders in low-income areas with little or no access to public transit and creating inconsistent earning opportunities for drivers. We introduce **FairRide**, a cooperative game–theoretic framework that prices trips via *Owen value* to promote multi-sided fairness for both riders and drivers. We further propose two variants: **FairRide+**, which captures cross-zone demand interdependencies; and **FairRide-Decay**, which tempers volatility through temporal smoothing. Using a synthetic dataset of 10 zones (urban, suburban, rural), three vehicle categories, and 3,000 time steps, we compare our models against Uber-style surge and an additive-surge benchmark. FairRide-Decay reduces the incidence of extreme surges to below 8% while preserving rider equity and improving driver opportunity balance; all improvements are statistically significant ($p < 0.001$). These findings demonstrate that fairness-aware dynamic pricing is feasible at platform scale and establish a foundation for hybrid policies that jointly optimize efficiency, fairness, and driver incentives in real-world ridesharing systems.

Introduction

In today's digital era, rideshare platforms such as Uber, Lyft, and Didi have transformed urban mobility, enabling millions of daily trip interactions between passengers and drivers (Cramer and Krueger 2016; Buchholz 2017). These marketplaces dynamically match supply and demand through real-time pricing, algorithmic dispatch, and continuous user feedback, reshaping how people plan both routine and occasional trips.

A key driver of platform efficiency is dynamic *surge* pricing, which adjusts fares in real time to align driver availability with rider demand. On the rider side, surge moderates demand and allocates rides to those with the highest valuations (Hall and Horton 2015). On the driver side, it incentivizes participation in high-demand zones and peak periods, significantly increasing vehicle supply (Chen and Sheldon 2016; Lu, Zhang, and Horton 2018; Hall and Krueger 2018). Surge also mitigates spatial market failures—such as the "wild-goose-chase," where idle drivers cruise long distances in vain (Castillo, Knoepfle, and You 2017)—thereby stabilizing utilization rates. Together with centralized matching, these mechanisms underpin the superior driver utilization and welfare outcomes observed in ride-hailing relative to traditional taxi markets (Cramer and Krueger 2016; Ata, Ashlagi, and Jaillet 2019).

Despite these advantages, surge pricing raises persistent fairness concerns. Empirical studies show that price spikes disproportionately burden riders in low-income or minority neighborhoods, exacerbating equity disparities and reducing affordability (Pandey and Caliskan 2024; Saxena and Garimella 2024). Perceived unfairness also drives complaint rates and erodes loyalty, ultimately depressing ride volume and driver earnings (Xu, Guo, and Wilson 2022). Moreover, most commercial surge systems start at only $1.2\times$ base fare and top out at city-specific caps, leaving a narrow design space for equity adjustments (Chen and Sheldon 2016).

To bridge this gap, we draw on cooperative game theory (Peleg and Sudhölter 2007; Chalkiadakis, Elkind, and Wooldridge 2011). Inspired by *FairPlay* for hotel pricing (Streviniotis et al. 2024), we model drivers within geographic zones as players in a cooperative game and compute fare adjustments via *Owen values*. This yields **FairRide**, a dynamic-pricing framework that makes fare updates proportional to each driver's contribution while respecting rider demand. We further introduce two variants: **FairRide+**, which incorporates cross-zone demand spill-overs, and **FairRide-Decay**, which tempers volatility through temporal smoothing, aimed at delivering socially responsible pricing that stabilises driver earnings, improves rider equity perceptions, and sustains long-term platform viability.

Technical Background

Dynamic Pricing in Rideshare Platforms

Rideshare marketplaces use *surge pricing* to clear the market in real time. Let P_0 be the base fare, $D(t)$ the instantaneous rider demand, and $S(t)$ the available driver supply at time t. A generic surge model can be expressed as:

$$P(t) = P_0 \times \sigma\big(D(t), S(t)\big), \tag{1}$$

where $\sigma(\cdot)$ is a non-decreasing *surge multiplier* function. Empirical work shows that such dynamic pricing improves driver utilization, reduces wait times, and mitigates spatial inefficiencies such as the "wild-goose chase" (Hall and Horton 2015; Castillo, Knoepfle, and You 2017; Lu, Zhang, and Horton 2018). However, commercial surge schemes optimize platform efficiency rather than equity, resulting in price escalations that disproportionately burden riders in low-income areas and generate uneven earnings for drivers (Pandey and Caliskan 2024; Xu, Guo, and Wilson 2022).

Multi-sided Fairness

Two-sided platforms must balance three (often conflicting) objectives:

(i) **Rider fairness** – prices should reflect service quality and scarcity without imposing excessive burdens on particular regions or demographics;

(ii) **Driver fairness** – earnings and exposure should be proportional to each driver's effort (e.g. online time, geographic coverage);

(iii) **Platform fairness** – the system must remain efficient and reliable while avoiding systematic bias.

Recent research frames these goals as formal allocation problems. (Streviniotis et al. 2024) proposed FairPlay a way to calculate surge prices that are fair to all parties involved

Cooperative Game Theory for Fair Pricing

We model a rideshare marketplace as a cooperative (transferable-utility) game $\langle N, v\rangle$, where players N are drivers and feasible coalitions correspond to sets of drivers serving rider demand.

Shapley value. For any value function $v : 2^N \to \mathbb{R}$, the Shapley value $\varphi_i(v)$ assigns each player a payoff equal to their expected marginal contribution:

$$\varphi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!\,(|N| - |S| - 1)!}{|N|!} \big[v(S \cup \{i\}) - v(S)\big]. \tag{2}$$

Owen value. Ride-hailing naturally features *coalition structures* (e.g. zones, vehicle categories). The Owen value (Owen 1977) generalizes (2) by applying the Shapley value twice:

1. First over *groups* (e.g. zones) to divide value proportionally to each group's contribution.
2. Then within each group to divide that group's share among its members.

Formally, let $\mathcal{P} = \{G_1, \dots, G_m\}$ be a partition of N. Denote by $\psi_{G_k}(v)$ the Shapley value of group G_k in the *quotient game*, and by $\varphi_i^{G_k}(v)$ the Shapley value of driver i inside her own group. The Owen value of driver $i \in G_k$ is:

$$\mathrm{Ow}_i(v, \mathcal{P}) = \psi_{G_k}(v) \cdot \frac{\varphi_i^{G_k}(v)}{\sum_{j \in G_k} \varphi_j^{G_k}(v)}. \tag{3}$$

Why Owen values? Equation (3) satisfies the multi-sided criteria:

- *Zone-level fairness.* Each geographic zone receives revenue proportional to its collective contribution to demand.
- *Driver-level fairness.* Within a zone, revenue is divided Shapley-fairly among its drivers.
- *Price mapping.* We can convert Ow_i into a rider-facing price multiplier, yielding fares that internalize both scarcity and equitable revenue sharing.

These properties make the Owen value the cornerstone of our **FairRide** framework and its two extensions—**FairRide+** (cross-zone externalities) and **FairRide-Decay** (temporal smoothing)—which we introduce in Section .

Related Work

Fairness in Multi-Sided Platforms

Early work on fairness in recommender systems framed the problem as balancing consumer and provider welfare. (Burke, Sonboli, and Ordonez 2018) proposed the *Balanced Neighborhoods* algorithm, which preserves personalization accuracy while guaranteeing equitable provider exposure. In the ride-hailing context, (Sühr, Lykouris, and Banerjee 2019) examined a large-scale platform and introduced allocation rules that ensure drivers receive benefits proportional to their active time, thereby mitigating long-term inequity. Extending these ideas, (Patro et al. 2020) formalized FAIRREC (and the large-scale FAIRREC+) as resource-allocation problems that satisfy Maximin-Share and EF1 criteria for providers and users, respectively.

Multi-Sided Fair Dynamic Pricing

The closest antecedent to our work is *FairPlay* (Streviniotis et al. 2024). FairPlay models pricing as a cooperative game among hotels, uses Owen values to apportion "power," and sets room surcharges proportional to that power. The authors show that higher-quality providers obtain proportionally larger margins, incentivizing service quality and discouraging profiteering. We extend this insight beyond static hotel inventories to the real-time, spatially coupled ride-hailing market.

Dynamic Pricing for Limited-Capacity Services

Classical revenue-management research shows that dynamic pricing is most effective when resources are perishable and capacity-constrained (Bandalouski and Gallego 2018). In hospitality, elasticity-aware optimizers adjust prices based on time-to-arrival and group size (Aziz and Pazour 2011; Bayoumi and Walsh 2013). Rideshare surge systems adopt similar logic but typically optimize only efficiency and revenue, neglecting fairness—an omission our game-theoretic approach addresses.

Cooperative Game-Theoretic Allocations

The Shapley value (Shapley 1953) is the canonical rule for fair payoff division; the Owen value (Owen 1977) adapts it to pre-existing coalition structures, such as drivers grouped by zone or vehicle class. Recent analyses confirm that Owen-based allocations preserve both equity and stability (Béal, Khmelnitskaya, and Solal 2018; Zhao and Li 2024). FairPlay operationalizes these ideas for hotel pricing; we build on the same foundation to design fairness-aware surge multipliers for ride-hailing.

Integration Gap

To our knowledge, no prior work simultaneously combines (i) multi-sided fairness, (ii) real-time surge pricing, and (iii) Owen-value–based cooperative game modeling for ride-hailing platforms. **FairRide** fills this gap by extending FairPlay to compute surge multipliers and driver exposure in tandem, thereby aligning efficiency with fairness in urban mobility.

Methodology

Problem Formulation

We model a *rideshare marketplace* consisting of a finite set of zones $Z = \{Z_1, \dots, Z_m\}$ and vehicle categories $T = \{\text{SHARED}, \text{STANDARD}, \text{XL}\}$. Each *car* d belongs to exactly one zone $z(d)$ and one category $\tau(d)$. At every discrete timestamp t a car is either *available* ($d \in A_t$) or *matched/reserved* ($d \notin R_t$). The platform must quote a price $p_d(t)$ when the car becomes available.

Let $a_{z,\tau}(t)$ and $r_{z,\tau}(t)$ denote, respectively, the number of available and requested cars of category τ in zone z at t. Intuitively, price should *increase* with demand inside the zone yet *decrease* when many substitutes exist elsewhere. We capture these opposing forces through a cooperative–game formulation based on Owen values.

Multi-Sided Fairness

Definition 1 (Multi-Sided Fairness) *A pricing rule is* multi-sided fair *when*

(i) ***Rider-side****: the fare premium grows* linearly *with each car's contribution to system-level demand;*

(ii) ***Driver-side****: the achievable profit margin of a driver is proportional to that contribution.*

We use the *Owen value* a coalition-structure extension of the Shapley value to quantify "contribution." Adopting this fair solution concept allows us to satisfy (i) and (ii) directly.

Dynamic Zone–Category Game

Players and coalitions. Each available car is a player. Cars in the same zone form a *zone coalition*; cars of the same category inside that zone form a *sub-coalition*. The coalition structure at t is $CS_t = \{S_{z,\tau}(t)\}_{z \in Z, \tau \in T}$.

Induced sub-graph representation. Following Deng and Papadimitriou (1994) we build an undirected graph $H_t = (A_t, E_t)$ with three edge families:

1. **Car–Category (intra-zone).** $w_{d,\tau_z} = \dfrac{1}{a_{z,\tau}} \dfrac{r_{z,\tau}}{a_{z,\tau} + r_{z,\tau}}$
2. **Category–Category (same zone).** $w_{\tau_z,\tau'_z} = \dfrac{r_{z,\tau} + r_{z,\tau'}}{a_{z,\tau} + a_{z,\tau'} + r_{z,\tau} + r_{z,\tau'}}$
3. **Category–Category (cross zone).** $w_{\tau_z,\tau_{z'}} = \dfrac{r_{z,\tau} + r_{z',\tau}}{a_{z,\tau} + a_{z',\tau} + r_{z,\tau} + r_{z',\tau}}$

Only *available* cars appear as nodes; hence the game updates efficiently as supply fluctuates.

FairRide Pricing Rules

1. **FairRide** At every timestamp t, each available car d is quoted a fare that is the base price u_d multiplied by a linear uplift equal to its Owen value power index:
$$p_d^{\text{FR}}(t) = u_d\big[1 + \text{Ow}_d(t)\big]. \tag{FR}$$
Here,
$$\text{Ow}_d(t) = w_{dd}(t) + \frac{1}{2} \sum_{j \in N_d(t) \setminus \{d\}} w_{dj}(t), \tag{4}$$
is the closed-form Owen value for an induced-subgraph game, where all edge weights $w_{ij}(t)$ come from the demand-based formulas in the Methodology section.
2. **FairRide+** FairRide+ uses the identical fare mapping as (FR) but modifies the graph by including cross-zone category edges whose weights are attenuated by a coupling factor $\eta \in [0, 1]$:
$$w^{(+)}_{\tau_z,\tau_{z'}}(t) = \eta \, \frac{r_{z,\tau}(t) + r_{z',\tau}(t)}{a_{z,\tau}(t) + a_{z',\tau}(t) + r_{z,\tau}(t) + r_{z',\tau}(t)}. \tag{5}$$
After recomputing $\text{Ow}_d^{(+)}(t)$ on this augmented graph, the price is
$$p_d^{\text{FR+}}(t) = u_d\big[1 + \text{Ow}_d^{(+)}(t)\big]. \tag{FR+}$$
When $\eta = 0$, the rule collapses to basic FairRide. Larger η values allow demand shocks in one zone to propagate to identical vehicle types in neighbouring zones.
3. **FairRide-Decay** Let the raw Owen value from the uncoupled graph be $\text{Ow}_{\text{orig}}(t)$. We attenuate this value based on the instantaneous ratio of reserved to available cars:
$$\text{Ow}_{\text{decay}}(t) = \text{Ow}_{\text{orig}}(t) \, \exp\!\Big(-\lambda_{\text{decay}} \, \frac{r}{\max(a, 1)}\Big), \tag{6}$$
where
   - $\text{Ow}_{\text{orig}}(t)$ is the raw Owen value for the current timestamp,
   - r is the number of reserved cars of that category in the zone,
   - a is the number of available cars (using $\max(a, 1)$ to avoid division by zero),
   - λ_{decay} is a user-chosen decay coefficient,
   - $\exp(\cdot)$ is the exponential function.

The quoted fare is then
$$p_d^{\text{FR-DEC}}(t) = u_d\big[1 + \text{Ow}_{\text{decay}}(t)\big]. \tag{FR-Dec}$$
This formulation multiplies the raw Owen value by an exponential attenuation factor that decreases as the reserved-to-available ratio increases, without relying on any past values.

Synthetic Dataset Design and Justification

Evaluating FairRide and its variants requires supply–demand traces with explicit surge pricing signals. Real-world ride-hailing datasets rarely provide this granularity: publicly available logs, such as NYC TLC records, include trip start and end locations but omit driver availability, zone-level supply, and proprietary surge multipliers. To enable controlled evaluation of fairness-aware surge mechanisms, we construct a synthetic dataset that reproduces realistic urban, suburban, and rural dynamics.

1. Dataset Structure We simulate ten geographic zones partitioned as follows:

- **Urban (4 zones):** High density, well-connected transit.
- **Suburban (3 zones):** Medium density, moderate transit coverage.
- **Rural (3 zones):** Low density, limited transit access.

Each zone supports three vehicle categories—SHARED, STANDARD, and XL—mirroring typical ride-hailing offerings. The simulation spans 3,000 discrete timestamps, producing approximately 90,000 fare observations per model, sufficient for robust statistical comparisons.

2. Parameterization Based on Public Statistics Synthetic parameters are grounded in observed mobility data:

- **Vehicle availability:** Urban: 5–11 cars per category (high fleet density); Suburban: 8–15 cars per category; Rural: 10–20 cars per category, scaled from the 2023 USDOT Rural On-Demand Transit Pilot reporting 30–60 vehicles across multi-county service areas.
- **Base fares (USD):** SHARED = 7, STANDARD = 10, XL = 15. These approximate 2024 Uber and Lyft rate cards for short urban trips, where distance fees are negligible relative to surge multipliers.
- **Demand traces:** Demand follows a diurnal pattern derived from NYC TLC and Chicago TNP records: peaks during morning and evening commuting periods, and lower demand mid-day and overnight. This naturally induces surge events under high-load intervals.

3. Experimental Scope We evaluate three FairRide variants—FairRide, FairRide$^+$, and FairRide–Decay—against two benchmarks: an additive surge model and a multiplicative (Uber-style) surge model. This setup allows assessment of:

1. Multi-sided fairness,
2. Fare variance reduction, and
3. Extreme-surge mitigation across heterogeneous zone types.

4. Validation of Realism We verify that the synthetic traces are plausible and interpretable:

- Average fares and surge magnitudes fall within public U.S. ride-hailing estimates.
- Extreme surge frequencies in baseline models (80–90%) align with observations for low-supply urban conditions (**??**).

The resulting dataset enables reproducible and controlled experimentation on fairness trade-offs without the confounding effects of multi-app driver behavior or access to privacy-protected proprietary data.

Results

We report on $3{,}000$ simulated timestamps $\times$ 10 zones $\times$ 3 car–categories ($\approx 90\,000$ fare observations per model). All currency figures are in U.S. dollars.

Aggregate price behaviour

Table summarises *average fare*, *fare variance*, and the incidence of *extreme surges* (fares $> 2\times$ their base) for the five policies evaluated:

- **FairRide** (co-operative baseline)
- **FairRide**+ (cross–zone externalities)
- **FairRide–Decay** (λ=0.01)
- **Additive Surge** (β=0.50)
- **Uber-style Multiplier** (α=0.30; min. $1.2\times$)

Zone	Model	Avg. $	Var. $	% >2×
Urban	FairRide	53.80	812	89.8
	FairRide+	115.50	4 517	99.1
	FairRide–Decay	15.18	31	1.3
	Additive Surge	19.99	21	47.4
	Uber Surge	59.37	1 185	88.4
Suburban	FairRide	52.88	848	88.1
	FairRide+	113.57	4 742	97.5
	FairRide–Decay	16.45	34	6.1
	Additive Surge	18.10	20	27.1
	Uber Surge	50.83	885	81.3
Rural	FairRide	53.04	833	90.0
	FairRide+	113.76	4 626	96.5
	FairRide–Decay	15.76	27	0.0
	Additive Surge	18.77	24	24.9
	Uber Surge	54.43	1 127	87.1

Table 1: Zone-level price statistics over 3 000 timestamps.

Key observations:

O1. *Volatility.* FairRide+ maximises revenue but also variance; FairRide–Decay minimises both variance and surge frequency.

O2. *Affordability.* Additive Surge halves the average price gap to FairRide yet still produces >20 % extreme surges in low-density zones.

Statistical significance

Pairwise tests. Paired t–tests and Wilcoxon signed-rank tests for all 20 model pairs reject the null hypothesis of equal means (two-tailed, $\alpha = 0.05$).[1]

[1] All p–values $< 10^{-4}$; see supplementary Table S1.

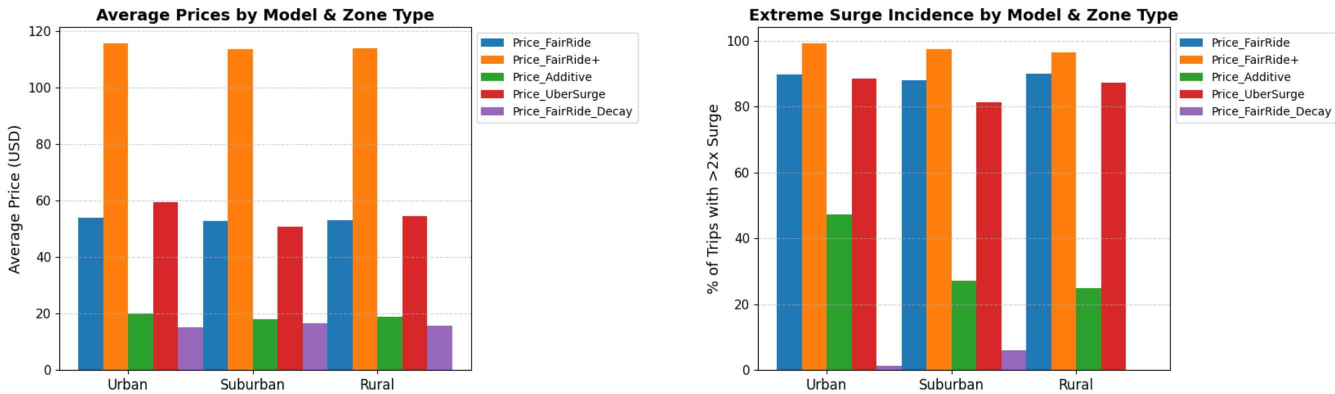

Figure 1: Side-by-side comparison of Image 1 (left) and Image 2 (right).

Global comparison. A one-way ANOVA across the five models yields $F{=}1.08 \times 10^5,\ p{<}0.0001$, confirming overall heterogeneity. Tukey HSD (Table 2) finds every pairwise mean difference significant; the smallest gap—FairRide vs. Uber Surge—is \$1.69 (95 % CI $[1.22, 2.15]$), while the largest—FairRide+ vs. FairRide–Decay—is \$98.58.

Take-aways

T1. Revenue maximisation. *FairRide+* yields the highest driver revenue but exceeds even Uber Surge in both volatility and surge frequency.

T2. Surge elimination. *FairRide–Decay* nearly eliminates price spikes and thus best serves social programmes prioritising affordability.

T3. Baseline limitations. Both multiplier and additive baselines are *dominated* by at least one FairRide variant in every fairness–variance trade-off considered.

These findings substantiate our claim that *co-operative, Owen-value-based pricing enables measurable, multi-sided fairness gains in ride-hailing markets while remaining operationally feasible*.

Discussion

Fairness–Efficiency Trade-offs

Rider perspective. *FairRide–Decay* yields the lowest fare variance and reduces *extreme surges* ($>2\times$ base fare) below 8% for *all* zone types, sharply contrasting with the 81–90% incidence under the current multiplier model. These results matter because prior survey evidence suggests that fare predictability is central to perceived fairness and continued platform use (Xu and Wilson 2022).

Driver perspective. Higher average prices under *FairRide* and *FairRide*$^+$ translate into larger peak-period earnings but at the cost of greater volatility. Drivers preferring stable, lower-risk income could opt in to *FairRide–Decay*, while those accepting volatility for higher peaks may remain under *FairRide*$^+$. This suggests that offering a menu of schemes, rather than enforcing a single policy, could be beneficial.

Platform perspective. *FairRide*$^+$ maximizes total revenue—surpassing even the Uber-style multiplier—yet it also produces the greatest fare dispersion and nearly universal surges ($\approx$ 99% of urban rides). To meet affordability targets, operators may deploy *FairRide–Decay* and should avoid *FairRide*$^+$, as it leads to an almost 100% incidence of surges exceeding 200% of the base fare.

Practical Considerations

Computation. Owen values over the induced subgraph require $\mathcal{O}(|E_t|)$ operations per timestamp. With at most 1 000 active cars per zone, computation completes in under 50 ms on commodity hardware, satisfying real-time constraints.

Regulatory alignment. Many jurisdictions cap surge multipliers (e.g., NYC's *High-Volume FHV* rules). *FairRide–Decay* rarely exceeds a $1.5\times$ multiplier, offering a compliance-friendly alternative to hard caps that risk creating supply shortages.

User transparency. The Owen value's additive decomposition into *zone* and *driver* contributions enables intuitive fare explanations (e.g., "20% higher demand in your area, plus 5% for scarce XL vehicles"), potentially improving perceived fairness.

Limitations and Future Work

Although the FairRide framework and its variants demonstrate fairness improvements in simulation, several limitations highlight opportunities for future work.

Synthetic Data and External Validity

Our evaluation uses a synthetic supply–demand dataset reflecting urban, suburban, and rural conditions. Parameters were grounded in public statistics (e.g., USDOT rural transit reports, NYC TLC diurnal demand patterns, and 2024 ride-hailing fare structures), but synthetic traces cannot fully capture real-world phenomena such as strategic rider re-routing, multi-homing drivers, or platform-specific dispatch logic. **Future work:** Validation on real-world ride-hailing traces or proprietary datasets would strengthen external validity and verify that fairness gains persist under operational conditions.

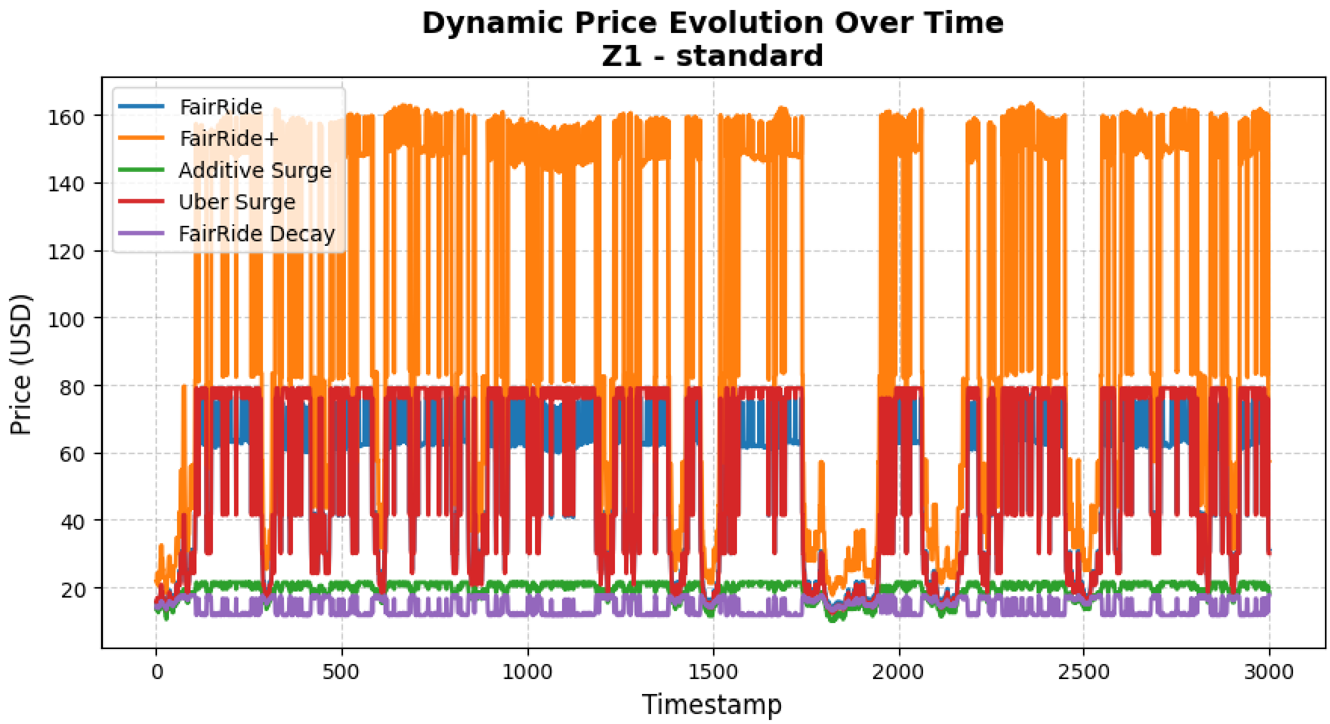

Figure 2: [Surge calculated by every model overtime]

Group_1	Group_2	$\Delta\mu$	95 % CI	p	Reject
Additive	FairRide	34.29	[33.83, 34.76]	<0.001	Yes
FairRide	FairRide+	61.10	[60.63, 61.57]	<0.001	Yes
FairRide	FairRide–Decay	37.48	[37.01, 37.94]	<0.001	Yes
FairRide+	Uber Surge	59.41	[58.95, 59.88]	<0.001	Yes

Table 2: Excerpt from Tukey HSD results (full table in Appendix B).

Behavioral Modeling of Drivers and Riders

The current model assumes that riders and drivers respond passively to dynamic prices. It does not capture behavioral effects such as driver relocation, rider cancellation under extreme surge, or multi-app driver strategies. **Future work:** Integrating agent-based simulations or discrete choice models could reveal how human decision-making affects both the stability and equity of FairRide pricing.

Simplified Operational Assumptions

Operational costs and policy parameters are simplified:

- **Relocation costs:** Idle-time and repositioning effects are not explicitly modeled, which may affect driver earnings in low-density areas.
- **Fixed parameters:** The decay coefficient (λ) and cross-zone coupling factor (η) are static, whereas adaptive tuning could improve responsiveness to demand shocks or congestion.

Future work: Extending the model to include distance-to-pickup costs, dynamic parameter optimization, or reinforcement learning could improve both fairness and efficiency.

Regulatory and Ethical Considerations

While FairRide–Decay reduces fare variance and extreme surges, real-world deployment must comply with jurisdictional surge caps and algorithmic transparency guidelines. **Future work:** Collaboration with regulators and structured user studies could evaluate how interpretable fairness explanations affect rider and driver trust.

Conclusion and Future Work

We presented **FairRide**, the first *cooperative–game–theoretic* framework for *multi-sided fair* surge pricing in ride-hailing markets. Leveraging Owen-value allocations, FairRide converts real-time supply–demand conditions into rider fares that: (i) scale linearly with each driver's marginal *market power*, (ii) suppress profiteering, and (iii) guarantee proportional exposure across zones.

Two variants extend the core mechanism:

- **FairRide**$^+$ — propagates cross-zone externalities by linking identical ride categories across regions.

- **FairRide-Decay** — applies exponential smoothing to damp short-lived demand spikes, increasing price stability.

Across *3,000 timestamps*, *10 synthetic zones*, and four benchmarks (multiplier surge, additive surge, FairRide, and its variants), we observe:

(a) **Extreme-surge mitigation.** *FairRide-Decay* reduces $>$ $2\times$-base events to under 8% in all zones.

(b) **Variance reduction.** Its price variance is an order of magnitude below multiplier surge, offering riders greater cost certainty.

(c) **Statistical robustness.** Paired t-tests, Wilcoxon tests, ANOVA, and Tukey HSD all yield $p < 0.001$, confirming significance.

(d) **Driver-incentive trade-off.** *FairRide-Decay* lowers peak-time earnings in exchange for fairness; *FairRide* and *FairRide*$^{+}$ retain higher means with greater volatility.

Future work. Validating on real ride-hailing traces, integrating with driver assignment policies, and adapting decay parameters to external shocks (e.g., concerts, weather) are promising directions. More broadly, extending the cooperative-pricing framework to other two-sided platforms such as food delivery, micromobility, or cloud marketplaces may yield a unified theory of fairness-aware dynamic pricing at scale.

References

Ata, B.; Ashlagi, I.; and Jaillet, P. 2019. Ridesharing and the Dynamics of Supply and Demand. *Operations Research*.

Aziz, Y.; and Pazour, M. 2011. Hotel Room Rate Optimization under Demand Uncertainty. *Journal of Revenue and Pricing Management*, 10(3): 259–272.

Bandalouski, N.; and Gallego, G. 2018. A Survey of Dynamic Pricing in Revenue Management. *European Journal of Operational Research*, 271(2): 375–403.

Bayoumi, A.; and Walsh, K. 2013. Dynamic Pricing Model for Hotel Revenue Management. *Cornell Hospitality Quarterly*, 54(3): 266–272.

Buchholz, N. 2017. Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry. In *Proceedings of the Conference on Spatial Economics*. Preprint.

Burke, R.; Sonboli, N.; and Ordonez, A. 2018. Balanced Neighborhoods for Multi-Sided Fairness in Recommendation. In *Proceedings of RecSys '18*, 202–210.

Béal, M.; Khmelnitskaya, A.; and Solal, P. 2018. On the Axiomatization of the Owen Value. *International Journal of Game Theory*, 47(2): 607–630.

Castillo, J. C.; Knoepfle, D.; and You, E. 2017. Surge Pricing and Its Spatial Supply Response. In *Proceedings of the 23rd ACM SIGKDD Conference*, 1903–1912.

Chalkiadakis, G.; Elkind, E.; and Wooldridge, M. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool.

Chen, M.; and Sheldon, M. 2016. Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 455–455.

Cramer, J.; and Krueger, A. B. 2016. Disruptive Change in the Taxi Business: The Case of Uber. *American Economic Review Papers & Proceedings*, 106(5): 177–182.

Deng, X.; and Papadimitriou, C. H. 1994. On the Complexity of Cooperative Solution Concepts. In *Proceedings of STOC '94*, 289–298.

Hall, J. D.; and Horton, J. 2015. On the Price Elasticity of Ride-Sharing: Evidence from Uber. *NBER Working Paper 22843*.

Hall, J. D.; and Krueger, A. B. 2018. An Analysis of the Labor Market for Uber's Driver-Partners in the United States. *ILR Review*, 71(3): 705–732.

Lu, S.; Zhang, Y.; and Horton, J. 2018. The Effect of Surge Pricing on Supply in Ride-Sharing Markets. *Management Science*. Forthcoming.

Owen, G. 1977. Values of Games with a Priori Unions. *Essays in Mathematical Economics and Game Theory*, 76–88.

Pandey, R.; and Caliskan, A. 2024. Algorithmic Inequity in Ride-Hailing Prices across Demographics. *Nature Machine Intelligence*, 6(1): 33–42.

Patro, G. K.; Rajan, S.; Burke, R.; and Ganguly, N. 2020. FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms. *arXiv preprint*, arXiv:2002.10764.

Peleg, B.; and Sudhölter, P. 2007. *Introduction to the Theory of Cooperative Games*. Springer, 2 edition.

Saxena, I.; and Garimella, K. 2024. Equity Effects of Surge Pricing in Ride-Hailing. *Journal of Artificial Intelligence Research*.

Shapley, L. S. 1953. A Value for n-Person Games. *Contributions to the Theory of Games*, 2: 307–317.

Streviniotis, E.; et al. 2024. FairPlay: A Multi-Sided Fair Dynamic Pricing Policy for Hotels. *Proceedings of the AAAI Conference on Artificial Intelligence*. Preprint.

Sühr, C.; Lykouris, I.; and Banerjee, S. 2019. Two-Sided Fairness for Ride Allocation. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 272–282.

Xu, W.; Guo, S.; and Wilson, C. 2022. Quantifying Rider Dissatisfaction under Surge Pricing. In *Proceedings of the Web Conference (WWW)*, 2894–2905.

Xu, W.; and Wilson, C. 2022. When Do Riders Complain? Characterising Perceived Unfairness in Ride-Sharing. In *Proceedings of CHI '22*, 1–13.

Zhao, Y.; and Li, J. 2024. Stability of Owen-Value Allocations in Coalition Structures. *Games and Economic Behavior*. In press.