Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Who Pays for Scarcity? Pay Floors, Surge Pricing and Waiting Times in New York City Ride-Hailing

Patrick Nyenkan

Patrick.nyenkan@ttu.edu

Texas Tech University

Research Article

Keywords: surge pricing, minimum pay regulation, pay floor, pass-through, event study, ride-hailing

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11722750/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Who Pays for Scarcity? Pay Floors, Surge Pricing and Waiting Times in New York City Ride-Hailing

Patrick Nyenkan
Department of Agricultural and Applied Economics, Texas Tech University
November 2026

Abstract

Dynamic pricing transfers scarcity rents, but to whom? New York City regulates a per-trip minimum for ride-hail driver pay, so the answer depends on where trips sit relative to that minimum. Using 239,470,448 completed high-volume for-hire trips in 2024, we document four facts. First, in a restricted sample covering about 76% of trips, mostly shorter and cheaper ones, recorded pay clusters tightly around a benchmark reconstructed from the TLC's published rates: in April, 63.9% of Uber and 40.0% of Lyft trips are paid within $0.05 of it, and 64.2% and 67.7% within $0.10. The platform ranking reflects the tolerance, because Lyft pays on a schedule about 0.4% above the TLC rates; these are proximity measures, not verified legal binding or compliance. Second, recorded driver pay moves with the base fare far less than one for one: the conditional within-cell slope is about 0.26 for Uber and 0.20 for Lyft. Third, a censored regression of pay on the floor and the fare puts the floor near-binding on about 60% of sampled Uber trips and, using Lyft's own estimated pay schedule, on about 77% of sampled Lyft trips, with roughly half of each additional fare dollar reaching the driver above the floor. Fourth, around 93 NBA game endings, the arena-core-minus-outer gap rises by $7.48 to $18.41 in fares, $6.72 to $10.37 in recorded driver pay and 2.2 to 2.9 minutes in pickup waits relative to two hours earlier, with larger changes for Uber; but 47% to 67% of the fare change is already present before the ending, the step from the quarter hour before to the quarter hour after is $3.87 to $8.47 in fares, and time-of-day drift is not ruled out, so we read these as co-movement around a predictable pulse, not as causal effects. A stylized floor-plus-share model shows how a pay floor can make driver pay flat in the fare for some trips. The public records do not reveal surge multipliers, a validated trip-level legal floor or platform net revenue. A matched non-event-day comparison and a sample-flow audit are specified in Appendix A and are not yet reported.

JEL classification: L11, L51, J30, R41, D43

Keywords: surge pricing; minimum pay regulation; pay floor; pass-through; event study; ride-hailing

1. Introduction

When scarcity raises the rider's price, who gets the extra money, and who bears the cost of the wait? The question matters for three live policy debates: whether surge prices should be capped, whether driver pay standards raise or compress earnings, and whether platforms use algorithmic pay to widen their margin when demand is tight. Existing work answers pieces of it. Castillo (2025)

quantifies the welfare effects of surge pricing in a structural model of a single platform. Hall, Horton and Knoepfle (2023) show that fare changes re-equilibrate waits and driver earnings. Rosaia (2025) studies competition between platforms in New York using proprietary data. Binns et al. (2025) and Liu et al. (n.d.) examine driver pay and tipping under algorithmic pricing, from driver-donated United Kingdom records and from New York trip records respectively.

We contribute descriptive evidence, from public trip records, on how a regulated pay floor shapes the sharing of price variation. We do not estimate a structural model, a platform take rate, rider incidence or a causal effect of surge pricing. The paper reports four sets of estimates from 2024 data: the share of trips whose recorded pay lies near a reconstructed pay benchmark, the within-cell association between driver pay and fare, a censored regression of pay on the floor and the fare, and an event study of NBA game endings. Quantities that need further estimation (a trip-level kink test with cell fixed effects, a take-rate decomposition and rider incidence) are set out as extensions in Section 8 and are not reported. Four findings follow.

1. **Many trips are paid near a pay-floor benchmark, and how many depends on the tolerance.** In the restricted sample, April proximity to a benchmark reconstructed from the TLC's published 2024 rates is 63.9% of Uber and 40.0% of Lyft trips within \$0.05, and 64.2% and 67.7% within \$0.10; full-year shares within \$0.05 are 56.5% and 34.0% and within 2% are 60.1% and 77.4%. These are proximity shares, not estimates of the fraction of trips legally constrained by the pay standard, and the sample covers about 76% of trips, mostly shorter and cheaper ones (Section 4). The platform ranking at \$0.05 reflects Lyft's pay schedule, which sits about 0.4% above the TLC rates (Section 6.1).
2. **Pay moves with the fare far less than one for one.** The conditional slope of recorded driver pay on base fare, within pickup zone and 15-minute cell, is about 0.26 for Uber and 0.20 for Lyft. A stylized floor-plus-share model (Section 3) shows how a floor can produce this pattern, but neither the proximity mass nor the average slope identifies a kink or the payment rule in the platforms' algorithms.
3. **A censored regression puts the floor near-binding on most sampled trips.** In random samples of March to December trips, the floor is near-binding on about 60% of Uber trips and, using Lyft's own estimated pay schedule, on about 77% of Lyft trips. Above the floor, roughly half of each additional fare dollar reaches the driver (latent slope 0.51 to 0.53 for Uber and 0.56 to 0.60 for Lyft). The implied overall slopes, about 0.20 to 0.21 for Uber and 0.13 for Lyft, are below the fixed-effects slopes, so the model does not reproduce them exactly. These are model-based shares, not legal binding shares, and their pickup-date-clustered standard errors are small for Uber (about 0.01 on the latent slope) and larger for Lyft (about 0.03) (Section 6.3).
4. **Event-time evidence on scarcity pricing.** Around 93 NBA game endings, fares, recorded driver pay and pickup waits in the arena zone all rise relative to nearby zones, and 47% to 67% of the bin +1 fare change is already present at bin −1. Because of pretrends and time-of-day drift, we report co-movement, not causal effects.

We do not claim to be the first paper on surge pricing or on New York take rates. The contribution is the joint description of the floor, the pay-on-fare slope and event-time adjustment, with public data and with the limits stated in Section 9. Appendix A specifies the audits and replication procedures that would validate the benchmark and the event design; they are not yet part of the reported evidence. Section 2 describes the regulation, Section 3 the framework, Section 4 the data, Section 5 the empirical strategy, Section 6 the results, Section 7 robustness, Section 8 extensions and policy implications, and Section 9 the limitations.

1.1 Related literature

Surge pricing, market clearing and driver supply. Castillo (2025) estimates the welfare effects of surge pricing from Uber data for Houston and reports a total gain of 2.15% of gross revenue, with riders gaining 3.57%, drivers losing 0.98% and the platform losing 0.50%. Hall, Horton and Knoepfle (2023) show that after Uber-initiated fare increases drivers' hourly earnings return to about their earlier level within roughly two months, as more hours and fewer trips per hour offset the higher fare, and that shorter waits only partly compensate passengers; two of the authors are Uber employees, as the paper discloses. Castillo, Knoepfle and Weyl (2025) study matching and pricing in ride-hailing, and Rosaia (2025) studies competition between platforms in New York. Cohen et al. (2016) estimate consumer surplus from Uber data. On the supply side, Cachon, Daniels and Lobel (2017) analyze surge pricing on a platform with self-scheduling capacity, and Garg and Nazerzadeh (2022) study driver-side surge pricing. Chen and Sheldon (2015), a working paper whose first author heads economic research at Uber, find that drivers work longer when surge prices are high, with hours elasticities of about 0.15 to 0.17 in ordinary least squares specifications and about 0.50 with instruments, and that a discontinuity design supports a causal reading. Lu, Frazier and Kislev (2018), also Uber-affiliated, exploit an outage of the surge heatmap for a subset of drivers in ten cities and find that seeing the map explains 10% to 60% of drivers' self-positioning and raises their revenue on surged trips by up to 70%. Yan et al. (2020), all of whom worked at Uber, review matching and dynamic pricing and show that jointly choosing prices and rider waiting can reduce price variability while raising utilization and throughput; they note that demand spikes at the end of a concert raise waits and lower utilization, the setting our event study uses. We do not estimate welfare, supply or equilibrium effects; we describe how a regulated pay floor and a predictable demand pulse show up in fares, pay and waits.

Driver pay regulation in New York. Parrott and Reich (2018) designed the earnings standard that the Taxi and Limousine Commission (TLC) adopted. Koustas, Parrott and Reich (2020), in a report, find that 96% of trips met the standard in February 2019 and 99% in the rest of 2019, that the company commission fell from 15% in June 2018 to 12.5% in June 2019, and that average waits fell from 5.82 to 4.79 minutes over the same months, with effects bundled with the vehicle cap and the congestion charge. A UCLA Labor Center brief developed for the New York Taxi Workers Alliance (Konishi et al., 2023) reports the companies' share of passenger fares rising from 9.0% in February 2019 to 20.7% in April 2022, and the share of rides on which the company took 30% or more of the fare rising from about 9% to 29%. Angrist, Caldwell and Hall (2021) provide

a benchmark for driver labor supply, with an elasticity of about 1.2 and about 1.8 for drivers on lease contracts. We add the share of 2024 trips paid at or near the pay benchmark, from public data on every trip.

Take rates, algorithmic pay and tips. Binns et al. (2025) analyze 1.5 million trips donated by 258 United Kingdom drivers around Uber's move to dynamic pricing. Hourly pay fell, pay became less predictable and more unequal across drivers, and the company's surplus per on-trip hour rose by 38%, although the mean driver share of the fare stayed near 75% while its median fell to 71%. For New York, Liu et al. (n.d.) use 2.0 million Uber trips from Manhattan south of 110th Street (April to October 2023, peak hours) and find that a 1% increase in a measure of pricing aggressiveness, the ratio of the posted fare to a meter-based fare, is associated with a 0.028% decrease in the tip ratio, which averages 4.7%; they attribute the decline to the opacity of the algorithm. Zhang, Sang and Wu (2026) study the effects of the 2025 congestion charge on ride-hailing and transit in New York and find that base fares and platform revenue per trip rose, especially for Uber, while driver pay gains were limited. Because fare and pay are net of different items in the public records, and because the records exclude tips, we do not measure the take rate and describe the fare-minus-pay gap as an accounting wedge.

Fairness, perceptions and regulation of algorithmic pricing. Dolata and Schwabe (2024) analyze the public discourse, 1,205 Reddit comments, 838 tweets and news reports, after the five-fold surge that followed the Brooklyn subway shooting of April 12, 2022, and show that algorithmic prices shape the notions of fairness people apply. Saxena, Zhang and Shahabi (2025) document disparate impact of ride-hailing prices on disadvantaged neighborhoods in Chicago and propose fair pricing mechanisms. Han, Gao and Deng (2018) discuss algorithmic price discrimination and antitrust enforcement in China's digital economy. Leccese (2024), in an unpublished working paper on a ride-sharing tax in Chicago, finds that price increases differed across community areas in line with their access to public transit and fell more heavily on areas with larger minority populations. Broader reviews of algorithmic pricing (Chirita and Chirita, 2024; Grace and Okoh, 2025) stress both efficiency gains and risks to fairness, transparency and competition, and Raza et al. (2026) note that open trip-level data are scarce in India, where surge is capped at twice the base fare. New York's public records are an exception, and our descriptive exercise depends on them.

2. Institutional Background

The TLC adopted its driver pay standard on December 4, 2018, with rates applying from January 1, 2019 and compliance visible from February 2019 (NYC TLC, 2018; Koustas et al., 2020). Minimum pay per trip equals the per-mile rate times trip miles plus the per-minute rate times trip minutes, each divided by the company's utilization rate, with no fixed component. For trips that end outside the city, the rule assigns a 50% utilization rate to all companies for the time and miles driven outside the city, which raises the applicable rates (NYC TLC, 2018). After the utilization adjustment, Konishi et al. (2023) report $1.088 per mile and $0.495 per minute in February 2019,

$1.103 and $0.502 in February 2020, and $1.161 and $0.529 in March 2022. **Our empirical work uses a benchmark built from the TLC's published 2024 rates (Section 4).** Prior evidence on compliance and company shares (Section 1.1) is descriptive, and we treat it as motivation, not as identification.

2024 rates. The TLC adjusts the per-mile and per-minute rates each year with the CPI-W. TLC Industry Notice 24-02 (February 14, 2024) set a 3.49% increase effective March 1, 2024 for all trips dispatched by high-volume for-hire services, which in 2024 were Lyft and Uber. Including the utilization rate, the minimum rates for non-wheelchair-accessible trips became $1.360 per mile and $0.583 per minute, and the out-of-town rates $1.576 per mile and $0.675 per minute. For a 7.5-mile, 30-minute trip the in-city rates imply a minimum of about $27.69 (7.5 × $1.360 + 30 × $0.583).

Upfront pricing. Since platforms moved to upfront rider prices, the rider fare and driver pay are set by separate algorithms and the old linear commission no longer describes the pay rule; Binns et al. (2025) document the same change for the United Kingdom. This is why Section 3 allows the platform's share to vary with the scarcity state, and why a constant-share assumption is a maintained benchmark, not a premise.

Sample year. New York's central business district tolling program, which Zhang, Sang and Wu (2026) study, began on January 5, 2025. The 2024 sample therefore precedes the charge and is not affected by it.

3. Model

The model is deliberately small and stylized. Its job is to show how a pay floor can make recorded driver pay flat in the fare for some trips, and to fix the notation for the accounting relationships in Propositions 2 and 3. It is not an estimated payment algorithm and does not replace the structural work of Castillo (2025) or Rosaia (2025).

3.1 Primitives

A trip has miles m and minutes t. The platform posts a rider fare p. The state of scarcity is θ, which raises demand relative to supply. A rider's generalized price is

$$G = p + v \cdot w$$

where w is the wait and v the value of time. Driver pay is

$$y = max\{ F(m,t), s(\theta) \cdot p \}$$

Here F(m,t) is a stylized floor schedule and s(θ) is a hypothetical fare-linked pay fraction. The expression is not an estimated Uber or Lyft payment algorithm. The TLC benchmark in Section 4 uses published rates already adjusted for utilization; dividing those adjusted rates by utilization again would double-count that adjustment. Recorded base passenger fare and recorded driver pay exclude different items, so the model's p − y is not a measured platform margin.

3.2 Where the floor binds

The floor binds when $s \cdot p < F$, that is, when the fare is low relative to the trip's miles and minutes. It does not bind when the fare is high, so the floor is a feature of low-price trips, not of surge. At the floor the modeled gap $p - y$ rises one-for-one with the fare, and the modeled take rate τ rises with p ($\tau = 1 - F/p$, so $d\tau/dp = F/p^2$); neither is a measured platform margin (Section 4). Above the kink τ is flat if s is constant and rises or falls with p according to how $s(\theta)$ responds.

Proposition 1 (kinked pass-through). Let $p^* = F / s$. Driver pay is flat in p for $p < p^*$ and has slope s for $p > p^*$. The take rate is increasing and concave in p below p^* and equal to $1 - s$ above it when s is constant.

Derivation. Below p^*, $y = F$, so $dy/dp = 0$ and $\tau = 1 - F/p$, which has first derivative $F/p^2 > 0$ and second derivative $-2F/p^3 < 0$. Above p^*, $y = s \cdot p$, so $dy/dp = s$ and $\tau = 1 - s$.

Proposition 2 (conditional comparative static). In the hypothetical non-floor regime, $\tau = 1 - s(\theta)$, so a decline in s as scarcity increases would raise the modeled take rate. This is a conditional prediction only: the public trip records do not identify $s(\theta)$, a platform take rate, or the causal effect of scarcity.

Proposition 3 (model accounting, not identified incidence). For a hypothetical scarcity change $d\theta$, the rider's generalized price changes by $dG = dp + v \cdot dw$. Within the simplified model, dy is driver pay per trip and $dp - dy$ is the residual of the stylized posted fare. The observed TLC fare-minus-pay difference is not platform net revenue, and neither welfare incidence nor causal pass-through is estimated.

$$I_rider = dG/dp, \quad I_driver = dy/dp, \quad I_platform = 1 - dy/dp$$

The model identities are conditional accounting relationships. The waiting-time component of generalized rider cost is not a financial transfer, but it is not the only welfare-relevant adjustment: driver time, trip selection, matching and quantity responses also matter.

3.3 Supply and what we do not identify

The supply side is hourly earnings and utilization, not per-trip pay (Hall et al., 2023). Public high-volume for-hire vehicle (HVFHV) records have no driver identifiers, so we do not estimate a driver-supply elasticity. We benchmark against Angrist, Caldwell and Hall (2021): about 1.2, and about 1.8 for drivers on lease contracts. Chen and Sheldon (2015) report lower hours elasticities, about 0.15 to 0.50 depending on specification, and Lu, Frazier and Kislev (2018) show that drivers' spatial response to surge depends on seeing it, so the supply response to a pulse is both uncertain in size and dependent on information. Any statement about driver welfare is therefore a statement about per-trip pay, not about hourly earnings.

What the evidence can test. A concentration of recorded pay near the reconstructed benchmark and a low conditional fare–pay slope are consistent with, but do not uniquely identify, the stylized floor-plus-share model. Neither a legally binding floor share nor an above-kink pay fraction is estimated.

4. Data and Measurement

The core data are the NYC TLC HVFHV trip records for January to December 2024, twelve monthly Parquet files, plus the TLC taxi zone lookup (NYC TLC, 2024b). The 12 files hold 239,470,448 trips, all dispatched by Uber (HV0003) or Lyft (HV0005); Juno and Via do not appear in 2024. Field definitions follow the TLC data dictionary dated March 18, 2025 (NYC TLC, 2025). The variable base_passenger_fare is before tolls, tips, taxes and fees; driver_pay excludes tolls and tips and is net of commission, surcharges and taxes. The variable on_scene_datetime applies to accessible vehicles only and is not used.

Table 1. Variable construction

Variable	Construction	Note
Fare	base_passenger_fare	Before tolls, tips, taxes and fees
Driver pay	driver_pay	Excludes tolls and tips; net of commission, surcharges and taxes. Fare and pay are net of different items, so fare minus pay is an accounting gap, not a take rate
Wait	pickup_datetime minus request_datetime, in minutes	Includes matching and dispatch delay; negative waits are flagged as a data-quality issue and left out of mean waits
Pay floor (regulatory-rate benchmark)	Trip miles times a per-mile parameter plus trip minutes times a per-minute parameter: $1.314 and $0.564 in January and February, $1.360 and $0.583 in March to December. Monthly averages of $14.30 to $16.21 for Uber and $13.88 to $16.26 for Lyft	March to December values match TLC Industry Notice 24-02; January and February values are derived by removing the 3.49% increase and are reproduced by the Uber pay envelope (Table 2c) but not checked against the earlier notice; a comparison measure, not a calculation of each trip's legally required pay
At-floor flag	driver_pay within 5 cents, 10 cents, 25 cents or 2% of the benchmark	Sensitive to the tolerance, so several are reported (Table 2b)

Sample. Uber and Lyft; non-shared, non-wheelchair-accessible, non-Access-A-Ride matched trips; positive fare, pay, miles and time (an estimated passenger-price index is also trimmed to [0.2, 10] but is not otherwise used here). About 1% of trips (2,316,353 of 239,470,448) have a pickup before the request. They are flagged as a data-quality exclusion and left out of mean waits: 0.5% to 1.0% of Uber trips and 1.5% to 2.1% of Lyft trips in each month, roughly 173,000 to 251,000 trips per month. Mean wait among valid trips ranges from 4.2 to 5.3 minutes for Uber and from 4.6 to 6.0 for Lyft across months.

What the data cannot do. Only completed trips are observed, so cancellations and unmet demand are missing and rider elasticities are bounded, not point-identified. There are no driver identifiers. No surge multiplier is reported, and tips are excluded from driver pay, which matters because tips respond to pricing aggressiveness (Liu et al., n.d.).

Event panel. Trip records were aggregated into 16,022,633 unique platform × pickup-zone × 15-minute cells, holding 239,469,582 completed trips, which is 866 fewer than the 239,470,448 in the monthly files (504 Uber and 362 Lyft); the difference has not yet been reconciled. A total of 4,203

repeated aggregation keys arose at monthly boundaries and were combined using count-weighted aggregation. A few cells fall outside calendar 2024 (the first interval begins on 31 December 2023 at 23:00 and the last on 1 January 2025), which we attribute to trips whose pickup time lies outside their file month; they are retained as recorded. Repeated cell keys do not establish duplicate trip records. The study examines 93 NBA game endings (53 Madison Square Garden, 40 Barclays Center) over 21 event-time bins (−8 through +12); bin 0 contains the ending. The event panel has 36,246 rows and 39 missing outcome cells. Core and outer zones are both potentially affected by the event; the outer zones are spatial comparators, not untreated controls. Twenty-two games in 11 overlapping pairs are flagged, leaving 71 non-overlapping games.

Floor sample and selection. The floor summary contains approximately 136.1 million Uber and 46.1 million Lyft trips (about 76% of the original platform totals), whereas full-year fixed-effects regressions use 135,451,811 and 44,887,925 observations. The stated restrictions are non-shared, non-wheelchair-accessible and non-Access-A-Ride matched service, positive fare/pay/distance/duration, and an estimated passenger-price-index trim to [0.2, 10]. This subsample is not representative. In April, mean driver pay in the floor sample is $16.34 for Uber and $16.51 for Lyft, against $19.68 and $20.17 in the full monthly file, which implies that the excluded trips average roughly $30 and $33 in pay. Floor shares therefore describe shorter and cheaper trips and should not be read as shares of all trips. A sequential count of losses by restriction and an included-versus-excluded comparison are not available, so neither the composition nor the selection bias from the excluded approximately 24% can be assessed. Because the price-index trim is built from the fare, the slope sample is also selected on the regressor of interest, and a slope estimated without the trim has not been run. No excluded observation should be described as unlawful or noncompliant on this basis.

Pay-floor measures. We distinguish two measures. The empirical envelope is a second-percentile quantile regression of driver pay on trip miles and minutes without an intercept, estimated by platform and month on random samples of 200,000 trips, excluding drop-offs at Newark Airport and in unrecorded zones: Q(0.02) of pay given miles and minutes = a(p,m) · miles + b(p,m) · minutes, where p is the platform and m the month. The regulatory-rate benchmark uses the per-mile and per-minute parameters in Table 1. The March to December values are the in-city rates for non-wheelchair-accessible trips in TLC Industry Notice 24-02 (NYC TLC, 2024a), which include the utilization rate. The January and February values are those rates before the 3.49% increase: dividing by 1.0349 gives $1.314 and $0.563, against the $0.564 used for the per-minute rate, a difference of a tenth of a cent. They are derived from the March rates and not taken from the notice in force at the time. The Uber pay envelope reproduces them to within 0.1% (Table 2c), which supports them, although they should still be checked against the notice. The benchmark applies the in-city rates, whereas trips ending outside the city carry higher statutory rates, so such trips, if included, would sit above it. It is a comparison measure and not a trip-by-trip legal calculation, and the regulatory formula may need further adjustments. Floor shares in Tables 2, 2b and 3 are measured against the benchmark. The envelope is used as a month-by-month cross-check on the benchmark (Table 2c) and, for Lyft, as the floor in the censored regression of Section 6.3.

5. Empirical Strategy

The strategy has four parts, all estimated on 2024 data: floor prevalence and bunching (Section 5.1), the within-cell slope of pay on fare (Section 5.2), a descriptive event study of NBA game endings (Section 5.3), and a censored regression of pay on the floor and the fare (Section 5.4).

5.1 Floor prevalence and bunching

For each month and platform we compute the share of trips whose driver pay lies within 5 cents of the reconstructed floor, the share within 2%, and the share paid more than 5 cents below it (Tables 2 and 3). For April we also report the shares within 10, 25 and 50 cents (Table 2b), because the choice of tolerance changes the comparison between platforms. For April 2024 we also tabulate pay minus the floor in 5-cent bins (Figure 1).

Under the stylized model, exact bunching at a correctly measured floor could arise when the constraint binds. Here the benchmark is approximate and binning and tolerance choices matter, and Lyft's pay sits within about ten cents above the benchmark, so a five-cent window cuts through its mass; the observed proximity shares are descriptive and do not establish either excess mass relative to a smooth counterfactual or actual floor binding.

5.2 Within-cell slope of pay on fare

For each platform we report the conditional association from a linear specification with recorded driver pay (USD per completed trip) as the dependent variable, base passenger fare (USD per trip) as the regressor of interest, trip distance (miles) and duration (minutes) as controls, and pickup-zone × 15-minute fixed-effect cells. April point estimates are 0.2647 for Uber and 0.1841 for Lyft (494,140 and 165,772 trips; 43,389 and 28,234 fixed-effect cells). Full-year point estimates and reported confidence intervals appear in Table 2. Earlier analysis describes standard errors as clustered at the zone-time-cell level, but the original variance-estimation code and exact cluster construction have not been independently verified. Therefore the reported intervals should not be treated as fully audited inference; a replication table must report the precise cell identifier, number of clusters and estimator.

The reported coefficient is a conditional within-cell association, not a causal pass-through rate. A low coefficient could arise from floor constraints, separate pay and fare algorithms, measurement differences, selection, or remaining trip-composition variation. The model-implied slope for floor-binding trips is not estimated separately.

5.3 Event study design

We use the endings of 93 NBA games at Madison Square Garden and Barclays Center as predictable demand pulses, in the spirit of the end-of-event demand spikes that Yan et al. (2020) describe. For each game g, platform p, outcome Y and bin k we compute the core-minus-outer difference D_{gpk}, then report $\Delta_{gpk} = D_{gpk} - D_{gp,-8}$, averaged across games, with game-level standard errors and t-based 95% confidence intervals. Outcomes are fare, driver pay, wait and completed trips per zone (trips divided by the number of zones in the group). We also report

pre/post contrasts (bins +1 to +8 against −8 to −1, bin 0 excluded) and a pretrend diagnostic (bins −4 to −1 against −8 to −5), and repeat the analysis without the 22 overlapping games.

Interpretation limits. Pre-ending differences are substantial and the pretrend diagnostic is positive and significant in all twelve fare, pay and wait comparisons. Re-referencing to bin −8 does not eliminate ordinary clock-time changes in the core-minus-outer gap. We therefore report event-time co-movement, not the causal effect of an NBA game ending or an observed surge multiplier. Non-overlapping-game, alternative-reference, leave-one-out and end-time comparisons are descriptive robustness checks, not substitutes for matched non-event-day counterfactuals. No matched non-event-day estimate is reported.

5.4 Censored regression

Under the stylized model, driver pay is $y = \max\{F, s \cdot p\}$. To estimate how often the floor binds and what share of an extra fare dollar reaches the driver above it, we fit a censored (Tobit) regression of driver pay on the fare, trip miles and trip minutes: $y = \max\{F, y^*\}$, with $y^* = \beta0 + \beta1 \cdot \text{fare} + \beta2 \cdot \text{miles} + \beta3 \cdot \text{minutes} + \varepsilon$, where F is the floor, ε is normal and its standard deviation varies with the floor, $\ln \sigma = \gamma0 + \gamma1 \ln F$. A trip is treated as censored when its pay is within a tolerance of 5, 10 or 25 cents of F. We report the observed share of trips near the floor, the model-implied share $\Pr(y^* < F)$ averaged over trips, the latent fare slope $\beta1$, and the implied overall slope $\beta1 \cdot \Pr(y^* > F)$, which is comparable with the fixed-effects slope of Section 5.2.

The floor is the regulatory-rate benchmark or, as a robustness check and for Lyft, the platform's own estimated pay envelope (Section 4). The model is estimated by maximum likelihood on random samples of 200,000 trips per platform, drawn from 1,000,000 non-shared, non-wheelchair-accessible, non-Access-A-Ride trips from March to December, the months with published rates, with positive fare, pay, distance and duration and without Newark Airport and unrecorded drop-off zones; the price-index trim is not applied. Month-by-month estimates use 100,000 trips each. The regression has no zone-by-time fixed effects, assumes normal errors and a linear latent pay rule, and its standard errors are clustered by pickup date (306 clusters), using a sandwich estimator with per-trip scores computed by finite differences and the delta method for the model share and the implied slope. They are conditional on the floor and ignore estimation error in the envelope and the choice of subsample, and the month-by-month fits report only their dispersion across months. The regression is a descriptive check on the stylized model, not a structural test or an estimate of a platform take rate.

6. Results

6.1 The pay floor and the slope of pay on fare

Table 2. Summary statistics and pass-through by platform, 2024

Statistic	Uber	Lyft
Trips	179,125,798	60,344,650
Mean base fare, $	26.49	25.40

Statistic	Uber	Lyft
Mean driver pay, $	20.38	19.66
Pay ÷ fare	0.77	0.77
Mean wait, valid trips, minutes	4.51	5.14
Trips with negative wait	0.72%	1.70%
Share within 5 cents of the floor	56.5%	34.0%
Share within 2% of the floor	60.1%	77.4%
Share paid more than 5 cents below the floor	0.12%	0.0005%
Within-cell slope of pay on fare, full year [95% CI]	0.262 [0.261, 0.262]	0.200 [0.198, 0.201]
Within-cell slope, April only [95% CI]	0.265 [0.256, 0.273]	0.184 [0.172, 0.196]
Full-year estimation sample, trips	135,451,811	44,887,925

Note: averages are trip-weighted means of the monthly values. Floor shares use the floor sample (Section 4). Slopes are conditional associations between driver pay and base fare, controlling for trip miles, trip minutes and fixed effects (Section 5.2).

Tolerance matters more than the platform difference. A substantial share of restricted-sample trips clusters near the reconstructed benchmark. Over the full year, 56.5% of Uber and 34.0% of Lyft trips lie within $0.05, and 60.1% and 77.4% within 2%. In April, the share of Uber trips within $0.05, $0.10 and $0.25 is 63.9%, 64.2% and 67.0%, and the share of Lyft trips is 40.0%, 67.7% and 78.9% (Table 2b). Lyft's pay therefore sits within about ten cents above the benchmark, and a five-cent window cuts through that cluster. At wider tolerances Lyft has more mass near the benchmark than Uber, so the five-cent contrast should not be read as a difference in practice. These figures should not be interpreted as legally binding-floor shares.

Table 2b. Share of April trips within a tolerance of the benchmark

Platform	Within $0.05	Within $0.10	Within $0.25	Within $0.50
Uber	63.95%	64.24%	67.02%	70.50%
Lyft	40.02%	67.74%	78.94%	80.64%

Note: floor sample, April 2024 (Section 4); deviations are measured as driver pay minus the benchmark, in either direction. Within 2% of the benchmark: Uber 67.1%, Lyft 78.4%. The shares in Figure 1 and its source output that use the plotted range as denominator (Uber 68.5% and Lyft 41.6% within $0.05) differ from these because the denominators differ.

These patterns are compatible with the stylized model but do not by themselves test its distinctive kink or identify the share parameter s; Section 6.3 takes this up with a censored regression. The within-cell fare coefficients (approximately 0.26 for Uber and 0.20 for Lyft) describe conditional association, not floor-driven attenuation or causal pass-through. Alternative pay algorithms and trip composition remain plausible explanations.

The composition moves over the year (Table 3). Uber's share within 5 cents is 65% in July and August and falls to 37% in December, while average pay stays at 76% to 79% of the average fare.

Lyft's February–March step. The share of Lyft trips within 5 cents of the benchmark rises from about 2% in February to about 41% in March and stays between 37% and 43% for the rest of the

year, while the share within 2% is already 74% to 77% in January and February and stays between 74% and 82% through November. The TLC's 3.49% rate increase took effect on March 1, 2024, so the timing coincides with the regulatory change. The estimated pay envelope explains the step (Table 2c). Uber's envelope reproduces the benchmark in every month, including January and February: its per-mile and per-minute rates are within 0.04% and 0.11% of the benchmark rates. Lyft's envelope sits above the benchmark in every month, by 0.40% to 0.50% per mile, and per minute by 0.93% in January and February and by about 0.22% from March. On a trip of 5 miles and 20 minutes, that is about 14 cents above the benchmark before March and about 5 cents after, so the cluster of Lyft trips crosses the five-cent window while the share within 2% hardly moves. The Lyft envelope rose by about 2.6% per minute against about 3.4% in the TLC rate, so Lyft's pay schedule moved closer to the benchmark when the rates rose. The step therefore reflects the position of Lyft's own pay schedule relative to the benchmark and the width of the tolerance, not a mis-scaled January and February benchmark, and it is not evidence of a change in the legal floor or of compliance.

Uber below the benchmark. In April, 3,312,541 Uber trips (29.8% of 11,118,350) are recorded below the benchmark, but only 14,602 of them (0.13% of trips) are more than 5 cents below and the average shortfall among below-benchmark trips is about 3 cents. The signed distribution is close to symmetric around zero (Figure 1). This is more consistent with rounding or a slightly misstated rate parameter than with underpayment, although the aggregate output does not by itself establish either. Uber's estimated pay envelope has a per-minute rate 0.1% below the benchmark rate in every month (Table 2c), about a cent on a 20-minute trip, which is the order of magnitude of the average shortfall. Six of 3,938,941 Lyft trips are below the benchmark. We draw no compliance inference, and trip-level signed deviations, statutory scope, utilization and any platform-specific adjustments must be reconciled before legal or behavioral interpretation.

Table 2c. Estimated pay envelope relative to the TLC benchmark rates

Platform	Months	Envelope, $ per mile	Envelope, $ per minute	Ratio to TLC rates (mile; minute)
Uber	Jan–Feb	1.3145	0.5635	1.0003–1.0004; 0.9989–0.9991
Uber	Mar–Dec	1.3604	0.5824	1.0002–1.0004; 0.9989–0.9992
Lyft	Jan–Feb	1.3206	0.5693	1.0050; 1.0093
Lyft	Mar–Dec	1.3655	0.5843	1.0040–1.0041; 1.0021–1.0023

Note: 2nd-percentile quantile regression of driver pay on trip miles and minutes without an intercept, by platform and month, on random samples of 200,000 trips per month; ranges are across the months in each period. TLC benchmark rates are $1.314 per mile and $0.564 per minute in January and February (derived) and $1.360 and $0.583 from March.

Table 3. Monthly pay-to-fare ratio and share of trips near the floor, 2024

Month	Uber pay ÷ fare	Uber within 5c, %	Uber within 2%, %	Lyft pay ÷ fare	Lyft within 5c, %	Lyft within 2%, %
Jan	0.76	55.5	59.9	0.77	2.3	77.5
Feb	0.79	57.4	60.4	0.78	2.0	74.2

Month	Uber pay ÷ fare	Uber within 5c, %	Uber within 2%, %	Lyft pay ÷ fare	Lyft within 5c, %	Lyft within 2%, %
Mar	0.76	58.4	62.2	0.78	40.6	77.1
Apr	0.76	64.0	67.1	0.79	40.0	78.4
May	0.76	61.3	64.2	0.78	41.2	81.9
Jun	0.77	58.6	61.6	0.76	41.3	79.7
Jul	0.77	65.5	68.9	0.77	43.0	79.9
Aug	0.76	65.4	69.0	0.78	43.0	80.4
Sep	0.77	54.3	58.0	0.78	40.0	76.3
Oct	0.77	52.8	56.1	0.79	40.0	78.8
Nov	0.79	49.8	52.8	0.76	41.8	76.3
Dec	0.78	37.1	42.9	0.74	37.5	69.0

Note: pay ÷ fare is average driver pay over average base fare for all trips in the monthly file; the floor shares use the floor sample (Section 4).

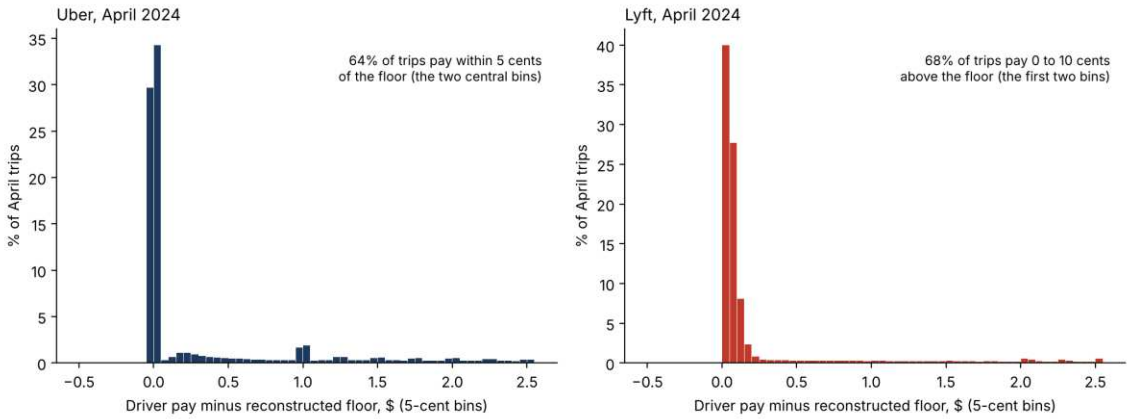


Figure 1. Driver pay minus the reconstructed floor, April 2024. Share of April trips in 5-cent bins, by platform. Uber has two central bins on either side of zero, which together hold 64% of trips, and secondary spikes about every 25 cents above $0.95. Lyft pays at or just above the floor, with 68% of trips 0 to 10 cents above it.

6.2 Event study: NBA game endings

At the first full bin after the ending, the core-minus-outer fare gap is $7.48 to $18.41 higher than at bin −8, and pickup waits are 2.2 to 2.9 minutes longer (Table 4). Uber's changes are larger than Lyft's at both arenas. These movements are predictable pulses, unlike the unforeseen emergency surge examined by Dolata and Schwabe (2024), so they speak to routine scarcity pricing rather than to crisis pricing.

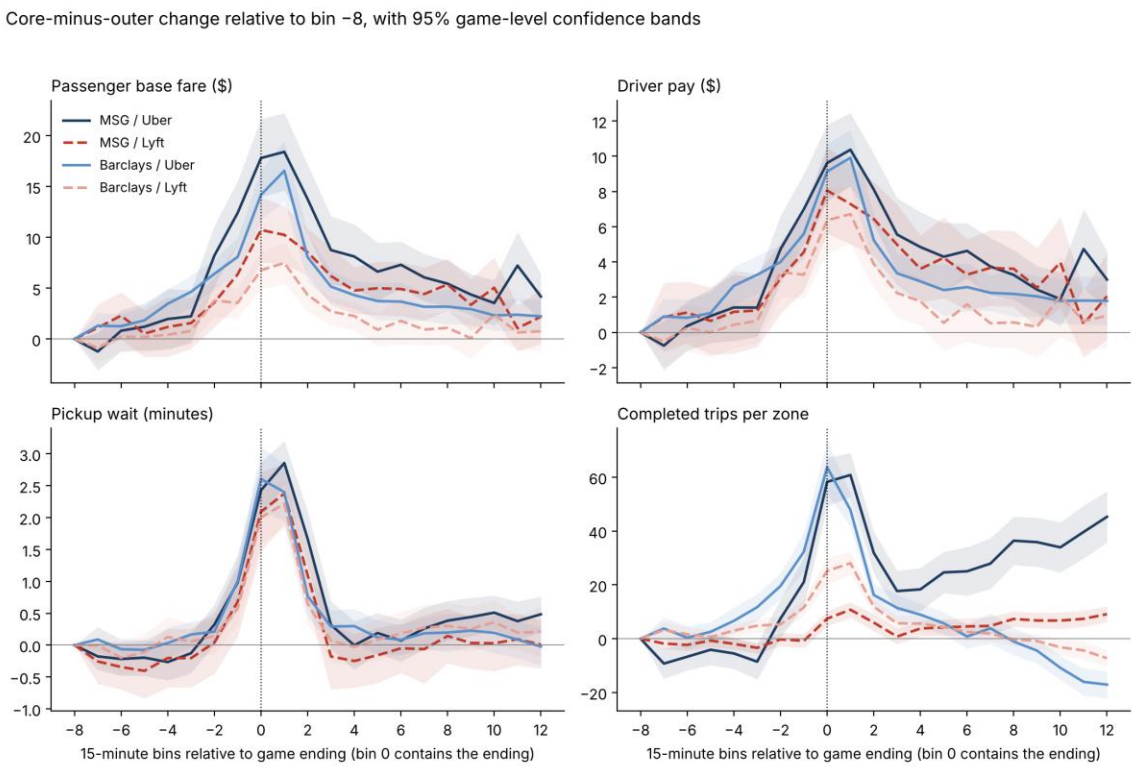


Figure 2. Event-study estimates for all 21 bins. Core-minus-outer change relative to bin −8, by arena and platform, with 95% game-level confidence bands. Fares, pay and waits build smoothly over the last hour before the ending and peak in bins 0 and +1; MSG Uber trips per zone stay elevated through bin +12.

Table 4. Change in the core-minus-outer gap at bin +1, relative to bin −8 (95% CI, game-level)

Arena / platform	Games	Fare, $	Driver pay, $	Wait, min	Trips per zone
MSG / Uber	53	18.41 [14.61, 22.22]	10.37 [8.29, 12.45]	2.86 [2.51, 3.20]	60.83 [52.65, 69.01]
MSG / Lyft	53	10.25 [7.42, 13.08]	7.29 [5.02, 9.56]	2.38 [1.90, 2.85]	10.75 [7.81, 13.69]
Barclays / Uber	40	16.55 [13.60, 19.50]	9.92 [8.33, 11.50]	2.40 [1.96, 2.85]	47.83 [41.71, 53.94]
Barclays / Lyft	40	7.48 [5.49, 9.47]	6.72 [5.05, 8.38]	2.22 [1.84, 2.60]	28.07 [23.86, 32.27]

Much of this is already present before the game ends. Between 47% and 67% of the bin +1 fare change, and between 26% and 42% of the wait change, is there at bin −1 (Table 5). Measured from bin −1, the event step in fares is $6.00 at MSG Uber, not $18.41. We therefore read the bin +1 versus −8 changes in Table 4 as including the pre-ending ramp and treat the −1 to +1 step as the better guide to the change at the ending itself. The zone definitions matter too. The MSG core zone (Penn Station/Madison Sq West) contains Penn Station, and the outer ring includes Times Sq/Theatre District, so late-night rail arrivals and Broadway let-outs can move the core-minus-outer gap independently of the game, while the Barclays core zone (Park Slope) is a large residential zone that may dilute the arena effect.

Table 5. Share of the bin +1 change already present at bin −1

Arena / platform	Fare gap at −1, $	Share of fare change at +1	Fare change, −1 to +1, $	Wait gap at −1, min	Share of wait change at +1
MSG / Uber	12.41	67%	6.00	0.98	34%
MSG / Lyft	6.38	62%	3.87	0.68	29%
Barclays / Uber	8.09	49%	8.47	1.01	42%
Barclays / Lyft	3.55	47%	3.93	0.57	26%

The pretrend diagnostic says the same thing (Table 6). At MSG Uber the pre/post fare contrast ($6.92) is almost the size of the pretrend ($7.00), and at Barclays Uber the pretrend ($4.87) exceeds the pre/post contrast ($3.34).

Table 6. Pre/post contrasts and pretrend diagnostic (game-level SE in parentheses)

Arena / platform	Fare pre/post, $	Fare pretrend, $	Wait pre/post, min	Wait pretrend, min
MSG / Uber	6.92 (1.05)	7.00 (1.06)	0.90 (0.13)	0.55 (0.12)
MSG / Lyft	4.53 (0.63)	2.22 (0.77)	0.62 (0.14)	0.40 (0.15)
Barclays / Uber	3.34 (0.59)	4.87 (0.41)	0.51 (0.07)	0.43 (0.08)
Barclays / Lyft	2.23 (0.38)	2.30 (0.34)	0.63 (0.10)	0.28 (0.08)

Driver pay rises with fares, but not one for one. The ratio of the pay change to the fare change at bin +1 is 0.56 to 0.90, and the fare-minus-pay gap widens by $0.77 to $8.04 (Table 7). These ratios are not pass-through estimates: trip mix moves fare and pay together, and fare minus pay is an accounting gap, not a take rate. The full-year within-cell slopes in Table 2 (Uber 0.26, Lyft 0.20) are lower than these event ratios, which is consistent with the event ratios mixing trip composition with pay-rule pass-through. Because tips are excluded from driver pay, the pay changes here are lower bounds on driver receipts per trip, and the sign of the tip response at high prices is uncertain (Liu et al., n.d.).

Table 7. Fare, pay and their gap at bin +1, relative to bin −8

Arena / platform	Fare change, $	Pay change, $	Pay ÷ fare	Fare − pay, $
MSG / Uber	18.41	10.37	0.56	8.04
MSG / Lyft	10.25	7.29	0.71	2.96
Barclays / Uber	16.55	9.92	0.60	6.64
Barclays / Lyft	7.48	6.72	0.90	0.77

At bin +12 the MSG Uber fare gap remains $4.18 above its bin −8 level (reported CI $1.92 to $6.44), and trips per zone are +45.3; Barclays Uber trips per zone are −17.1. Such persistence is consistent with, but does not prove, late-night baseline drift, differential local activity or lingering event-related changes. A matched non-event-day analysis has not been estimated, so these explanations cannot be separated.

The bin +1 averages are not driven by a few games (Table 8). Between 85% and 98% of games show a positive fare change and 92% to 100% a positive wait change. Dropping any one game moves the mean fare change by less than $1, and the median is within about $2.20 of the mean. Re-referencing to the average of bins −8 to −5 instead of bin −8 barely changes the estimates: MSG Uber fare goes from $18.41 to $18.22 and the share already present at bin −1 stays at 67%.

Table 8. Fare change at bin +1 across games, relative to bin −8

Arena / platform	Mean, $	Median, $	10% trimmed mean, $	Games with a positive change	Leave-one-out range of the mean, $
MSG / Uber	18.41	16.28	17.67	96%	17.74 to 19.16
MSG / Lyft	10.25	9.12	9.64	85%	9.74 to 10.71
Barclays / Uber	16.55	16.41	15.98	98%	15.87 to 17.14
Barclays / Lyft	7.48	6.63	7.31	88%	7.08 to 7.75

Persistence is a late-night pattern (Table 9). The bin +12 persistence comes from the 72 games that end at 21:30 or later. At MSG the Uber fare gap at bin +12 is +$5.55 for late games and −$0.97 to −$1.16 for earlier ones, and trips per zone are +54.8 against −1.3 for afternoon games. That is consistent with time-of-day drift in the core-minus-outer gap, which a matched non-event-day comparison (Appendix A4) would be needed to separate from the game effect. The pre-ending ramp does not disappear in earlier games: at bin −1 the Uber fare gap is already 49% to 54% of the bin +1 change in afternoon games and 61% in MSG early-evening games, but only 28% in Barclays early-evening games. Those groups have 4 to 7 games each, so they are suggestive only. Weekday games show larger changes than weekend games (MSG Uber $20.80 against $12.38 at bin +1).

Table 9. Uber fare change by game-end time, relative to bin −8 (bins −1, +1 and +12, $; trips per zone at +12)

Arena	Game ends	Games	Fare at −1	Fare at +1	Fare at +12	Trips per zone at +12
MSG	Before 19:00	7	5.48	10.07	−0.97	−1.3
MSG	19:00 to 21:30	4	2.44	3.97	−1.16	26.5
MSG	21:30 or later	42	14.52	21.18	5.55	54.8
Barclays	Before 19:00	5	3.79	7.75	1.79	−5.3
Barclays	19:00 to 21:30	5	5.54	19.44	−0.68	−9.3
Barclays	21:30 or later	30	9.23	17.54	2.79	−20.3

6.3 Censored regression: how often the floor binds and the pass-through above it

We start with the data and no model. Figure 3 plots, for each platform, the share of trips within 10 cents of the floor against the ratio of the base fare to the floor, in 20 equal-count bins of a random sample of 1,000,000 trips from March to December; the floor is the TLC benchmark for Uber and Lyft's estimated pay envelope for Lyft. For Uber the share is 75% to 77% up to a ratio of about 1.2, falls to 64% in the 1.62 to 1.69 bin and to 31% in the 1.87 to 2.03 bin, and is below 2% above

2.3, while mean pay minus the benchmark stays between $0.41 and $0.51 up to a ratio of about 1.3 and rises to $8.76 in the top bin. For Lyft the share is 93% to 94% up to a ratio of about 1.3, falls to 77% in the 1.56 to 1.63 bin and to 45% in the 1.84 to 2.03 bin, and is 7% in the top bin, while mean pay minus the envelope stays between $0.18 and $0.22 up to a ratio of about 1.3 and rises to $4.50. This is the flat-then-rising pattern of Proposition 1, although the transition is gradual and not a sharp kink, as would happen if the pay share varies across trips or the floor is measured with noise.

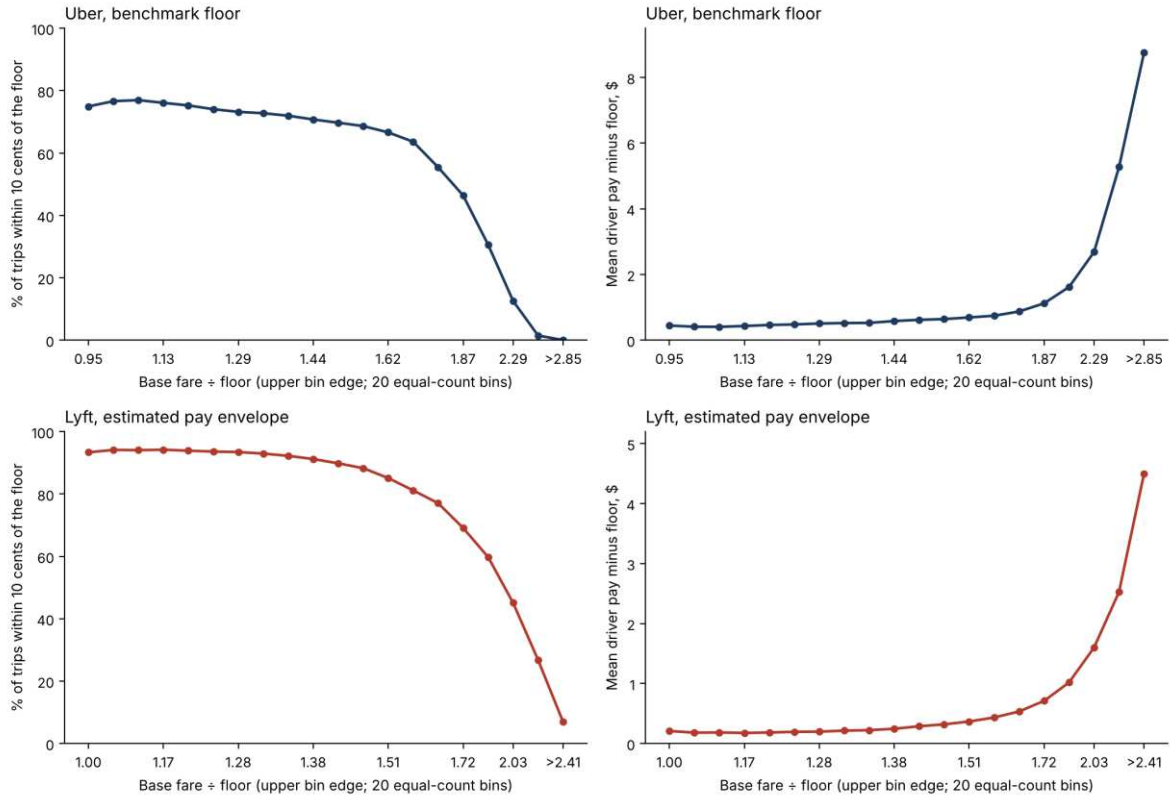


Figure 3. Proximity to the floor against the fare-to-floor ratio, March–December 2024. Twenty equal-count bins of a random sample of 1,000,000 trips per platform; the upper bin edge is shown. The floor is the TLC benchmark for Uber and Lyft's estimated pay envelope for Lyft. Left: share of trips within 10 cents of the floor. Right: mean driver pay minus the floor.

Table 10 reports the censored-regression estimates, with standard errors clustered by pickup date (306 dates). For Uber the results are the same under the benchmark and the estimated envelope and barely move with the tolerance: 58% to 61% of trips are near the floor, the model puts latent pay below the floor on about 60% of trips (standard error 0.8 points), and the latent fare slope is 0.51 to 0.53 (standard error 0.008), so that above the floor roughly half of each additional fare dollar reaches the driver, holding miles and minutes fixed. Weighted by the probability that a trip is above the floor, the implied overall slope is 0.20 to 0.21 (standard error 0.004), below the fixed-effects slope of about 0.26 by much more than sampling error. The floor-plus-share rule can therefore account for the low observed slope, and the model attributes somewhat more attenuation to the floor than the fixed-effects regression shows. The two are not like for like: the censored

model has no zone-by-time fixed effects and covers March to December with airport drop-offs removed, so we read the match as rough consistency and not as confirmation.

For Lyft the estimates depend on the floor. Against the TLC benchmark they vary with the tolerance, from 39% to 78% of trips near the floor and a latent slope of 0.27 to 0.57, because Lyft's pay sits just above that benchmark (Table 2c). Against Lyft's own estimated pay schedule they are stable: 78% to 79% of trips near the floor, a model share of about 77% (standard error 0.6 points), a latent slope of 0.56 to 0.60 (standard errors 0.025 to 0.031) and an implied overall slope of 0.13 (standard error about 0.008), against a fixed-effects slope of about 0.20. Setting the envelope at the 1st, 2nd or 5th percentile of pay instead gives near-floor shares of 77% to 79%, model shares of 76% to 78%, latent slopes of 0.54 to 0.60 and the same implied slope of 0.13, so the result does not hinge on one envelope choice. Because the envelope is itself a lower envelope of observed pay, clustering at it is partly mechanical, so we describe it as Lyft's effective pay schedule and not as a legal minimum.

Table 10. Censored-regression estimates by platform, floor and tolerance, March–December 2024

Platform	Floor	Tolerance	Near floor	Floor binds (model)	Latent slope	Implied slope
Uber	TLC	$0.05	57.6%	59.6% (0.8)	0.513 (0.008)	0.207 (0.004)
Uber	TLC	$0.10	57.9%	59.7% (0.8)	0.514 (0.008)	0.207 (0.004)
Uber	TLC	$0.25	60.6%	61.4% (0.8)	0.526 (0.008)	0.203 (0.004)
Uber	Envelope	$0.05	57.6%	59.5% (0.8)	0.512 (0.008)	0.208 (0.004)
Uber	Envelope	$0.10	57.9%	59.7% (0.8)	0.514 (0.008)	0.207 (0.004)
Uber	Envelope	$0.25	60.4%	61.2% (0.8)	0.526 (0.008)	0.204 (0.004)
Lyft	TLC	$0.05	38.7%	53.8%	0.269	0.124
Lyft	TLC	$0.10	65.3%	70.0%	0.312	0.094
Lyft	TLC	$0.25	78.4%	77.0%	0.569	0.131
Lyft	Envelope	$0.05	77.6%	76.7% (0.6)	0.561 (0.031)	0.131 (0.009)
Lyft	Envelope	$0.10	78.1%	77.2% (0.6)	0.583 (0.027)	0.133 (0.008)
Lyft	Envelope	$0.25	78.8%	77.7% (0.6)	0.601 (0.025)	0.134 (0.008)

Note: Tobit with $\ln \sigma = \gamma 0 + \gamma 1 \ln F$, fitted by maximum likelihood on random samples of 200,000 trips per platform. Near floor is the share of trips within the tolerance of the floor; Floor binds is the model-implied share with latent pay below the floor, in percent; Latent slope is $\beta 1$; Implied slope is $\beta 1$ times the average probability that latent pay exceeds the floor. Standard errors in parentheses are clustered by pickup date (306 clusters) and are conditional on the floor; they are not computed for the Lyft rows that use the TLC benchmark. Fixed-effects slopes (Table 2): Uber 0.262 for the full year, Lyft 0.200.

Estimated month by month with each platform's own envelope (Table 11), Uber's share near the floor is 50% to 66% from March to November and falls to 38% in December, the same fall seen in Table 3, while its implied slope moves from 0.17–0.22 to 0.25. Lyft's share is 69% to 82%, also lowest in December, with a latent slope of 0.43 to 0.68. Each month uses 100,000 trips, so we read the ranges and not individual months. Across months, the mean latent slope is 0.501 (standard error 0.012) for Uber and 0.561 (0.020) for Lyft, and the mean implied slope is 0.201 (0.007) and 0.120 (0.007).

Table 11. Month-by-month censored-regression estimates, estimated pay envelope as the floor, 10-cent tolerance

Month	Uber near floor	Uber latent slope	Uber implied slope	Lyft near floor	Lyft latent slope	Lyft implied slope
Mar	58.9%	0.516	0.205	78.1%	0.543	0.121
Apr	64.2%	0.526	0.182	78.9%	0.599	0.133
May	62.1%	0.526	0.194	82.3%	0.434	0.076
Jun	58.9%	0.495	0.197	80.2%	0.565	0.108
Jul	65.8%	0.549	0.180	79.7%	0.528	0.102
Aug	65.8%	0.524	0.171	80.5%	0.680	0.112
Sep	54.9%	0.516	0.222	76.6%	0.549	0.126
Oct	53.6%	0.441	0.195	79.2%	0.615	0.133
Nov	50.0%	0.462	0.217	76.6%	0.550	0.126
Dec	37.7%	0.453	0.251	69.3%	0.546	0.161
Mean	57.2%	0.501	0.201	78.1%	0.561	0.120
SE of mean	2.7	0.012	0.007	1.1	0.020	0.007

Note: 100,000 trips per platform and month; columns as in Table 10. The last two rows give the mean across the ten months and its standard error (the standard deviation across months divided by √10), in percentage points for the near-floor columns.

What these estimates do and do not show. They are model-based and conditional on normal errors, a linear latent pay rule and the choice of floor; they omit cell fixed effects, and their standard errors, clustered by pickup date, are conditional on the floor and ignore estimation error in the envelope and the choice of subsample. They support the reading that, for a large share of trips, recorded pay is set at or very close to a schedule that matches the TLC rates for Uber and sits slightly above them for Lyft, and that above that schedule roughly half of each additional fare dollar reaches the driver. They do not identify the platform take rate, rider incidence or the fraction of trips that are legally constrained, and the censored-regression shares are not legal binding shares.

7. Robustness and Validation

Robustness checks are reported where the results appear. Floor shares are shown at tolerances of 5 cents, 10 cents, 25 cents, 50 cents and 2% (Tables 2, 2b and 3), and the benchmark is cross-checked against each platform's estimated pay envelope month by month (Table 2c). The censored regression is run at three tolerances and two floors (Table 10) with pickup-date-clustered standard errors, by month (Table 11), and with Lyft's envelope re-estimated at the 1st, 2nd and 5th percentiles. The event-study estimates are checked against excluding the 22 overlapping games, a reference window of bins −8 to −5, dropping any single game, and the time of the game ending (Section 6.2, Tables 8 and 9). Four procedures are specified in Appendix A but are not yet implemented in the reported results: a sequential sample-flow audit, trip-level validation of the pay benchmark against the statutory calculation, documented clustered inference for the fixed-

effects slope, and a matched non-event-day comparison. The interpretations in Section 6 are worded to hold without them.

8. Extensions and Policy Implications

Policy implications are conditional model predictions rather than measured policy effects. If the stylized floor-plus-share rule held and trip characteristics, service quantities and platform responses were fixed, changing a floor would mechanically affect payments on constrained trips while a fare cap could leave those payments unchanged. The reported proximity shares do not identify which trips are actually constrained, and behavioral responses could change these predictions.

Price caps. In the hypothetical model, a binding cap can reduce above-floor fare-linked payments, whereas floor-constrained payments are mechanically unaffected at fixed trip characteristics. Actual platform pricing and driver-compensation algorithms are not observed, and supply, demand and matching responses are outside the analysis. The event-study pay-to-fare-change ratios are not estimates of a payment share.

Pay standards. Within the stylized model, an increase in the floor mechanically raises compensation for a trip that remains floor-constrained with fixed characteristics. The observed benchmark-proximity shares (in April, 63.9% of Uber and 40.0% of Lyft trips within five cents, and 64.2% and 67.7% within ten cents) are not verified treatment shares and cannot be used to quantify the aggregate effect of a regulatory change, and the model-based shares in Section 6.3 depend on normal errors and a linear latent pay rule, so they do not identify it either.

Platform margins and distributional concerns. Because fare and pay are net of different items, the data do not identify the platform's take rate. The gap between average fare and average pay, about 23% of the average fare on both platforms over the year, is an accounting wedge that also contains taxes and fees. Evidence from the United Kingdom that dynamic pricing widened pay dispersion across drivers (Binns et al., 2025), and from Chicago that prices fall unevenly across neighborhoods (Saxena et al., 2025), indicates distributional margins that trip-level New York data could address. Transparency in how fares and pay are set is a recurring policy lever in this literature (Han et al., 2018; Grace and Okoh, 2025; Liu et al., n.d.).

Extensions. Five pieces of work would turn the descriptive results into estimates of incidence:

- Extending the censored regression of Section 6.3 with zone-by-time fixed effects and a test of the restriction that the kink sits at F/s, on trip-level data by platform and month.
- A decomposition of changes in the take rate into floor, share-rule and composition components.
- Instrumenting a fare measure with event and weather shocks to address measurement error, and estimating rider incidence over a range of values of time, with the Uber–Lyft substitution pattern as an extension.

- A matched non-event-day control for the NBA event study, with a placebo on non-game days.
- Counterfactual price caps with a simple matching function (Castillo, Knoepfle and Weyl, 2025), and extensions to rain and subway disruptions and to the 2019 to 2024 rate changes.

These extensions need the trip-level records and, for the event designs, zone-by-15-minute panels for ordinary days. Appendix A specifies the matched non-event-day design and the audits; neither these procedures nor the extensions listed here are part of the evidence reported.

9. Limitations

Only completed trips are observed, so cancellations and unmet demand are missing. There are no driver identifiers, so pay is per trip and not hourly earnings, and tips are excluded from pay. The sample is one year, platform pay rules are not directly observed, and the floor is a benchmark built from the TLC's published in-city rates, not a trip-by-trip legal calculation, with the composition of the roughly 24% of trips outside the floor sample not yet audited, and with January and February parameters derived from the March rates, which the Uber pay envelope supports but a TLC notice has not confirmed. The within-cell slope is a conditional association. The censored regression assumes a linear latent pay rule with normal heteroskedastic errors, omits cell fixed effects, uses random subsamples with pickup-date-clustered standard errors that are conditional on the floor and, for Lyft, depends on an estimated pay envelope, so its shares are model-based and not legal binding shares. For the event study: a positive, significant pretrend in all twelve fare, pay and wait comparisons; a reference bin (−8) that absorbs time-of-day drift; outer zones that may respond to the arena themselves, including the Times Sq/Theatre District zone in the MSG outer ring; and no matched non-event-day control. Several of the studies we draw on use data from other cities or countries, come from authors affiliated with Uber, or are working papers, so we use them for context and benchmarks, not as estimates for New York. The claims in this paper are limited to what survives these constraints.

10. Conclusion

The 2024 NYC TLC records show substantial clustering of recorded driver pay near a reconstructed regulatory-rate benchmark, with a platform ranking that depends on the tolerance, low conditional within-cell associations between passenger base fare and driver pay, a censored-regression pattern in which the floor is near-binding on most sampled trips and roughly half of each additional fare dollar reaches the driver above it, and marked fare, pay and waiting-time co-movement around 93 NBA game endings. These results do not identify actual surge multipliers, a platform take rate, the fraction of legally floor-binding trips or causal effects of game endings or regulation. Stronger incidence or policy claims would require, at a minimum, a complete exclusion audit, trip-level validation of the regulatory benchmark, documented clustered inference and a matched non-event comparison; each of these is specified in Appendix A and none is yet part of the evidence reported here.

Appendix A. Replication procedures and audit specifications

Status of these procedures. Procedures A1 to A5 are specifications. They have not been run on the data behind the reported results, and the file names identify the planned scripts. A6 describes the envelope check and censored regression that were run for Sections 6.1 and 6.3. Results from the remaining procedures will replace the corresponding caveats in Sections 4, 6 and 9.

A1. Sample-flow audit (01_sample_flow_audit.py). Starting from 179,125,798 Uber and 60,344,650 Lyft completed trips, the script tabulates sequential removals by platform and month for service category, missing or nonpositive fare, pay, miles and minutes, drop-off scope and the price-index trim, in the original order, and compares retained and excluded trips on fare, pay, distance, duration, wait, airport, weekend and night shares. It also reports trips whose pickup month differs from the file month and exact-duplicate counts.

A2. Regulatory-floor validation (02_floor_validation.py). For each trip the script computes the signed difference between recorded pay and the benchmark in cents, under the rate schedule in force on the pickup date and, as a diagnostic, under both the pre- and post-March schedules, separately for trips ending inside and outside the city. It reports shares below, equal to and within one and five cents of the benchmark, daily shares around 1 March 2024, a fixed-sample re-scoring of February and March trips and a rounding diagnostic. These diagnostics are designed to separate a benchmark artifact from a change in payment practice in Lyft's February–March jump.

A3. Fixed-effects replication (04_fixed_effects_slope.py). The script estimates, by platform, the within-cell regression of recorded driver pay on base fare, trip miles and trip minutes, with a cell defined as pickup zone by 15-minute timestamp. It reports the number of observations, cells and clusters, the coefficient, the standard error and the 95% interval under cell-, pickup-zone- and pickup-date-clustered variance estimators; the sandwich estimator and small-sample correction are documented in the script.

A4. Event-time counterfactual (03_event_study_matched_controls.py). Each game is matched to up to six non-game windows at the same arena, weekday and New York local clock time (shifts of one to six whole weeks in either direction), excluding windows that overlap another game, fall on a listed arena event or a federal holiday, or leave the data range. The script reports game-level

matched estimates (game minus the mean of its controls) beside the unmatched estimates, the pretrend and pre/post contrasts in matched form, the share of the bin +1 change present at bin −1, and a placebo in which one control window is treated as a pseudo-game.

A5. Sources and code. Statutory rates and effective dates used by the scripts are stored in config/rates.json and cited to TLC Industry Notice 24-02 (NYC TLC, 2024a). The January and February 2024 parameters will be taken from the notice in force at that time and not derived from the March rates. All scripts read only the public TLC trip files, the TLC zone lookup and the list of game-ending times.

A6. Envelope check and censored regression (run for Sections 6.1 and 6.3). Samples are random draws (DuckDB reservoir sampling, seed 42) from the public monthly files: 1,000,000 trips per platform from March to December 2024 for Tables 10 and 11, and 200,000 trips per platform and month for Table 2c. Restrictions: non-shared, non-wheelchair-accessible and non-Access-A-Ride service, positive fare, pay, miles and minutes, and drop-off zones other than Newark Airport (1) and unknown zones (264, 265); the price-index trim is not applied. The benchmark uses \$1.360 per mile and \$0.583 per minute. The envelope is a 2nd-percentile quantile regression of pay on miles and minutes with no intercept, estimated on a 300,000-trip subsample for Tables 10 and 11 and by month for Table 2c. The Tobit log-likelihood treats trips with pay at most the tolerance above the floor as censored, with latent mean $\beta 0 + \beta 1$ fare $+ \beta 2$ miles $+ \beta 3$ minutes and log standard deviation $\gamma 0 + \gamma 1 \ln F$; it is maximized by BFGS from OLS starting values on 200,000-trip subsamples (100,000 for monthly fits). Figure 3 uses 20 equal-count bins of fare divided by the floor (the benchmark for Uber, the estimated envelope for Lyft) and a 10-cent tolerance, drawn from the same 1,000,000-trip samples. Standard errors for Table 10 use a sandwich variance clustered by pickup date (306 clusters), with per-trip scores and the Hessian computed by finite differences and the delta method for the model share and the implied slope; the Lyft envelope is re-estimated at the 1st, 2nd and 5th percentiles as a robustness check. Table 11 reports the mean and standard error across months.

Data Availability

All trip data are the public NYC Taxi and Limousine Commission High-Volume For-Hire Vehicle Trip Records for 2024 and the TLC taxi zone lookup (NYC TLC, 2024b). Game-end times are the wall-clock timestamps of the final play in ESPN's historical play-by-play records, converted to New York time, for 93 games with no duplicate game identifiers. Code for the envelope check and the censored regression (Section 6.3 and Appendix A6) will be released with the paper; the other Appendix A procedures are specified and not yet implemented.

Declaration

The author received no external funding for this study and declares no competing interests.

References

Angrist, J. D., Caldwell, S., & Hall, J. V. (2021). Uber versus taxi: A driver's eye view. *American Economic Journal: Applied Economics, 13*(3), 272–308. https://doi.org/10.1257/app.20190655

Binns, R., Stein, J., Datta, S., Van Kleek, M., & Shadbolt, N. (2025). Not even nice work if you can get it: A longitudinal study of Uber's algorithmic pay and pricing. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery. https://doi.org/10.1145/3715275.3732099

Cachon, G. P., Daniels, K. M., & Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management, 19*(3), 368–384.

Castillo, J. C. (2025). Who benefits from surge pricing? *Econometrica, 93*(5), 1811–1854. https://doi.org/10.3982/ECTA19106

Castillo, J. C., Knoepfle, D., & Weyl, E. G. (2025). Matching and pricing in ride hailing: Wild goose chases and how to solve them. *Management Science*. https://doi.org/10.1287/mnsc.2022.00096

Chen, M. K., & Sheldon, M. (2015). *Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform* [Working paper]. UCLA Anderson School of Management and University of Chicago.

Chirita, M., & Chirita, G. (2024). A comprehensive overview of deep learning for algorithmic pricing in ride-sharing platforms. *Annals of "Dunarea de Jos" University of Galati, Fascicle I: Economics and Applied Informatics, 30*(1). https://doi.org/10.35219/eai15840409404

Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). *Using big data to estimate consumer surplus: The case of Uber* (NBER Working Paper No. 22627). National Bureau of Economic Research.

Dolata, M., & Schwabe, G. (2024). Towards the socio-algorithmic construction of fairness: The case of automatic price-surging in ride-hailing. *International Journal of Human–Computer Interaction, 40*(1), 55–65. https://doi.org/10.1080/10447318.2023.2210887

Garg, N., & Nazerzadeh, H. (2022). Driver surge pricing. *Management Science, 68*(5), 3219–3235.

Grace, I., & Okoh, O. F. (2025). Digital platforms and algorithmic pricing: Investigating market efficiency and consumer welfare in the age of big data. *Malaysian E Commerce Journal, 9*(2), 26–34. https://doi.org/10.26480/mecj.02.2025.26.34

Hall, J. V., Horton, J. J., & Knoepfle, D. T. (2023). *Ride-sharing markets re-equilibrate* (NBER Working Paper No. 30883). National Bureau of Economic Research.

Han, W., Gao, Y., & Deng, A. (2018). Algorithmic price discrimination on online platforms and antitrust enforcement in China's digital economy. *The Antitrust Source*, August 2018. https://ssrn.com/abstract=3250075

Konishi, A., Ramakrishnan, V., Waheed, S., & Herrera, L. (2023). *Analysis of high volume for-hire vehicle data for New York City: Selected months, 2019–2022* [Research brief developed for the New York Taxi Workers Alliance]. UCLA Labor Center. https://escholarship.org/uc/item/5q09h352

Koustas, D., Parrott, J., & Reich, M. (2020). *New York City's gig driver pay standard: Effects on drivers, passengers, and the companies* [Report]. Center for New York City Affairs, The New School. https://escholarship.org/uc/item/3xb4152j

Leccese, M. (2024). *Taxing ride-sharing: Which neighborhoods pay more?* [Unpublished working paper, April 6, 2024].

Liu, X., Zhang, L., Ma, Y., Ma, W., & Teo, H.-H. (n.d.). *How algorithmic pricing aggressiveness affects tipping behaviour on two-sided platforms: Empirical evidence from Uber* [Unpublished manuscript submitted to an INFORMS journal] (SSRN Working Paper No. 5080000). Social Science Research Network.

Lu, A., Frazier, P. I., & Kislev, O. (2018). *Surge pricing moves Uber's driver-partners* (SSRN Working Paper No. 3180246). Social Science Research Network.

New York City Taxi and Limousine Commission. (2018). *Notice of promulgation of rules: Driver income and vehicle lease transparency* (adopted December 4, 2018). NYC TLC.

New York City Taxi and Limousine Commission. (2024a). *Industry Notice 24-02: Minimum high-volume for-hire driver pay rates to increase 3.49% beginning March 1, 2024*. NYC TLC.

New York City Taxi and Limousine Commission. (2024b). *TLC trip record data: High volume for-hire vehicle trip records, January–December 2024*. https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page

New York City Taxi and Limousine Commission. (2025, March 18). *Data dictionary: High volume for-hire vehicle trip records*. NYC TLC.

Parrott, J. A., & Reich, M. (2018). *An earnings standard for New York City's app-based drivers: Economic analysis and policy assessment* [Report]. Center on Wage and Employment Dynamics, University of California, Berkeley, and Schwartz Center for Economic Policy Analysis, The New School.

Raza, M. A., Shah, B., Singh, A., Goyal, A., & Upadhyay, N. (2026). Fare prediction, demand forecasting, and surge pricing in Indian ride-hailing: A case study of Bangalore. In *2026 International Conference on Computing, Sciences and Communications (ICCSC)*. IEEE. https://doi.org/10.1109/ICCSC67078.2026.11468610

Rosaia, N. (2025). Competing platforms and transport equilibrium. *Econometrica, 93*(6), 2235–2271. https://doi.org/10.3982/ECTA21773

Saxena, N. A., Zhang, W., & Shahabi, C. (2025). Unveiling and mitigating bias in ride-hailing pricing for equitable policy making. *AI and Ethics, 5*, 1549–1560. https://doi.org/10.1007/s43681-024-00498-3

Yan, C., Zhu, H., Korolko, N., & Woodard, D. (2020). Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics, 67*(8), 705–724. https://doi.org/10.1002/nav.21872

Zhang, Y., Sang, Y., & Wu, M. (2026). Congestion pricing in New York City: Effects on ride-hailing and transit. *Transportation Research Part A: Policy and Practice, 208*, 104966. https://doi.org/10.1016/j.tra.2026.104966