Platform Pricing and the Geography of Congestion Tax Incidence: Evidence from Ride-Hailing in Chicago

Keitaro Ninomiya*

Abstract

Does a geographically uniform congestion tax generate uniform incidence across local transportation markets? I study Chicago's ride-hailing congestion surcharge using public trip-level administrative records and a fuzzy regression-discontinuity design at the daily 10 p.m. tax boundary. Pass-through is incomplete in both markets: in a common January–February 2026 sample, 75.2% of the realized charge reaches the prices of urban trips and 60.6% reaches the prices of suburban trips. The platform's own fare falls when the charge applies, by about $0.29 per dollar of charge in both markets, so platform pricing absorbs part of the statutory burden. After the platform restructured its pricing in March 2026 that offset roughly doubles and pass-through falls to 39.1% for urban trips ($p = 0.012$ for the change) and 18.5% for suburban trips. Trip-volume effects are small and statistically insignificant in both regimes. The robust findings are therefore that a mobility tax need not be passed through fully, because the platform's own fare adjustment absorbs part of it, and that the size of that absorption moves with the platform's pricing structure rather than with the tax.

*Ninomiya: University of Illinois Urbana-Champaign, email: keitaro2@illinois.edu. As this is a single-authored manuscript, author order is not applicable. I received no financial support for this research. I am indebted to Dan Bernhardt, Andy Garin, Alex Bartik, and Russell Weinstein for guidance and comments on earlier drafts. IRB approval was not applicable because the study uses administrative trip records and involved no direct interaction with human subjects. This study is not an AEA-registered randomized controlled trial.

I. Introduction

A congestion tax may be uniform in statute yet uneven in economic incidence. Cities use taxes and fees on travel both to finance transportation systems (Anas, 2012) and to manage congestion (Almagro et al., 2024; Cook et al., 2025). Transportation demand, available alternatives, and pricing conditions vary across a metropolitan area, so the same statutory charge need not produce the same consumer-price increase in every local market. This paper asks whether a geographically uniform congestion tax generates geographically uniform incidence across urban and suburban ride-hailing trips.

I study the local response of urban travel markets to the tax, using trips served by ride-hailing platforms. Ride-hailing is a useful setting for this question because platform prices can adjust across locations and times. This flexibility differs from transportation services with largely fixed fare schedules. When a city adds a uniform charge, the price ultimately paid by a ride-hailing passenger reflects both the statutory charge and any contemporaneous adjustment in the underlying fare. The platform can also respond centrally by redesigning its price schedule—for example, by changing booking fees—rather than leaving prices to adjust only through local rider and driver behavior. Platform pricing can therefore mediate the incidence of a local mobility tax even when the tax schedule itself does not vary across covered locations. This centralized supply-side response is important when cities use mobility charges to finance transportation infrastructure: the effective burden, travel incentives, and potentially the stability of the tax base depend partly on a private platform's network-wide pricing decisions. The relevant distinction is therefore not only between the statutory and economic geography of the tax, but also between public infrastructure financing and the centralized private pricing system through which the charge reaches travelers.

I study tax incidence across multiple pricing regimes set by the operating platform, through the effect of the tax on prices and volume. Ride-hailing platforms in Chicago add a booking fee on top of the metered fare and the statutory charge. Between March and June 2026 the average fee more than doubled, and the share of trips carrying no booking fee at all fell from about one in ten to almost none[1]. The city introduced this time-and-location surcharge for trips involving a downtown zone in January 2020, and expanded it on January 6, 2026 by extending the first zone beyond the original downtown area and adding a second zone in Hyde Park. During the period I study, a $1.50 congestion charge applied to covered trips during the daily 6 a.m.–10 p.m. window. I use public administrative records covering

[1]The date of a trip predicts its fee better than the trip's distance or route does, the signature of a schedule the platform reset repeatedly over these months.

the universe of reported transportation network provider trips in Chicago. The records contain fares, final prices, additional charges, trip distance and duration, rounded timestamps, and pickup and drop-off geography.

The empirical design uses the daily discontinuity in tax exposure at 10 p.m. Because reported timestamps are rounded to 15-minute bins and observed charges do not follow assignment perfectly, the design is fuzzy rather than sharp. I reconstruct whether each trip should be exposed to the charge from its time, duration, and endpoint geography. The reconstructed time and geographic rules each classify observed charge status with approximately 99% accuracy. I then use cutoff assignment as an instrument for the additional charge actually recorded on the trip, comparing trips picked up within a narrow window on either side of 10 p.m. and excluding the ambiguous 10 p.m. bin.

I find incomplete pass-through in both markets at the immediate introduction of the congestion tax, with consistently lower point estimates for suburban trips. In the preferred $\pm$30-minute specification, passengers bear 75.2% of the realized charge on urban trips and 60.6% on suburban trips. This comparison holds the calendar window and specification fixed across markets by using trips from January 27–February 27, 2026. The 14.5-percentage-point urban–suburban gap is imprecisely estimated (SE 15.0 points; $p = 0.33$). The $\pm$15-minute window preserves the ordering. At $\pm$60 minutes the gap grows to 43.5 points and is statistically significant ($p = 0.006$), but a 2025 diagnostic indicates that this increase comes from the additional clock-time support rather than sample size alone. I therefore treat the wider-window result as sensitivity evidence rather than the preferred estimate.

The price response reflects adjustment in the price of a comparable trip rather than a change in which trips occur. Differences in community-area origin–destination composition account for under two percent of the price discontinuity in either market, so the response happens within local travel markets, not across them. The adjusting margin is the platform's own fare: the reported base fare falls by \$0.285 per dollar of realized charge on urban trips and \$0.295 on suburban trips before March, and by \$0.655 and \$0.592 after March. One minus the urban fare offset is 0.715 before March and 0.345 after, against measured pass-through of 75.2% and 39.1%, so the offset accounts for the pass-through estimates and its near-doubling accounts for their decline. Trip distance and duration move by at most 0.03 miles and sixteen seconds per dollar of charge, too little to explain price movements of this size. The public records contain no surge indicator and do not separately identify a fixed base-fare component, so I cannot say which part of the fare moves.

The trip-volume estimates are also small: before March, volume falls by 2.0% in the urban

market and 1.8% in the suburban market; after March, the estimated declines are 0.1% and 1.3%, respectively. None is statistically distinguishable from zero, so I do not use the volume estimates to construct demand elasticities.

The platform's observed pricing structure changed in March 2026, and pass-through is markedly lower afterward. Estimated on the same specification as the pre-March comparison, urban pass-through falls from 75.2% to 39.1% and suburban pass-through from 60.6% to 18.5%. Testing both regimes jointly, the urban decline is 36.1 percentage points (standard error 14.3 points; $p = 0.012$), while the suburban decline of 42.2 points is imprecise ($p = 0.12$). Over the same months the booking fee the platform adds beyond the statutory charge rose from \$2.66 to \$6.14 per trip on average, so a smaller share of an unchanged statutory charge reached consumer prices as the platform's own fee grew. The regime is not randomly assigned, so this is a contrast between two periods rather than an estimate of the causal effect of the pricing change.

The paper's primary contribution is to urban economics, transportation, and local public finance. It shows that the effective consumer burden of a mobility tax cannot necessarily be inferred from its statutory schedule: platform pricing is part of the incidence mechanism. The direct comparison also demonstrates the importance of holding time and specification fixed when attributing different estimates to geography. These findings complement the literature on using road-user charges to finance transportation and on how the use of the resulting revenue affects efficiency and incidence (Small, 1992; de Palma and Lindsey, 2007; Lindsey, 2012; Anas, 2012; Parry, 2009). The paper also contributes to research on ride-hailing tax incidence. Existing work studies ride-hailing as a local tax base, the effects of ride-hailing on congestion, and the average incidence and distributional consequences of ride-hailing taxes (Agrawal and Zhao, 2023; Cairncross et al., 2021; Tarduno, 2021, 2025; Leccese, 2024). I complement this literature by estimating local price responses to the same charge, directly testing their equality in a common period, and decomposing adjustment in the platform-reported fare.

The remainder of the paper describes Chicago's policy and the administrative data, presents the fuzzy-RD design and its assignment validation, reports the price and fare results, and concludes.

II. Institutional Setting and Data

A. Chicago's congestion surcharge

Chicago taxes trips supplied by transportation network providers (TNPs), including Uber, Lyft, and Via. In January 2020, the city introduced an additional time-and-location congestion surcharge for trips beginning or ending in a designated Downtown Zone. The policy was subsequently modified, including an extension to weekend trips in 2025. These earlier changes provide institutional context, but the paper's main estimates concern the geographic expansion implemented on January 6, 2026.

The 2026 policy expanded the first congestion zone beyond the original downtown area and created a second, geographically separate zone in Hyde Park. Figure 1 shows the original and expanded boundaries. During the analysis period, a $1.50 congestion charge applied when a trip began or ended in a covered zone during the daily active window from 6 a.m. through 10 p.m. The headline design focuses on weekdays at the 10 p.m. endpoint, where otherwise similar trips on opposite sides of the boundary face different assigned tax exposure.

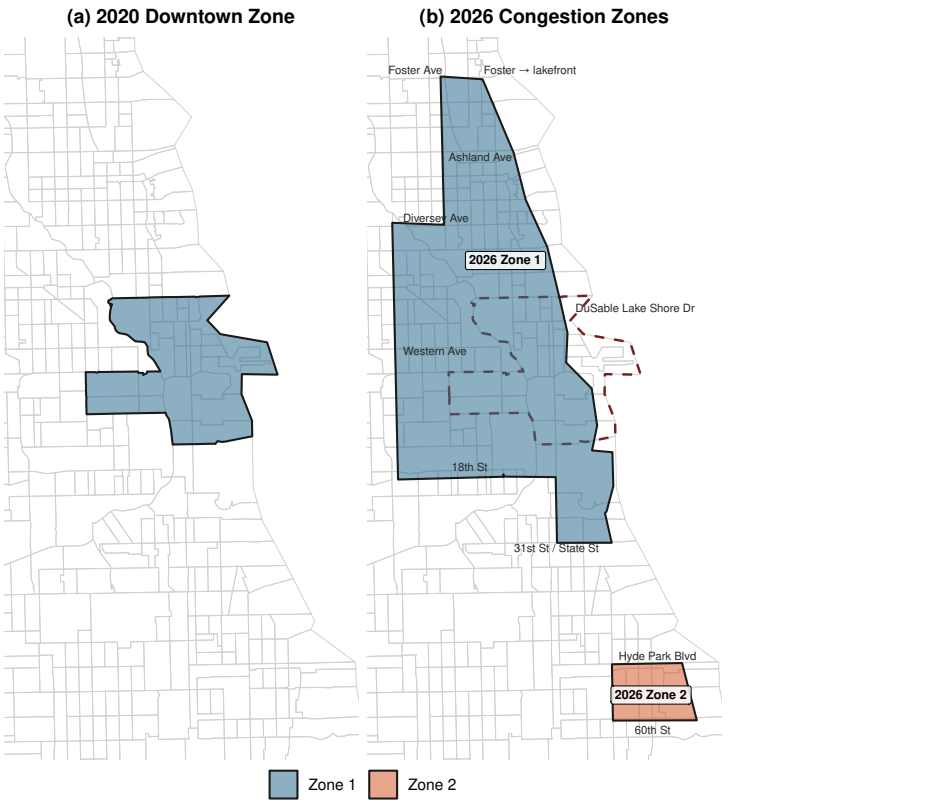
The statutory charge is uniform across the newly covered geography, but the final price need not increase by the same amount everywhere. A passenger's final price contains the underlying fare and reported additional charges. The platform can adjust the fare even when the statutory charge is fixed. The empirical analysis compares urban and suburban trips to determine whether this pricing margin generates different local consumer incidence.

B. Administrative trip records

I use public trip-level records from the Chicago Department of Business Affairs and Consumer Protection. The data cover reported TNP trips in Chicago and identify pickup and drop-off time, pickup and drop-off census tract or community area when disclosure rules permit, trip distance, trip duration, the platform-reported fare, additional charges, and final trip price. Reported timestamps are rounded to 15-minute intervals, and some locations are reported only at the community-area level.

The fare and additional-charge fields must be interpreted carefully. The fare is the underlying platform-reported fare component used in the decomposition below. The additional-charge field can combine statutory taxes and fees with other additions; it is not a pure measure of surge pricing. For this reason, I use the distribution of observed additional

Figure 1. Chicago Congestion-Tax Zones



Notes. Panel (a) shows the January 2020 Downtown Zone. Panel (b) shows Congestion Zone One and Congestion Zone Two following the January 6, 2026 expansion, with the 2020 boundary shown for reference. Boundaries follow Chicago Municipal Code §3-46-020 and the BACP TNP Congestion Zone Boundaries notice. Census-tract outlines are shown for geographic reference. Navy Pier and McCormick Place are excluded from Zone One under the ordinance.

charges to validate reconstructed tax assignment and use the realized additional charge as the endogenous treatment in the fuzzy-RD analysis. Final price is the consumer-facing outcome.

The public records do not include driver identifiers, driver earnings, the division of revenue between a platform and a driver, or platform margins. The analysis can establish whether final prices rise by less than the realized charge and whether the underlying fare adjusts. It cannot attribute incomplete pass-through specifically to drivers or to the platform. The records also do not identify the transportation alternative, if any, chosen by a passenger whose trip does not occur.

The headline pre-March samples run from January 27 through February 27. The urban price sample uses the statutory-grid comparison from 2025, while the suburban sample uses newly charged trips from 2026. Post-March price, fare, and composition estimates use March 6 through June 30, 2026. The trip-volume design instead uses March 2 through June 30 in 2025 and 2026. These sample differences are stated in each results table because the first stage is estimated separately for every outcome and sample.

III. Empirical Strategy

A. The daily tax boundary

The empirical design exploits the end of the congestion-charge window at 10 p.m. on weekdays. For a trip involving a covered zone, the statutory schedule changes discontinuously at this clock boundary. The analysis compares trips assigned to the charged endpoint with trips assigned to the uncharged endpoint in a $\pm$15-minute neighborhood. Results are reported in *charge-on orientation*: the sign is normalized so that a positive first stage denotes an increase in the realized additional charge.

The public timestamps create two assignment complications. First, pickup and drop-off times are rounded to 15-minute bins, so a timestamp alone does not perfectly locate the relevant endpoint on one side of the statutory boundary. I combine the rounded timestamps with reported trip duration to reconstruct the endpoint time and exclude ambiguous boundary observations under the sample rule. Second, a trip is geographically covered when either relevant endpoint lies inside a congestion zone. I reconstruct this status from official zone geometry and retain high-certainty census tracts. These rules determine assignment; they do not condition the analysis sample on the realized charge.

B. Validation of assignment

Tables 1 and 2 compare reconstructed assignment with tax status revealed by clean values of the observed additional-charge field. The time-and-location rule correctly classifies 99.19% of verification trips in both the 2026 and 2020 samples. The geographic rule correctly classifies 99.81% of the 2026 verification sample and 98.70% of the 2020 sample. The approximately 99% classification rates show that time, duration, and official zone geometry recover the statutory assignment with high accuracy.

Table 1. Validation of the Time-and-Location Tax Assignment Rule

	Assigned taxed	Assigned untaxed	Overall
Panel A: 2026 reform			
Correctly classified (%)	99.96	99.08	99.19
Verification trips	2,248	15,353	17,601
Panel B: 2020 reform			
Correctly classified (%)	99.35	98.92	99.19
Verification trips	164,991	96,919	261,910

Notes. The table compares tax status assigned from trip time and endpoint location with the tax status revealed by the reported additional-charge component. The 2026 verification sample covers high-certainty Zone 1 and Zone 2 trips from January 27 through February 27, 2026, including weekends, and conditions on the clean statutory grid: \$2.73 denotes taxed and \$1.23 denotes untaxed. The 2020 verification uses single, non-airport trips after January 6, 2020, with pickup and drop-off tracts classified as fully inside or outside the congestion zone; \$4.83 denotes taxed and \$3.08 denotes untaxed. Because reported timestamps are rounded to 15-minute bins, the 2020 verification excludes trips whose relevant in-zone pickup or drop-off event lies within 7.5 minutes of a statutory clock cutoff.

Figure 2 displays the discontinuity in reported additional charges at the daily boundaries. The visible jump confirms that the statutory endpoint shifts realized charge exposure. Imperfect compliance and the combined nature of the additional-charge field nevertheless mean that the treatment is not deterministic, motivating a fuzzy rather than sharp RD.

C. Fuzzy-RD specification

Let Z_i indicate assignment to the charged side of the 10 p.m. boundary, let C_i denote the additional charge recorded on trip i, and let Y_i denote an outcome. The first stage is

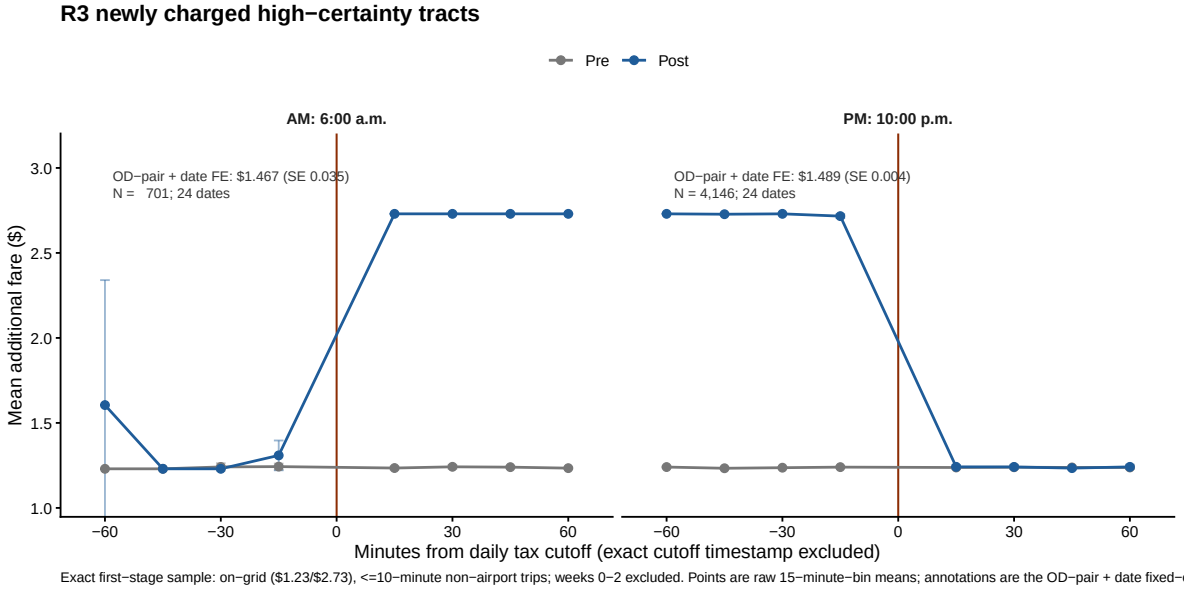
$$C_i = \pi Z_i + \alpha_{d(i)} + \mu_{od(i)} + u_i, \tag{1}$$

Table 2. Validation of the Geographic Tax Assignment Rule

	Assigned taxed	Assigned untaxed	Overall
Panel A: 2026 reform			
Correctly classified (%)	100.00	99.23	99.81
Verification trips	508,454	166,931	675,385
Panel B: 2020 reform			
Correctly classified (%)	99.57	96.59	98.70
Verification trips	419,432	172,867	592,299

Notes. The table validates a geographic assignment constructed solely from official congestion-zone geometry. A tract is classified as inside when at least 99% of its area lies within the applicable congestion zone and as outside when no more than 0.1% lies within it; trips involving an intermediate-overlap tract are excluded. A trip is assigned taxed if either endpoint is inside and assigned untaxed if both endpoints are outside. To isolate geography from clock-time classification, both reconstructed endpoint events must lie at least 7.5 minutes inside the active 6 a.m.–10 p.m. interval. Observed status is measured on the clean additional-charge grids: \$2.73/\$1.23 in 2026 and \$4.83/\$3.08 in 2020. The 2026 sample covers all days from January 27 through February 27; the 2020 sample uses citywide single, non-airport trips after January 6. The geographic rule does not use the empirically estimated tract-compliance tiers, avoiding a circular validation against charges.

Figure 2. Realized Additional Charges at the Daily Tax Boundaries



Notes. Mean additional charges by 15-minute clock bin around the daily charge boundaries for high-certainty 2026 geography. The figure documents the first-stage discontinuity in the observed additional-charge field. The headline estimates compare trips picked up within a fixed window on either side of the weekday 10 p.m. boundary, excluding the 10 p.m. bin itself, and report effects in charge-on orientation.

where $\alpha_{d(i)}$ denotes pickup-date fixed effects and $\mu_{od(i)}$ denotes origin–destination-pair fixed effects when applicable. The second stage is

$$Y_i = \beta \widehat{C}_i + \alpha_{d(i)} + \mu_{od(i)} + \varepsilon_i. \tag{2}$$

The coefficient β is the local effect of one dollar of realized additional charge for trips whose charge is shifted by cutoff assignment. For final price, it is the pass-through rate. For the fare, distance, and duration outcomes, it is the change in that outcome per dollar of realized charge. In the preferred urban–suburban test, standard errors are clustered at the pickup-date × OD-pair level, allowing arbitrary dependence among trips serving the same OD market on the same date. Other specifications use the clustering level indicated in their table notes. Price and fare outcomes are adjusted for miles and duration as indicated in the table notes.

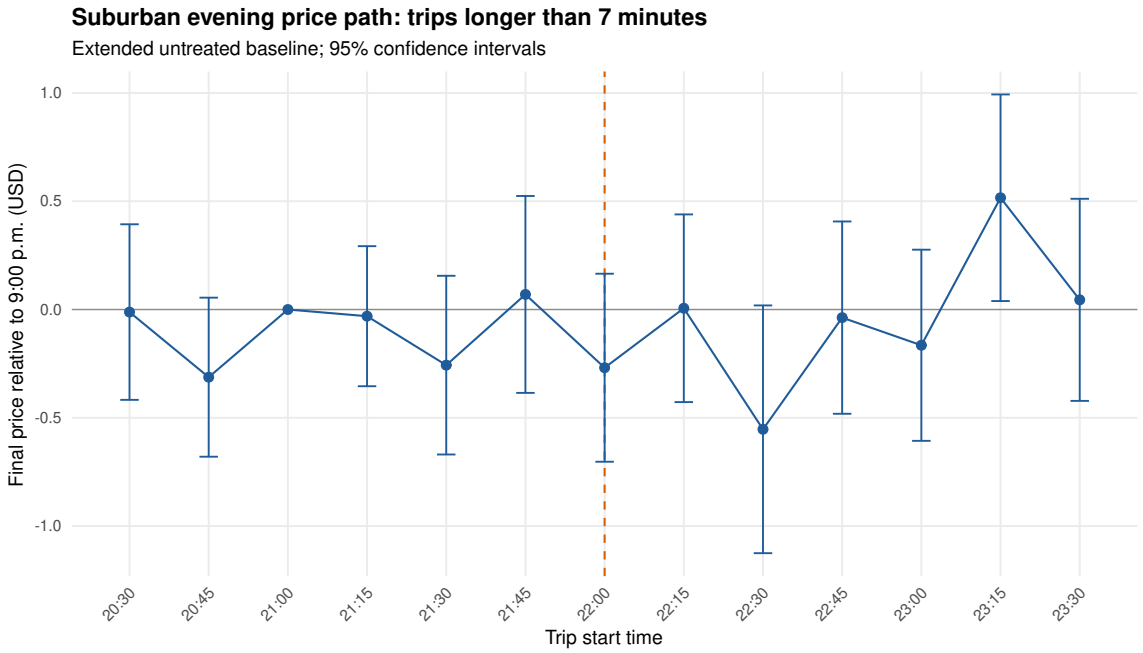
The estimation window is a fixed number of clock minutes on either side of the boundary, with ±30 minutes as the preferred bandwidth and a ladder reported alongside it. Because reported timestamps are rounded to fifteen-minute bins, a trip stamped 10 p.m. may have begun on either side of the cutoff; after March the mean additional charge in that bin sits almost exactly midway between the adjacent taxed and untaxed bins. The 10 p.m. bin is therefore excluded throughout, so that charge assignment is unambiguous for every trip that enters.

Trip volume is estimated with a Poisson difference-in-discontinuities model on a zero-filled pickup-date × clock-minute × OD-pair panel. The estimator compares the 2026 change in the count discontinuity at the daily boundary with the discontinuity over the same calendar dates in 2025. I estimate it separately before and after the March 2026 pricing change. Because none of the spatial-market estimates is statistically distinguishable from zero, I treat volume as a null secondary outcome and do not construct a demand elasticity.

The identifying assumption is that, absent the change in charge assignment, potential outcomes would evolve smoothly across the 10 p.m. boundary. The comparison is local: it does not identify effects for other hours or locations. Figure 3 shows the pre-policy evening price path after the duration selection rule. There is no comparable mechanical jump at 10 p.m. before the 2026 expansion in the suburban sample. This diagnostic supports the use of the administrative boundary while leaving the usual RD smoothness assumption explicit.

The preferred spatial comparison stacks urban and suburban trips from the common January 27–February 27, 2026 window. Both samples use the same clock-time window, duration

Figure 3. Pre-Policy Evening Price Path



Notes. The figure reports the 2025 evening price path for suburban trips after the duration selection rule. Estimates are relative to the 9 p.m. reference bin and include date-by-pickup-tract and date-by-drop-off-tract fixed effects. It is a diagnostic for a mechanical pre-policy discontinuity, not an alternative treatment-effect estimator.

restriction, price adjustment, date and origin–destination fixed effects, and pickup-date × OD-pair-clustered inference. I instrument the realized charge and its interaction with a suburban indicator using cutoff assignment and its corresponding interaction. The interaction coefficient directly tests equality of urban and suburban pass-through. The preferred half-window is 30 minutes; I report 15- and 60-minute estimates to show sensitivity to local sample size and clock-time support. This specification changes geography while holding the calendar period and pricing regime fixed.

The empirical hierarchy is as follows. The contemporaneous interacted fuzzy RD is the preferred test of spatial heterogeneity. The community-area OD decomposition and trip-attribute estimates then distinguish geographic composition from changes in observed trip characteristics. Trip volume is a separate reduced-form Poisson difference-in-discontinuities outcome and is not divided by a charge first stage. Post-March estimates describe a different pricing regime, but comparisons across regimes do not identify the causal effect of the pricing change.

IV. Results

A. A contemporaneous test of spatial heterogeneity

Table 3 reports the preferred urban–suburban comparison and its bandwidth ladder. Every row uses January 27–February 27, 2026, the weekday 10 p.m. cutoff, common trip restrictions and price adjustment, and common fixed effects. The joint IV specification instruments the realized charge separately by geography and estimates both pass-through coefficients together.

In the preferred ±30-minute window, pass-through is 75.2% for urban trips and 60.6% for suburban trips. Both are precisely different from zero. The suburban-minus-urban difference is −14.5 percentage points with a standard error of 15.0 points ($p = 0.33$). The ±15-minute window preserves this ordering. At ±60 minutes the estimates are 98.3% and 54.8%, producing a 43.5-point difference that is significant under this clustering ($p = 0.006$) but not under date-level clustering ($p = 0.09$). I interpret the table as strong evidence of incomplete pass-through in both markets and as inconclusive on a spatial gradient.

Table 3. Contemporaneous Urban–Suburban Pass-Through Test

Half-window	Urban	Suburban	Difference	Difference SE	p-value
±15 minutes	65.8% (8.9)	50.7% (13.3)	−15.1 pp	16.0 pp	0.34
±30 minutes	75.2% (10.6)	60.6% (10.5)	−14.5 pp	15.0 pp	0.33
±60 minutes	98.3% (12.6)	54.8% (9.6)	−43.5 pp	16.0 pp	0.006

Notes. Joint fuzzy RD around the weekday 10 p.m. cutoff. Entries in the urban and suburban columns are pass-through estimates; standard errors in percentage points are in parentheses. Both markets use January 27–February 27, 2026, non-airport trips lasting at most ten minutes, a common adjustment of final price for miles and duration, and pickup-date and community-area OD-pair fixed effects. Cutoff assignment and its interaction with the suburban indicator instrument for the realized charge and its interaction. Standard errors are clustered by pickup-date × community-area OD pair. The 15-, 30-, and 60-minute samples contain 17,790, 35,409, and 69,907 usable observations and 1,381, 1,779, and 2,225 clusters, respectively. Difference is suburban minus urban.

The larger gap at ±60 minutes could reflect either greater precision from additional trips or different pricing conditions farther from the cutoff. To distinguish these explanations, I repeat the bandwidth ladder over the same January–February dates in 2025. The urban sample is comparably large in 2025, but its pass-through estimate does not rise with the bandwidth: it is 16.7% at ±15 minutes, 49.1% at ±30 minutes, and 29.9% at ±60 minutes. The sharp increase in the 2026 urban coefficient therefore cannot be attributed to sample size alone. Because it appears when the outer 30–60-minute observations enter, I retain ±30

minutes as the preferred local specification and treat ±60 minutes as a sensitivity result.

The null on the gap is not an artifact of the level at which I cluster. Table 4 reports the same ±30- and ±60-minute contrasts under four clustering choices. The point estimates are identical by construction; only the standard errors move. At ±30 minutes the difference is insignificant under all four, with p between 0.16 and 0.54. At ±60 minutes the difference is significant when clustering on date × OD pair or on the OD pair, but not when clustering on the pickup date, which is the level at which the running variable actually varies. This reinforces treating the ±60-minute gap as a sensitivity result rather than a finding.

Table 4. Urban–Suburban Difference under Alternative Clustering

Half-window	Clustering	Clusters	Difference SE	p-value
±30 minutes	date × OD pair	1,779	15.0 pp	0.33
	OD pair	172	10.4 pp	0.16
	date	24	23.5 pp	0.54
	two-way date + OD	24	20.6 pp	0.49
±60 minutes	date × OD pair	2,225	16.0 pp	0.006
	OD pair	205	14.2 pp	0.002
	date	24	24.6 pp	0.09
	two-way date + OD	24	23.4 pp	0.08

Notes. The suburban-minus-urban pass-through difference of Table 3 re-evaluated under four clustering levels. The estimate is −14.5 percentage points at ±30 minutes and −43.5 at ±60 minutes in every row; the sample, estimator, and fixed effects are unchanged and only the covariance estimator differs. The two-way rows use Cameron–Gelbach–Miller with the reference distribution set by the smaller of the two cluster counts.

B. Where the price response comes from

Two explanations compete for an incomplete price response: the mix of local travel markets may change at the cutoff, or the platform may adjust the fare it charges for a comparable trip. Table 5 separates them on the sample and instrument of Table 3.

Community-area OD composition explains essentially none of the price discontinuity. Adding community-area OD-pair fixed effects moves the pre-March suburban discontinuity from \$0.797 to \$0.784 and the urban discontinuity from \$1.047 to \$1.050. Composition therefore accounts for under two percent of either response, and the price response occurs within local travel markets rather than across them.

The platform's own fare is the adjusting margin. Instrumenting the realized charge, the platform-reported base fare falls by \$0.285 per dollar of charge on urban trips and by \$0.295

on suburban trips before March, and by \$0.655 and \$0.592 respectively after March. These offsets account for the pass-through estimates almost exactly: one minus the urban offset is 0.715 before March and 0.345 after, against measured pass-through of 75.2% and 39.1%. The near-doubling of the fare offset between regimes is therefore the proximate reason pass-through fell, and it is present in both markets rather than in the suburban market alone.

Observed trip attributes move very little. Distance changes by at most 0.03 miles per dollar of charge and duration by at most 0.27 minutes, or sixteen seconds. Several duration coefficients are statistically distinguishable from zero, but they are far too small to generate price movements of the size in the first two columns, so the response is a change in the price of a comparable trip rather than a shift toward different trips. The records contain no surge indicator and do not separately identify a fixed base-fare component, so I cannot say which part of the fare moves; variation in surge or other dynamic-pricing components is one possible channel among several. These intensive-margin checks also cannot rule out selection in which trips occur.

Table 5. Price Response, Fare Offset, and Trip Attributes

		Price discontinuity (\$)		Per dollar of realized charge		
Market	Regime	Raw	Within OD	Base fare (\$)	Distance (mi)	Duration (min)
Urban	Pre-March	1.047 (0.151)	1.050 (0.152)	−0.285 (0.102)	−0.010 (0.005)	0.070 (0.015)
Suburban	Pre-March	0.797 (0.165)	0.784 (0.152)	−0.295 (0.086)	0.007 (0.012)	0.109 (0.045)
Urban	Post-March	0.422 (0.122)	0.435 (0.122)	−0.655 (0.092)	−0.018 (0.004)	0.075 (0.012)
Suburban	Post-March	0.078 (0.128)	0.256 (0.121)	−0.592 (0.165)	0.030 (0.012)	0.272 (0.061)

Notes. Standard errors clustered by pickup-date × community-area OD pair are in parentheses. The price-discontinuity columns are reduced-form jumps in final price adjusted for miles and duration, with pickup-date fixed effects only (Raw) and with pickup-date and community-area OD-pair fixed effects (Within OD). The last three columns instrument the realized charge with cutoff assignment and include both sets of fixed effects. All rows use ±30 minutes around the weekday 10 p.m. cutoff, non-airport trips of at most ten minutes, and exclude the ambiguous 10 p.m. pickup bin; both markets use 2026 dates in both regimes. Sample sizes are 31,657 and 3,781 trips before March and 104,644 and 17,422 after. The suburban post-March raw discontinuity is close to zero, so the ratio of its two price columns is not informative and is not reported.

C. Effects on trip volume

Table 6 reports trip-volume effects separately before and after the March 2026 pricing change. I estimate Poisson difference-in-discontinuities models on zero-filled pickup-date × clock-minute × OD-pair panels. Before March, the estimated effect is a 2.0% decline in urban volume and a 1.8% decline in suburban volume. After March, the corresponding estimates are declines of 0.1% and 1.3%. None is statistically distinguishable from zero. The results

therefore do not establish a larger suburban contraction, and their imprecision makes demand elasticities uninformative.

Table 6. Trip-Volume Effects Before and After the March 2026 Pricing Change

Fee regime	Urban	Suburban
Pre-March	−2.0% (3.6 pp)	−1.8% (4.4 pp)
Post-March	−0.1% (1.9 pp)	−1.3% (2.4 pp)

Notes. Percentage effects from Poisson difference-in-discontinuities regressions at the weekday 10 p.m. endpoint; delta-method standard errors are in parentheses. The pre-March comparison is January 27–February 27, 2026 versus the same dates in 2025. The post-March comparison is March 2–June 30, 2026 versus the same dates in 2025. Regressions include calendar-month and OD-pair fixed effects, cluster by pickup date, and exclude trips longer than ten minutes.

D. Pass-through after the March 2026 pricing change

The platform's observed pricing structure changed in March 2026. Table 7 re-estimates pass-through on each side of that break with the estimator of Table 3: the same joint IV, trip restrictions, price adjustment, fixed effects, and clustering, with only the date window moved. The pre-March window is January 27–February 27, 2026 and the post-March window is March 6–June 30, 2026. Both exclude the 10 p.m. pickup bin, which fifteen-minute timestamp rounding leaves ambiguous between the taxed and untaxed sides.

Pass-through is substantially lower in the later regime. In the ±30-minute window the urban estimate falls from 75.2% to 39.1% and the suburban estimate from 60.6% to 18.5%. Estimating both regimes jointly, so that the contrast is a linear restriction within one covariance matrix rather than a comparison of two separate fits, the urban decline is 36.1 percentage points with a standard error of 14.3 points ($p = 0.012$). The suburban decline is 42.2 points but is not precisely estimated (standard error 27.3 points; $p = 0.12$). At ±60 minutes the urban decline is 51.4 points ($p < 0.001$). The instrument remains strong in both markets after March, with first-stage F statistics of 1,014 for urban trips and 53 for suburban trips at ±30 minutes.

The spatial comparison remains null in the later regime. The post-March suburban-minus-urban difference is −20.6 percentage points with a standard error of 27.3 points ($p = 0.45$). Neither regime therefore yields a statistically distinguishable urban–suburban gap, while the change over time within each market is large and, for urban trips, precisely estimated.

The decline coincides with a change in the platform's own price schedule rather than in the statutory charge, which was unchanged at $1.50 throughout. Table 8 shows that over

Table 7. Pass-Through Before and After the March 2026 Pricing Change

Half-window	Market	Pre-March	Post-March	Change	p-value
±30 minutes	Urban	75.2% (10.6)	39.1% (9.7)	−36.1 pp (14.3)	0.012
	Suburban	60.6% (10.5)	18.5% (25.2)	−42.2 pp (27.3)	0.12
±60 minutes	Urban	98.3% (12.6)	46.9% (9.1)	−51.4 pp (15.5)	0.001
	Suburban	54.8% (9.6)	35.6% (16.2)	−19.2 pp (18.8)	0.31

Notes. Pass-through estimates from the joint fuzzy RD of Table 3; standard errors in percentage points are in parentheses. The pre-March window is January 27–February 27, 2026 and the post-March window is March 6–June 30, 2026. Both regimes use non-airport trips lasting at most ten minutes, a common adjustment of final price for miles and duration applied within regime, and pickup-date and community-area OD-pair fixed effects; the 10 p.m. pickup bin is excluded throughout. The Change column and its p-value come from a single pooled specification in which the realized charge is instrumented separately by market and regime, with regime-specific community-area OD-pair effects, so the contrast is tested within one covariance matrix. Standard errors are clustered by pickup-date × community-area OD pair. The ±30-minute samples contain 35,409 pre-March and 122,037 post-March observations; the ±60-minute samples contain 69,907 and 241,125.

the post-March window the booking fee the platform adds beyond the statutory charge rose from $2.85 to $6.22 per trip on average, while the share of trips paying nothing beyond the statutory charge fell from 8.6% to 0.9%. The share of the statutory charge reaching consumer prices fell as the platform's own fee grew, which is what one expects when a platform can absorb part of a charge into a second price component it controls.

Table 8. The Platform's Booking Fee After March 2026

Month	Trips	Mean fee	Median fee	Share paying no fee
March 2026	694,416	$2.85	$2.30	8.6%
April 2026	510,172	$3.68	$3.07	7.2%
May 2026	447,855	$4.21	$3.54	6.6%
June 2026	488,184	$6.22	$5.19	0.9%

Notes. The booking fee is the portion of reported additional charges in excess of the statutory congestion charge. Monthly figures are trip-weighted averages of weekly estimates. "Share paying no fee" is the fraction of trips whose additional charges equal the statutory amount to within half a cent. Trips within fifteen minutes of the 6 a.m. and 10 p.m. cutoffs are excluded, because whether the statutory charge applies is ambiguous there and the fee is defined by subtracting it.

These estimates do not identify the causal effect of the pricing change itself. The regime is not randomly assigned, and the post-March window spans four months over which the platform's fee schedule continued to move; pickup-date fixed effects absorb that drift in the level, but the comparison remains a contrast between two periods rather than an experiment. What the table establishes is narrower: measured on one estimator, consumer pass-through of an unchanged statutory charge was significantly lower after the platform restructured its

pricing than before.

V. Conclusion

This paper finds evidence of incomplete consumer pass-through and adjustment in the platform-reported fare, but not a conventionally significant urban–suburban difference in the preferred specification. In the contemporaneous ±30-minute window, pass-through is estimated at 75.2% for urban trips and 60.6% for suburban trips. The direct suburban-minus-urban contrast is −14.5 percentage points with a standard error of 15.0 points ($p = 0.33$). The ordering is stable across bandwidths, but the larger, statistically significant ±60-minute gap is driven by the outer clock-time support rather than sample size alone. The evidence for a spatial gradient is therefore suggestive and bandwidth-sensitive.

The decomposition points to pricing within local travel markets rather than to their composition or to changes in observed trip attributes. Community-area OD composition accounts for under two percent of the price discontinuity in either market. The adjusting margin is the platform's reported base fare, which falls by roughly \$0.29 per dollar of realized charge in both markets before March and by \$0.59 to \$0.66 after March. Distance and duration move by at most 0.03 miles and sixteen seconds per dollar, too little to account for the price response.

Trip-volume responses are modest and statistically insignificant in both pricing regimes. Before March, urban volume falls by 2.0% and suburban volume by 1.8%; after March, the corresponding declines are 0.1% and 1.3%. On a single estimator, however, price pass-through is substantially lower after March: urban pass-through falls from 75.2% to 39.1% ($p = 0.012$ for the change) and suburban pass-through from 60.6% to 18.5% ($p = 0.12$). The spatial contrast remains statistically indistinguishable in both regimes, so the regime comparison establishes a change over time rather than spatial divergence or convergence.

The broader implication is narrower but more credible than a claim of geographic heterogeneity. The statutory charge need not be added one-for-one to the consumer price because the platform can adjust another price component. Determining whether that adjustment differs across local markets requires a contemporaneous, directly tested contrast; in this setting, the available contrast is too imprecise to establish such a difference.

References

Agrawal, D. R. and Zhao, W. (2023). Taxing Uber. *Journal of Public Economics*, 221:104862.

Almagro, M., Barbieri, F., Castillo, J. C., Hickok, N. G., and Salz, T. (2024). Optimal urban transportation policy: Evidence from chicago. NBER Working Paper 32185, National Bureau of Economic Research. Revised July 2026; revise and resubmit at Econometrica.

Anas, A. (2012). The optimal pricing, finance and supply of urban transportation in general equilibrium: A theoretical exposition. *Economics of Transportation*, 1(1–2):64–76.

Cairncross, F., Hall, J., and Palsson, C. (2021). Vancuber: The effect of ride-hailing on public transportation, congestion, and traffic fatalities. Working Paper.

Cook, C., Kreidieh, A., Vasserman, S., Allcott, H., Arora, N., Tomkins, A., Turkel, E., and van Sambeek, F. (2025). The network-wide effects of congestion pricing: Evidence from new york city. NBER Working Paper 33584, National Bureau of Economic Research. Revised August 2026.

de Palma, A. and Lindsey, R. (2007). Transport user charges and cost recovery. *Research in Transportation Economics*, 19:29–57.

Leccese, M. (2024). Taxing ride-sharing: Which neighborhoods pay more? *Journal of Regional Science*, 64(4):1393–1413.

Lindsey, R. (2012). Road pricing and investment. *Economics of Transportation*, 1(1–2):49–63.

Parry, I. W. H. (2009). Pricing urban congestion. *Annual Review of Resource Economics*, 1:461–484.

Small, K. A. (1992). Using the revenues from congestion pricing. *Transportation*, 19(4):359–381.

Tarduno, M. (2021). The congestion costs of uber and lyft. *Journal of Urban Economics*, 122.

Tarduno, M. (2025). Elasticities and tax incidence in urban ridesharing markets: Evidence from chicago. *Journal of Urban Economics*, 150:103809.