Equitable micromobility deployment mandates show benefits only when spatially targeted

Hale Erkan[1,*] and Ioannis Stamatopoulos[2]

Abstract

To reduce mobility inequities, governments require private operators to meet minimum visit frequencies, trip counts, or vehicle deployments in disadvantaged neighborhoods. Shared scooters are a case in point: by 2024, 36 United States programs had adopted minimum deployment thresholds. We evaluate 15 programs in 13 cities (2019–2023) with spatial regression discontinuity, analyzing 66 million vehicle snapshots and telemetry from a large operator. Most mandates cover large areas and have null effects on availability—not from noncompliance, but because thresholds rarely constrain deployment. But the two cities that use targeted, pocket-based mandate zones centered on transit stops—Baltimore and Denver—increase availability by +38.38 scooters per km^2 ($p < 0.01$) and +70.30 scooters per km^2 ($p < 0.01$), respectively, closing 98% and 85% of the preexisting gap between mandate areas and median central neighborhoods. In Baltimore, where we also observe ridership from a large operator, the mandate causes an additional +26.47 rides per km^2 ($p = 0.05$), closing 35% of the corresponding preexisting ridership gap.

1 Introduction

Structural mobility inequalities in the United States are stark. In the five largest cities—home to more than 50 million urban residents—historically disadvantaged neighborhoods have 56% fewer public transit stops per km^2 than central neighborhoods. As a result, a routine grocery trip by public transit takes 33% longer in a historically disadvantaged neighborhood. These inequalities persist in 2026 despite substantial public investment to address them, including $204 billion through the Bipartisan Infrastructure Law.[1]

One policy response has been to mandate minimum service levels (e.g., minimum bus frequencies or minimum taxi rides) in disadvantaged neighborhoods. Shared micromobility—electric scooters and bikes, together generating over 225 million trips in 2024 alone—has become a prominent domain for this policy response, with 36 regulatory programs across 34 United States cities by 2024 mandating minimum-deployment thresholds in disadvantaged areas.[2] Because these mandates apply within ad hoc administrative boundaries that often cut through disadvantaged neighborhoods, they create natural experiments for evaluating policy effectiveness by comparing outcomes on either side of the boundaries.

We use these boundary discontinuities to evaluate mandate effects on device availability in 11 cities—Atlanta, Baltimore, Chicago, Denver, Kansas City, Los Angeles, Louisville, Portland, Sacramento, San Francisco, and Seattle—covering 61% of the population in all mandate cities across the United States between 2020 and 2021. Further, using operational data from one of the largest active United States micromobility operators between 2020 and 2023, we evaluate mandate effects on ridership and service quality in 7 cities—Baltimore, Chicago, Washington, DC, Los Angeles, Minneapolis, Portland, and San Francisco, increasing our study's coverage to 64%. Our results provide the first large-scale evidence on whether minimum service mandates in urban transportation narrow structural mobility gaps.

We find that mandates show benefits only when targeted. On one hand, most mandates produce no detectable improvements in availability or ridership, not because operators fail to comply, but because broad mandate zones let them satisfy aggregate fleet-share thresholds by concentrating vehicles in already high-demand areas. On the other hand, Baltimore and Denver use small mandate zones centered on transit stops and increase availability by +38.38 and +70.30 scooters per km^2, respectively ($p < 0.01$ for both), closing 98% and 85% of the preexisting gap between mandate areas and median central neighborhoods. In Baltimore, where operator telemetry also lets us measure ridership, the mandate causes an additional +26.47 rides per km^2 ($p = 0.05$), closing 35% of the corresponding preexisting ridership gap.

2 Related literature and contribution

Shared micromobility supply is highly concentrated for two reasons. First, returns to deployment are highest in a few dense urban areas. For example, He et al. (2021) estimate that adding a bike-sharing station to central London's already thick network would generate up to ten times more ridership than adding a station to the periphery's thin network. Second, competition further intensifies that concentration. Erkan et al. (2025) find that scooter platforms in Kansas City generate a 531% oversupply relative to the social optimum as they compete for customer attention. Rosaia (2025) finds a parallel pattern in New York City ride-hailing, where competition between platforms leaves roughly 19% more driver idle time than a single operator would. Although this concentration is a natural market outcome, it conflicts with regulators' efforts to expand transportation access in historically underserved communities.

In response, a growing operations management literature studies policy interventions that can offset this concentration. Siddiq et al. (2022) argue that, in theory, subsidies for mixed-mode trips that connect underserved riders to public transit could increase transit adoption at zero government cost. Wu and Zhang (2025) estimate that equalizing bus frequencies across low- and high-income neighborhoods in Chicago would raise ridership in low-income areas by about 2 percentage points but lower it in high-income areas by about 5. Finally, Xu et al. (2025) provide a conceptual rationale for minimum-deployment mandates by arguing that observed demand in low-supply areas may understate true need. However, whether such policy interventions improve access and ridership in underserved neighborhoods remains an open empirical question. We provide the first large-scale empirical evidence on that question.

Methodologically, we bring spatial regression discontinuity to operations management. Regression discontinuity is a standard tool for policy evaluation in economics and political science; see, for example, Imbens and Lemieux (2008), Lee and Lemieux (2010), and Cattaneo and Titiunik (2022) for reviews. In operations, most applications exploit scalar cutoffs instead of geographic boundaries. For example, Calvo et al. (2019) use a federal budget threshold that triggers oversight of public projects, and Kadiyala et al. (2024) exploit a retailer policy under which the probability of receiving a gift-card promotion email jumps at purchase-recency thresholds. The closest design to ours is Kaaua and Virudachalam (2025), who identify the effect of commute distance by comparing nearby legislative districts on opposite sides of a state border. Their strategy shares the geographic intuition of spatial regression discontinuity, but identification comes from cross-border differences in commute distance rather than from a discontinuous change in treatment at a geographic boundary. In our setting, treatment changes discontinuously at the borders of equity zones. To our knowledge, ours is the first in operations management to apply a spatial regression discontinuity design.

3 Data

Equitable deployment mandates

Cities regulate shared micromobility through programs that define service areas, no-ride and low-speed zones, set fleet caps and rebalancing requirements, levy per-device fees, and require operators to share their data. These programs are enforced through fines and reviewed every 6–12 months. To date, there have been 239 such programs: 148 included access provisions, such as payment without a smartphone; 78 of these imposed minimum or maximum deployment thresholds; and 36 required minimum deployment in historically disadvantaged neighborhoods—henceforth, equitable deployment mandates (Brown and Howell, 2024).

We study 15 distinct programs with equitable deployment mandates across 13 cities in the United States. For our primary analysis we focus on 11 programs, spanning 11 cities between 2020 and 2021, for which we observe real-time scooter availability for all active operators (see programs marked with "+" on Table 1, and Figure 1). For a secondary analysis we focus on 9 programs, spanning 7 cities between 2020 and 2023, for which we observe detailed operational data from one of the largest active micromobility operators in the United States (programs marked with "∗" on Table 1).

Each mandate establishes geographic boundaries intended to encompass historically disadvantaged neighborhoods and sets minimum deployment thresholds within them. Boundaries often cut arbitrarily through neighborhoods (we demonstrate this point in Section 5), and may define a single contiguous area or multiple discrete subareas. Thresholds can vary across subareas. We extract all boundaries, areas/subareas, and thresholds as well as audit cadence and enforcement mechanisms for each equitable deployment mandate directly from municipal permits and compliance evaluations (See Appendix C).

Program	Deployment thresholds	Audit cadence & enforcement	Dates	% of city in mandate area
ATL 2021^{+}	3 subareas, each ≥2% of fleet	monthly; revocation	6/21–12/21	17.2
BAL 2021^{+*}	12 subareas, each 3 devices	daily (5am–9am); revocation	6/21–8/21	1.8
CHI 2020^{+*}	20 subareas, each ≥2.5% of fleet; jointly ≥50% of fleet	daily, twice; fines/suspension/revocation	8/20–12/20	43.0
CHI 2022*	10 subareas, each ≥4% of fleet; jointly ≥50% of fleet	daily, twice; fines/suspension/revocation	7/22–10/22	45.6
DC 2020*	18 subareas; jointly ≥400 devices	daily at 6am; revocation	8/20–12/20	15.1
DEN 2021^{+}	25 subareas, jointly ≥30% of fleet	daily; fines/fleet reduction/revocation	6/21–12/21	15.4
KC 2021^{+}	6 subareas, jointly ≥10% of fleet	NA; deployment restrictions	6/21–12/21	22.1
LA 2021^{+*}	≥20% of fleet	NA; suspension/revocation/non-renewal	6/21–12/21	13.7
LOU 2021^{+}	2 subareas, jointly ≥20% of fleet	NA; fines/remove devices	6/21–8/21	18.9
MSP 2023*	5 subareas, ≥2%, 7%, 7%, 7%, 8% of fleet	weekly & monthly; fines/suspension/revocation	5/23–11/23	6.4
PDX 2019*	≥20% of fleet	daily; fines/suspension/revocation	8/19–12/19	19.0
PDX 2021^{+*}	≥20% of fleet	daily; fines/suspension/revocation	8/21–12/21	19.0
SAC 2021^{+}	9 subareas, jointly ≥20% of fleet	daily; fleet reduction/suspension	6/21–12/21	51.8
SF 2021^{+*}	4 subareas: ≥5, 19, 47, 51 devices	hourly; fleet reduction/suspension/revocation	6/21–12/21	26.8
SEA 2021^{+}	4 subareas, ≥10% of fleet	NA; fleet reduction/revocation	8/21–12/21	28.6

Table 1. A plus sign (+) denotes availability data compiled for all operators; an asterisk (∗) denotes availability and downstream outcome data compiled for one of the largest active micromobility operators in the United States between 2020 and 2023.

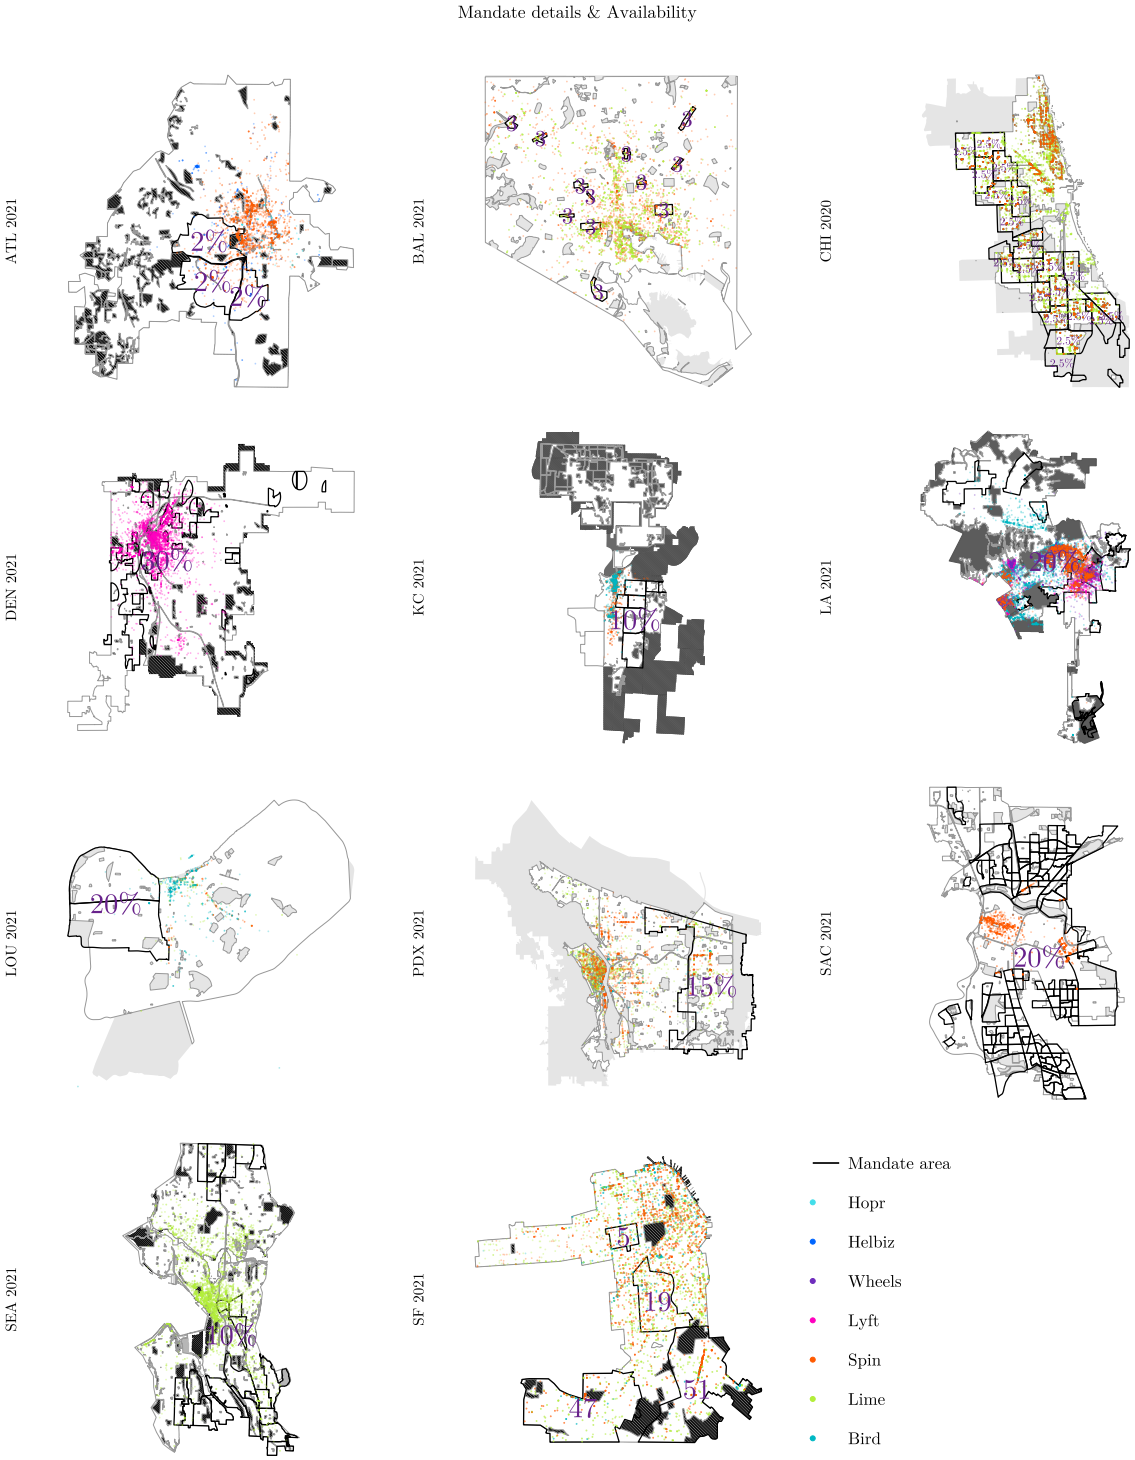


Figure 1. Mandate boundaries, areas/subareas, and 10 a.m. vehicle availability on each program's median deployment date. Excludes CHI 2020, DC 2020, MSP 2023, and PDX 2019, for which real-time scooter location data are unavailable. Availability looks qualitatively similar across dates and times of day (Figure A.1).

Real-time scooter locations

For our primary analysis, we use real-time, scooter-location data. We collected these data by querying, every 10 minutes, the publicly available APIs of all active operators in the cities we

study—Bird, Helbiz, Hopr, Lime, Lyft, Skip, Spin, Razor, and Wheels, together accounting for roughly 90% of the United States shared-scooter fleet during our study period.[3] (This continuous public oversight is unlike otherwise similar settings, such as the Italian minimum-service bus mandates in Lodi et al., 2016.) Each API pull yields a timestamped snapshot of all available devices in a city, including vehicle identifiers and GPS coordinates (latitude and longitude). Because operators frequently rotate vehicle identifiers—often every few minutes—we cannot reliably track individual scooters over time. Across all 11 programs, the raw data include 66,268,805 snapshots and 53,629 vehicles.

Status updates

For a secondary analysis, we use a proprietary dataset from a United States electric scooter operator with an average fleet share of 30% across the 7 cities we study. Each observation is a status update with a timestamp, vehicle ID, latitude and longitude, battery level, type (deployment, ride start, ride end, low battery, maintenance pickup), and, when applicable, a ride ID. We observe deployment updates when scooters are released to the street; user pick-up/drop-off updates when rides start and end; low battery updates when charge falls below 30%; and maintenance pickup updates when scooters are retrieved by the operator for servicing, charging, or rebalancing. Linking all updates over time, we reconstruct complete on-street histories which underpin all outcomes.

For illustration, a typical status sequence for one scooter in CHI 2020 on September 21, 2020, is as follows:

Update type	Lat	Long	Date	Time	Ride-id	Battery (%)
deployment	41.939	-87.754	20-09-21	03:13:10	-	92
ride_start	41.939	-87.754	20-09-21	22:41:24	A	88
ride_end	41.939	-87.756	20-09-21	22:57:52	A	74
ride_start	41.939	-87.756	20-09-22	11:55:57	B	72
ride_end	41.934	-87.742	20-09-22	12:02:00	B	67
ride_start	41.934	-87.742	20-09-22	18:35:26	C	67
ride_end	41.939	-87.756	20-09-22	18:43:43	C	55
ride_start	41.939	-87.756	20-09-25	13:10:09	D	50
ride_end	41.934	-87.742	20-09-25	13:16:42	D	36
low_battery	41.934	-87.742	20-09-29	02:38:48	-	30
maint_pick_up	41.934	-87.742	20-09-29	03:45:29	-	30
deployment	41.880	-87.662	20-09-30	16:03:36	-	92

Across all nine programs, between April 30, 2019 and December 26, 2023, these raw data comprise 6.47 million status updates, including 2.15 million rides, 1.65 million deployments and maintenance pickups, and 0.52 million low-battery signals.

Sample construction

To construct our final scooter-location sample we apply two filters. First, we exclude observations in non-rideable areas (e.g., water, highways, rail corridors, parks). Second, after computing the availability outcome (as described in Section 4) we trim the observations associated with the top

and bottom 1% of this outcome. Our final scooter-location sample includes 56,834,432 vehicle snapshots. Breaking it down by program, our final sample includes: 1,664,704 vehicle snapshots in ATL 2021; 1,447,661 in BAL 2021; 29,015,892 in CHI 2020; 1,889,236 in DEN 2021; 1,899,809 in KC 2021; 12,228,261 in LA 2021; 167,154 in LOU 2021; 2,306,303 in PDX 2021; 787,456 in SAC 2021; 2,294,273 in SEA 2021; and 3,133,683 in SF 2021.

To construct the final status-update sample, we apply three filters. The first two match those used for the scooter-location sample, except that we trim according to all outcomes rather than availability alone. We also exclude March–June 2020, a period of acute COVID-19 disruption, and retain data from July onward, when ridership returned to pre-pandemic levels. Our final status-update sample includes 4.60 million status updates, including 1.28 million rides, 1.54 million deployments and maintenance pickups, and 0.48 million low-battery signals. Breaking it down by program, our final sample includes: 834,911 status updates in BAL 2021; 525,058 in CHI 2020; 303,330 in CHI 2022; 603,836 in DC 2020; 665,568 in LA 2021; 247,215 in MSP 2023; 469,876 in SF 2021; 655,572 in PDX 2019; and 297,126 in PDX 2021.

4 Study design

Outcomes

Cities primarily evaluate shared micromobility programs through device availability, which we use for our primary analysis. Cities also target improvements in ridership and essential ridership—trips to and from essential services such as public transit stops and healthcare facilities. In a secondary analysis, we add ridership and essential ridership, as well as revenue per scooter, rebalancing per scooter, battery levels, and reaction time to low battery signals, while continuing to examine availability. We measure all outcomes daily for an arbitrary area within each city.

Formally, for each arbitrary area $\mathcal{A}$ in a given city on day t, we define our vector of $J = 7$ daily outcomes, $\boldsymbol{Y}_t(\mathcal{A}) = \big(Y_t^1(\mathcal{A}), \ldots, Y_t^J(\mathcal{A})\big)$, as follows:

1. **Availability**: scooters per km^2 in $\mathcal{A}$, computed hourly (7am to 11pm) and averaged over day t.
2. **Rides**: rides per km^2 starting in $\mathcal{A}$ on day t.
3. **Essential rides**: rides per km^2 starting in $\mathcal{A}$ on day t with origin/destination within 150m of health/education facilities or 20m of transit stops.
4. **Revenue**: revenue per scooter in $\mathcal{A}$ on day t, with revenue defined as \$1 + ride duration in minutes $\times$ per-minute rate.
5. **Rebalancing**: maintenance events (deployments and pickups) per scooter in $\mathcal{A}$ on day t.
6. **Battery**: average battery level of scooters at deployment in $\mathcal{A}$ on day t.

7. **Reaction time**: average time between "low battery" status updates and subsequent "maintenance pick-up" status updates for scooters in $\mathcal{A}$ on day t.

Although our primary analysis focuses on $Y_t^1(\mathcal{A})$, we present the identification and inference in full-vector notation, which nests the scalar case and accommodates the full outcome vector used in the secondary analysis.

Identification

Our identification strategy relies on the assumption that, sufficiently close to mandate boundaries, the distribution of outcomes would have been the same on either side absent the mandate (Imbens and Lemieux, 2008). This assumption is plausible given that administrative boundaries are typically ad hoc and cut through homogeneous neighborhoods, as we show on page 10. We also test for mandate effects deeper within mandate areas on page 12. We operationalize this assumption by comparing median outcome differences (over days) near the mandate boundary (Figure 2, left panel).

Formally, let $\mathcal{A}$ denote an arbitrary area in the city and $\mathcal{A}'$ its complement, so the two are disjoint and together span the city. Also, let $\mathcal{A}|_\epsilon \coloneqq \{x \in \mathcal{A} : \text{dist}(x, \mathcal{A}') < \epsilon\}$ be the coordinates in $\mathcal{A}$ within distance $\epsilon > 0$ of its boundary with $\mathcal{A}'$. Finally, let $\boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 0)$ and $\boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 1)$ denote the potential outcome vectors of $\mathcal{A}|_\epsilon$ under no treatment and treatment, respectively. The identification assumption is

$$\boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 0) \overset{d}{=} \boldsymbol{Y}_t(\mathcal{A}'|_\epsilon, 0) \text{ as } \epsilon \to 0, \tag{1}$$

where $\overset{d}{=}$ denotes equality in distribution. This assumption implies

$$\boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 1) - \boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 0) \overset{d}{=} \boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 1) - \boldsymbol{Y}_t(\mathcal{A}'|_\epsilon, 0) \text{ as } \epsilon \to 0. \tag{2}$$

For operationalization, for outcome $j = 1, \ldots, J$, we define the treatment effect

$$\tau_\epsilon^j(\mathcal{A}) \coloneqq \operatorname*{med}_t\big(Y_t^j(\mathcal{A}|_\epsilon) - Y_t^j(\mathcal{A}'|_\epsilon)\big) \tag{3}$$

and collect these effects in the vector $\boldsymbol{\tau}_\epsilon(\mathcal{A}) = \big(\tau_\epsilon^1(\mathcal{A}), \ldots, \tau_\epsilon^J(\mathcal{A})\big)$. Note that in practice we must use some $\epsilon > 0$—in particular, we use $\epsilon = 0.5$ km, except for BAL 2021 (0.1 km), DC 2020 (0.1 km), DEN 2021 and SF 2021 (0.2 km), and MSP 2023 (0.3 km) because their mandate-area interiors are very narrow—and our inference needs to account for this approximation.

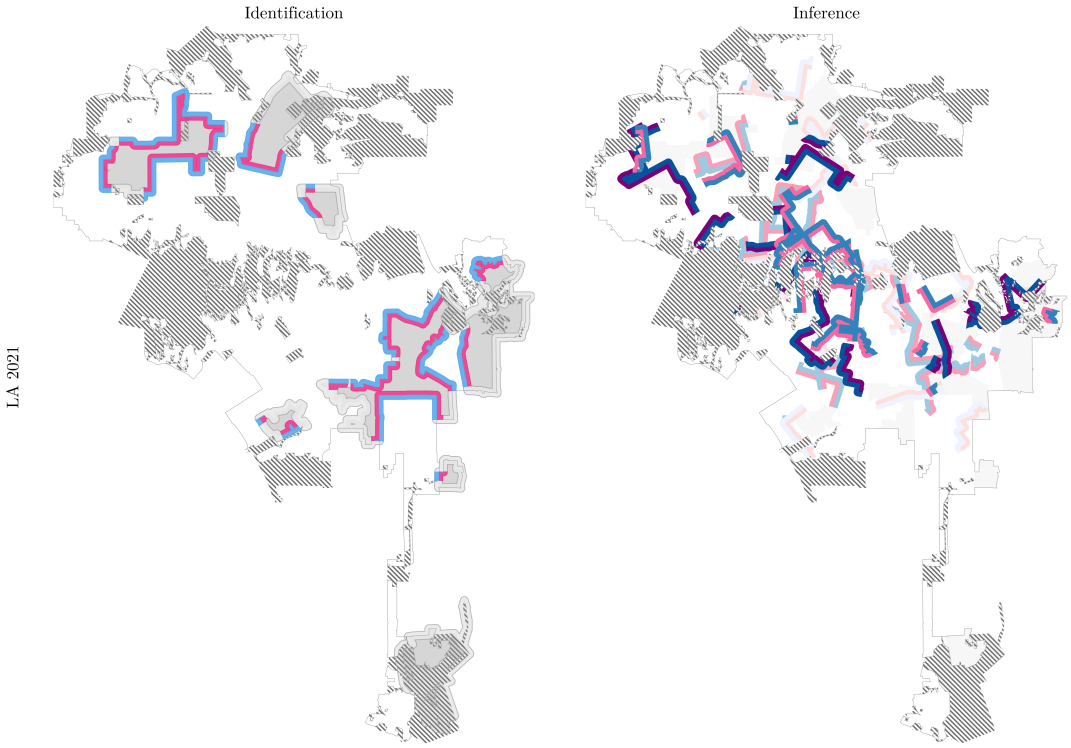


Figure 2. Identification, and inference for LA 2021. The left panel, *Identification*, illustrates the mandate area, $\mathcal{A}|_\epsilon$ (pink), and $\mathcal{A}'|_\epsilon$ (blue). The right panel, *Inference*, shows the randomly shifted boundaries used to construct placebo treatment effects. We draw spatial shifts uniformly and apply each shift separately to each boundary piece. Each pink–blue shaded combination corresponds to a different spatial shift of the mandate boundaries.

Inference

Our inference strategy relies on the null hypothesis that, sufficiently close to mandate boundaries, the outcome difference between the two adjacent bands is statistically indistinguishable from the corresponding difference across bands defined by a randomly drawn boundary (Figure 2, right panel). We operationalize this null hypothesis by randomly shifting the mandate boundaries within the city and computing the fraction of shifts whose "placebo" treatment effect is at least as large in absolute value as the effect observed at the true boundary. We also adjust for multiple-hypotheses testing (List et al., 2019; Romano and Wolf, 2005; Westfall and Young, 1993). Specifically, we first studentize (standardize by the placebo standard deviation) each outcome's test statistic across placebo draws, and then compare the maximum absolute studentized statistic across outcomes to the corresponding distribution of maxima across placebo draws. If we cannot reject, we stop. If we can, we drop the outcome corresponding to the largest absolute studentized statistic and repeat the process with the remaining outcomes.

Formally, let X be a uniformly distributed random spatial shift. The null hypothesis is

$$H_0 : \boldsymbol{Y}_t(\mathcal{A}|_\epsilon, 1) - \boldsymbol{Y}_t(\mathcal{A}'|_\epsilon, 0) \stackrel{d}{=} \boldsymbol{Y}_t((X+\mathcal{A})|_\epsilon, 0) - \boldsymbol{Y}_t((X+\mathcal{A})'|_\epsilon, 0), \text{ for sufficiently small } \epsilon. \quad (4)$$

For operationalization, we uniformly draw spatial shifts $x_1, \ldots, x_K$ and apply them to the man-

date area (Figure 2, right panel). We initialize the (adjusted) p-value for every outcome j to

$$\frac{1}{K}\sum_{k=1}^{K} \mathbf{1}(|\tau_\epsilon^j(\mathcal{A})/\sigma^j| \leq \max_{j'} |\tau_\epsilon^{j'}(x_k + \mathcal{A})/\sigma^{j'}|), \tag{5}$$

where σ^j is the standard deviation of the placebo distribution for outcome j generated by $\tau_\epsilon^j(x_k + \mathcal{A})$ across $k = 1, \ldots, K$. If the minimum p-value across outcomes is higher than 0.05, we stop. If it is lower than 0.05, we assign the outcome j^* that achieves this minimum the corresponding p-value, drop j^*, restrict the maximum in (5) to the remaining outcomes, and repeat the process. We continue until we are unable to reject or run out of outcomes.

To expand the analysis to all programs, we pool across all program-outcome combinations. For example, for the secondary analysis, which involves 9 programs and all 7 outcomes, we pool across the $9 \times 7 = 63$ studentized treatment effect estimates.

5 Design validation

Statistical power

To assess the power of our inference procedure, we check whether it can pick up availability effects of relevant magnitude—namely, closing the gap between the mandate area and the median central-neighborhood census tract (i.e., the heart of downtown). (Formally, for each program we set $\tilde{\tau} = \underset{t}{\text{med}}\big(Y_t^1(\overline{\mathcal{A}}) - Y_t^1(\mathcal{A}'_\epsilon)\big)$, where $\overline{\mathcal{A}}$ is the median-availability central-neighborhood census tract; then pool the $\tilde{\tau}$ across programs and apply the inference procedure described in Section 4.)

Across the 11 programs for which we observe every active operator—ATL 2021, BAL 2021, CHI 2020, DEN 2021, KC 2021, LA 2021, LOU 2021, PDX 2021, SAC 2021, SEA 2021, and SF 2021—these availability gaps, measured in scooters per km^2, are +32.76, +39.23, +55.00, +83.04, +234.67, +154.92, +14.02, +174.26, +80.98, +122.68, and +54.67.

Our inference procedure is well powered, detecting all gaps at the $p < 0.01$ level (Table A.2).

Covariate balance

Identification assumption (1) is not directly testable. We instead test whether several covariates that are unaffected by treatment are balanced across the boundaries. We use 6 such covariates, which we construct from SafeGraph points of interest in each program city: elementary and secondary schools, drinking places (e.g., bars and taverns), supermarkets, outpatient care centers, gas stations, and pharmacies.

We apply the same approach as for our main outcomes. That is, we operationalize identification by taking the coordinate-wise median of covariate differences at the boundary, and inference by randomly shifting the boundaries within the city, studentizing test statistics across placebo draws, and implementing the same step-down max-T procedure described above. (Formally, the

analysis is identical to that described in Section 4, except that we replace the 7-dimensional outcome vector $\boldsymbol{Y}_t$ with the 6-dimensional covariate vector $\boldsymbol{Z}$.)

We can reject the null of covariate balance (Table A.4) in only 1 out of 75 tests at the $p = 0.05$ level: BAL 2021 pharmacies ($p = 0.014$). For all remaining 74 tests, adjusted p-values exceed 0.16. For example, in LA 2021 the cross-boundary differences for the 6 covariates listed above are -0.072, 0.599, 0.054, 0.138, 0.244, and 0.240, none of which is statistically significant ($p = 1.00$ for all six).

6 Findings

Availability

Our primary analysis focuses on scooter availability. The core results use real-time location data, aggregate across all firms within each program, and restrict attention to availability as the outcome. Averaging across the 11 programs and adjusting inference for cross-program aggregation, we cannot reject the null of no aggregate effect on availability. The mean point estimate equals +4.46 available scooters per km^2 ($p = 0.52$). We can reject aggregate effects larger than +19.21 available scooters per km^2.

At the program level, only BAL 2021 and DEN 2021 generate economically and statistically significant positive effects on availability. The point estimate for BAL 2021 equals +38.38 scooters per km^2 ($p = 0.00$), and for DEN 2021 equals +70.30 scooters per km^2 ($p = 0.01$). To put these magnitudes in perspective, the initial supply gaps (described on page 10) were 39.23 scooters per km^2 in BAL 2021 and 83.04 in DEN 2021, effectively closing those gaps by 98.1% and 84.0%, respectively. Point estimates for all 11 programs are:

	Point estimate	**Adjusted p-value**
ATL 2021	−1.72	1.00
BAL 2021	+38.38	<0.01*
CHI 2020	−27.19	0.59
DEN 2021	+70.30	0.01*
KC 2021	−3.65	1.00
LA 2021	−19.34	0.99
LOU 2021	−4.01	0.75
PDX 2021	+3.11	1.00
SAC 2021	+7.73	0.74
SEA 2021	−23.48	1.00
SF 2021	+8.91	0.58

Table 2. Availability point estimates and p-values. Point estimates are measured in scooters per km^2 and asterisks denote $p < 0.05$. No-step-down Max-T inference yields identical conclusions (Table A.5).

Nested interior layers

The mandates' intended effects did not materialize deeper inside the boundaries either. Focusing on the programs with null results, we explore this possibility by partitioning each mandate area into nested interior layers, progressively moving inward from the boundary and sequentially replicating our identification and inference strategy (Figure 3).

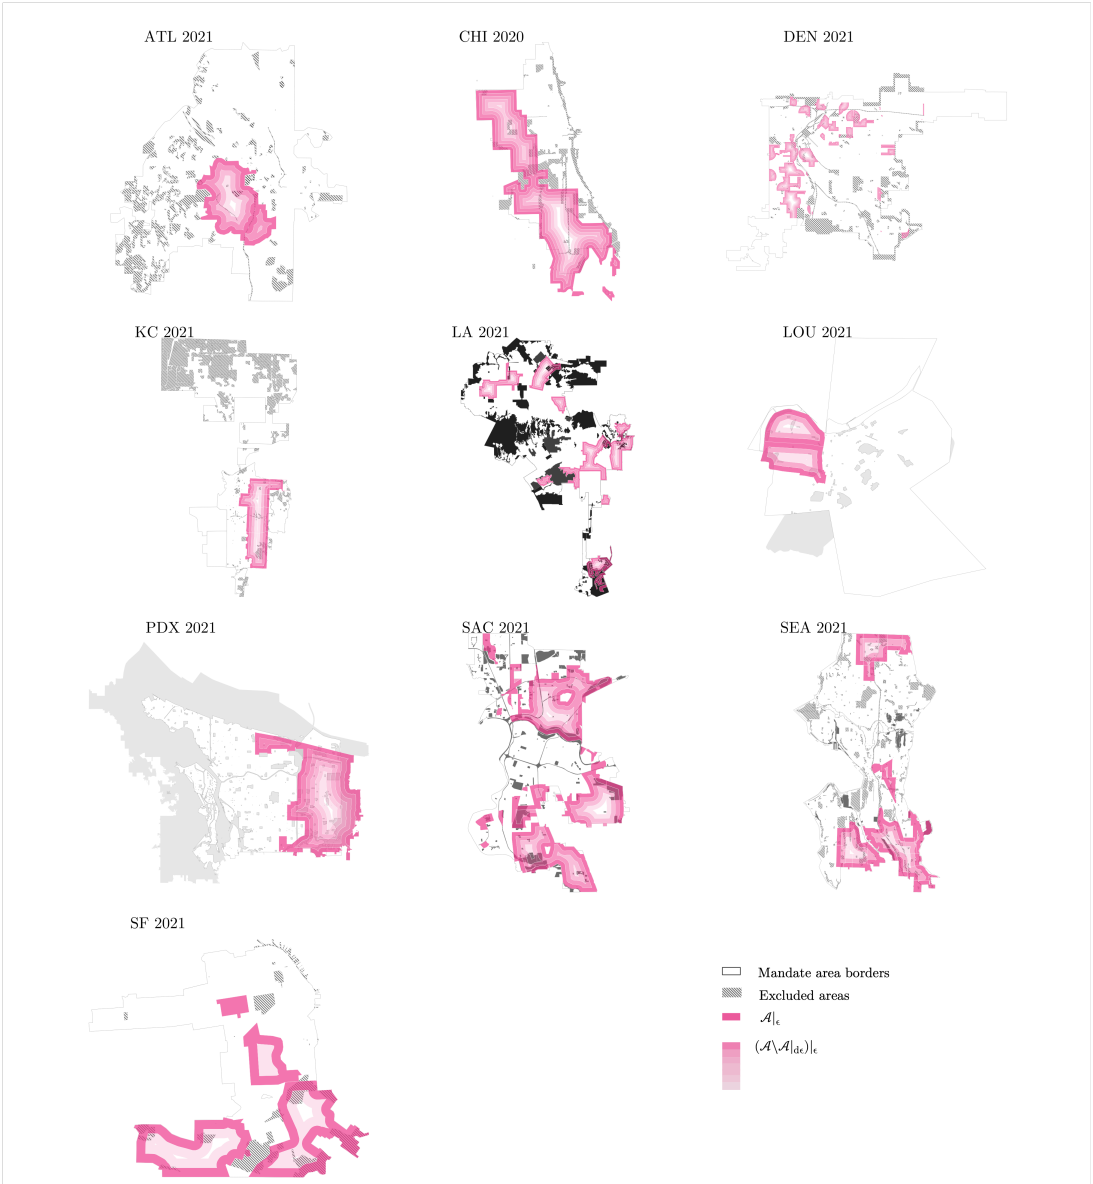


Figure 3. Nested interior layers for each program. Progressively lighter shades of pink denote layers at increasing depth d inside the mandate area: each layer d corresponds to the annulus $(\mathcal{A} \setminus \mathcal{A}|_{d\epsilon})|_\epsilon$, where $d = 1$ is the band closest to the mandate boundary and higher values of d move further inward.

Formally, for each depth $d = 1, \ldots, D$, we repeat the analysis in Section 4 replacing $\mathcal{A}$ with $\mathcal{A} \setminus (\mathcal{A}|_{d\epsilon})$. Note that for $d = 0$, $\mathcal{A} \setminus (\mathcal{A}|_{0\epsilon}) = \mathcal{A}$. For inference, we stratify programs by the number of layers they allow given their shape and size—SF 2021 allows 2; SEA 2021 allows 3; ATL 2021, KC 2021, LA 2021, and SAC 2021 allow 4; LOU 2021 allows 5; PDX 2021 allows 6; CHI 2020 allows 7—and compute p-values within each stratum. The results mirror the main

analysis: we do not find evidence of outcome improvements at the interior layers (Table A.6).

For instance, in CHI 2020, availability estimates per km^2 for depths $d = 1, \ldots, 6$ range from -2.17 to $+4.12$ with a minimum p-value of 0.71, rides per km^2 estimates range from -0.38 to $+0.17$ with a minimum p-value of 0.25, and essential-rides per km^2 estimates range from -0.09 to $+0.07$ with a minimum p-value of 0.19. We also do not observe any systematic trend (e.g., a gradual ramp up) in any outcome as we move further within the mandate areas (Figure A.2).

Additional outcomes

We next use proprietary status-update data from a large operator (with 30% fleet share in the cities we examine) to examine outcomes that require operational telemetry. This sample includes 9 programs: BAL 2021, CHI 2020, CHI 2022, DC 2020, LA 2021, MSP 2023, PDX 2019, PDX 2021, and SF 2021. We analyze 7 outcomes: availability, rides, essential rides, revenue per scooter, rebalancing per scooter, reaction time to low battery, and battery level at deployment. Note that of the two programs our main analysis detected effects for, we only observe BAL 2021—we do not observe DEN 2021.

Consistent with the main analysis, the only two statistically significant changes come from BAL 2021: an availability increase of +37.19 scooters per km^2 ($p < 0.01$) and a ridership increase of +26.47 rides per km^2 ($p = 0.05$). Among other programs, availability point estimates range between -1.31 and $+8.44$ (minimum adjusted p-value of 0.35), and ridership point estimates range between -0.89 and $+3.36$ (minimum adjusted p-value of 0.92).

For the remaining outcomes—essential rides per km^2, revenue per scooter, rebalancing per scooter, reaction time to low battery, and battery level at deployment—point estimates vary in sign across programs, and adjusted p-values are uniformly high (Table 3). For example, revenue per scooter ranges from -1.06 to $+0.68$ (minimum adjusted p-value of 1.00), rebalancing per scooter ranges from -0.09 to $+0.24$ (minimum adjusted p-value of 1.00), battery level ranges from -0.01 to $+0.31$ (minimum adjusted p-value of 1.00), and reaction time ranges from -0.03 to $+0.23$ (minimum adjusted p-value of 0.62).

	Availability		**Rides**		**Essential rides**		**Battery**		**Rebalancing**		**Revenue**		**Reaction time**	
BAL 2021	+37.19	0.001*	+26.47	0.050*	+3.71	0.976	+0.10	1.000	+0.24	1.000	+0.68	1.000	+0.01	1.000
CHI 2020	+0.61	1.000	−0.01	1.000	−0.08	1.000	+0.08	0.942	−0.05	1.000	−0.11	1.000	+0.09	0.620
CHI 2022	−0.08	1.000	−0.18	1.000	0.00	1.000	+0.03	1.000	−0.09	1.000	−0.95	1.000	−0.03	1.000
LA 2021	+1.13	0.920	+0.39	0.998	+0.17	1.000	+0.04	1.000	0.00	1.000	−0.32	1.000	0.00	1.000
MSP 2023	+8.44	0.352	+3.36	0.920	+0.84	1.000	+0.04	1.000	0.00	1.000	+0.01	1.000	+0.23	1.000
PDX 2019	−0.02	1.000	+0.12	1.000	−0.28	1.000	+0.31	1.000	0.00	1.000	0.00	1.000	+0.03	1.000
PDX 2021	+0.26	1.000	+0.26	1.000	−0.28	1.000	+0.31	0.996	0.00	1.000	0.00	1.000	+0.03	1.000
SF 2021	+1.88	1.000	+0.43	1.000	−1.06	1.000	−0.01	1.000	+0.01	1.000	−0.71	1.000	0.00	1.000
DC 2020	−1.31	1.000	−0.89	1.000	−1.82	1.000	−0.01	1.000	+0.05	1.000	−1.06	1.000	+0.01	1.000

Table 3. Point estimates and p-values for additional outcomes. See page 7 for these outcomes' definitions.

7 Mechanism discussion

The ineffectiveness of equitable deployment mandates was not driven by noncompliance. First, deployment generally met or exceeded the mandated thresholds, which therefore did not bind (Table A.7). This conclusion holds even when measuring compliance with median daily availability across operational hours (7 a.m.–11 p.m.) over each program period—a more stringent metric than one-time compliance snapshots, which most cities use. For example, DEN 2021 mandated a 30% fleet share and realized 34.76%; LA 2021 mandated 20% and realized 21.03%; and CHI 2020 mandated 2.5% in each of twenty subareas, all of which exceeded the threshold at the median, with availability ranging from 2.69% to 2.89%. PDX 2021 fell slightly below its 20% threshold, realizing 17.08%. Only KC 2021 and LOU 2021 showed substantial shortfalls, realizing 4.63% and 2.93% against thresholds of 10% and 20%, respectively. Second, compliance rates are uncorrelated with treatment effects (Figure 4, left panel).

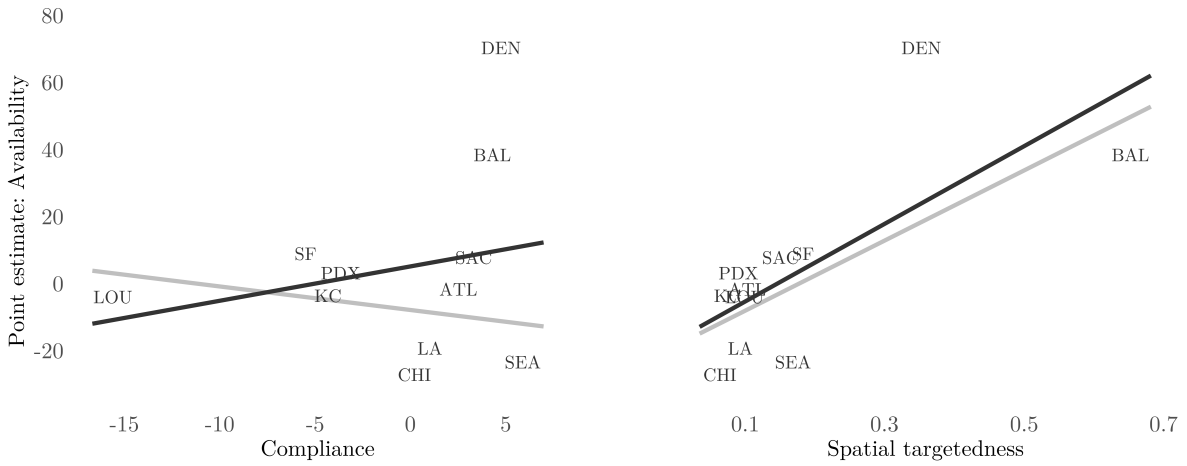


Figure 4. Availability estimates against compliance and targetedness. Each panel displays two OLS regression lines: a dark line fitted to all 11 cities, and a light line that excludes BAL 2021 and DEN 2021.

A different dimension separates the positive cases from the rest: spatial targetedness. To see this point, contrast LA 2021 (Figure 1, second row, third column) and BAL 2021 (Figure 1, first row, second column). LA 2021 required operators to deploy 20% of their fleet across 7 subareas, roughly located at Central LA/Boyle Heights, Del Rey, North Hollywood, San Pedro, South Central LA, Sylmar, and West Valley/Canoga Park. The operator fully complied, but did so by deploying a daily median of 28% of its fleet in Central LA—bordering Downtown LA and Hollywood, its two highest-traffic neighborhoods—and a median of 0% in the other 6 subareas. BAL 2021, on the other hand, required 3 devices in each of 12 small zones scattered along transit stops, such as bus routes, covering just 1.8% of the city. Operators complied by doing exactly that—no room for misalignment.

Indeed, BAL 2021 and DEN 2021 are clear outliers in targetedness. To quantify this pattern, we overlay a 100-meter grid on each city's service area and compute the fraction of cells that

touch the mandate area that also touch its complement. BAL 2021 has a ratio of 0.65, DEN 2021 has a ratio of 0.35, whereas all other programs' ratios range between 0.06 and 0.18. Moreover, targetedness ratios are positively correlated with treatment effects (Figure 4, right panel).

That targetedness drives success is consistent with Kabra et al. (2020). Studying Paris's docked bike-share system, they show that the ridership return to improving availability depends sharply on where the improvement occurs. In a counterfactual exercise, increasing bike availability by 10 percentage points at the four nearest stations to each point of interest generates the largest ridership gains near supermarkets and public transit stops (about +19% and +9% per station improved), well above the gains near cafes (+5.5%), and universities (+3.5%).

8 Conclusion

Taking a step back, our study shows that equitable deployment mandates are far from guaranteed to improve access in underserved communities, but they can succeed when spatially targeted. Of course, further investigation is needed to pin down all relevant design features. But, even leaving mechanisms aside, evaluating the effectiveness (or lack thereof) of equity mandates is crucial in itself: ineffective policies left unchecked waste scarce regulatory resources and mask ongoing service gaps.

Endnotes

1. See Appendix B for the comparison of grocery store trips in historically disadvantaged neighborhoods versus central neighborhoods via public transportation, and see the U.S. Department of Transportation's "Justice40 Fact Sheet" (link) for Bipartisan Infrastructure Law spending details.

2. See North American Bikeshare and Scootershare Association's "2024 Shared Micromobility State of the Industry Report" (link) for the 2024 trip data. See Brown and Howell (2024) and National Institute for Transportation and Communities' "Evaluating Equity Requirements in Shared Mobility Programs" (link) for an overview of equity mandates and regulatory programs. See Federal Transit Administration Circular 4702.1B (link), Denver RTD's "Service Policies & Standards" (link), and Maricopa Association of Governments' Transit Standards (Phase I) (link) for the federal guidelines and regional service standards.

3. See Transportation Research and Education Center (TREC) and Portland State University's report, "Taxing Shared Micromobility: Assessing the Global Landscape of Fees and Taxes and Their Implications for Cities, Riders, and Operators" (link) for information on market shares and operator fleet sizes during the study period. See the Mobility Data Specification (MDS), an open-source standard developed by the Open Mobility Foundation (link), for the standardized data-sharing specifications used by shared micromobility operators.

References

Brown A, Howell A (2024) Mobility for the people: Equity requirements in us shared micromobility programs. *Journal of Cycling and Micromobility Research* 2:100020. 4, 16

Calvo E, Cui R, Serpa JC (2019) Oversight and efficiency in public projects: A regression discontinuity analysis. *Management Science* 65(12):5651–5675. 3

Cattaneo MD, Titiunik R (2022) Regression discontinuity designs. *Annual Review of Economics* 14:821–851. 3

Erkan H, Agarwal A, Muthuraman K, Stamatopoulos I (2025) Competition for customer attention is creating oversupply pockets in US micromobility. *Operations Research* Forthcoming. 3

He P, Zheng F, Belavina E, Girotra K (2021) Customer preference and station network in the London bike-share system. *Management Science* 67(3):1392–1412. 3

Imbens GW, Lemieux T (2008) Regression discontinuity designs: A guide to practice. *Journal of Econometrics* 142(2):615–635. 3, 8

Kaaua D, Virudachalam V (2025) Going the distance: The impact of commute on gender diversity in public service. *Manufacturing & Service Operations Management* 27(2):607–623. 3

Kabra A, Belavina E, Girotra K (2020) Bike-share systems: Accessibility and availability. *Management Science* 66(9):3803–3824. 15

Kadiyala B, Özer Ö, Şimşek AS (2024) Dual value of delayed incentives: An empirical investigation of gift card promotions. *Manufacturing & Service Operations Management* 26(5):1806–1825. 3

Lee DS, Lemieux T (2010) Regression discontinuity designs in economics. *Journal of Economic Literature* 48(2):281–355. 3

List JA, Shaikh AM, Xu Y (2019) Multiple hypothesis testing in experimental economics. *Experimental Economics* 22:773–793. 9

Lodi A, Malaguti E, Stier-Moses NE, Bonino T (2016) Design and control of public-service contracts and an application to public transportation systems. *Management Science* 62(4):1165–1187. 6

Romano JP, Wolf M (2005) Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association* 100(469):94–108. 9

Rosaia N (2025) Competing platforms and transport equilibrium. *Econometrica* 93(6):2235–2271. 3

Siddiq A, Tang CS, Zhang J (2022) Partnerships in urban mobility: Incentive mechanisms for improving public transit adoption. *Manufacturing & Service Operations Management* 24(2):956–971. 3

Westfall PH, Young SS (1993) *Resampling-Based Multiple Testing: Examples and Methods for* p*-Value Adjustment* (New York: John Wiley & Sons), ISBN 978-0-471-55761-7. 9

Wu OQ, Zhang J (2025) Public transit time and fare design: Ridership maximization under income disparity Available at SSRN 5989535. 3

Xu A, Yan C, Goh CY, Jaillet P (2025) A locational demand model for bike-sharing. *Manufacturing & Service Operations Management* 27(3):897–916. 3

A Public transit stops in historically disadvantaged and central neighborhoods

We compile stop locations for Chicago, New York City, Los Angeles, Phoenix, and Houston from municipal GTFS feeds and GIS layers.[1] Each dataset was converted to WGS84 coordinates and formatted to a common schema (`id`, `lon`, `lat`), then spatially joined to census tracts. Central neighborhoods were identified from municipal open-data portals, and historically disadvantaged tracts were defined using the Justice40 Tracts November 2022 archive. For each tract, we calculate the number of stops, tract area (km^2), and stop density per km^2, and then aggregated results by central neighborhood versus disadvantaged tracts.

B Grocery–access analysis

For the large-population cities mentioned on page 2, we use the Justice40 disadvantaged communities map.[2] We use the corresponding mandate areas and downtown polygons from municipal open-data portals. Additional sources include U.S. Census block groups with total population, the U.S. Department of Agriculture's *SNAP Retailer Locator* (filtered to *Supermarket*), and the Google Directions and Routes Distance Matrix APIs for transit accessibility. The Supplemental Nutrition Assistance Program (SNAP) is the federal food-assistance program that subsidizes grocery purchases for low-income households; its retailer database provides the most comprehensive listing of authorized food stores nationwide.

1. Intersect block groups with the city boundary and apportion population by area share. Sample $N = 500$ origins per city, with probability proportional to the apportioned population.

[1] Chicago Transit Authority. "CTA System Information — List of 'L' Stops (CSV)" (link) (2025). Dataset: System Information — List of 'L' Stops; accessed 2025-09-28; Chicago Transit Authority and City of Chicago Data Portal. "CTA Bus Stops (Point Layer)" (link) (2025). Dataset: CTA_BusStops; accessed 2025-09-28; Metropolitan Transportation Authority. "MTA Static GTFS Feeds (Stops.txt Across Subway and Bus Feeds)" (link) (2025). GTFS stops for multiple feeds; accessed 2025-09-28; NYC Open Data and Metropolitan Transportation Authority. "Subway Stations (NYC Open Data)" (link) (2025). Dataset: MTA Subway Stations; accessed 2025-09-28; Los Angeles County Metropolitan Transportation Authority. "Metro Rail Lines and Stops (Point Layer)" (link) (2025). Also available via Los Angeles GeoHub; accessed 2025-09-28; City of Los Angeles GeoHub. "Bus Stop Benches (Point Layer)" (link) (2025). Bus stop points/amenities proxy; accessed 2025-09-28; Valley Metro. "Light Rail Stations (Point Layer)" (link) (2025). GIS/Hub resource; accessed 2025-09-28; City of Phoenix. "Bus Stops with Amenities (Point Layer)" (link) (2025). GIS layer; accessed 2025-09-28; Metropolitan Transit Authority of Harris County (Houston METRO). "METRORail Stations (Point Layer)" (link) (2025). GIS/Open data portal; accessed 2025-09-28; Metropolitan Transit Authority of Harris County (Houston METRO). "METRO Bus Stops (Point Layer)" (link) (2025). GIS/Open data portal; accessed 2025-09-28

[2] ArcGIS Online. "Justice40 Tracts November 2022 Version 1.0 (Archive)" (link) (2022). Accessed: 2025-09-28; The White House. "Justice40 Initiative" (link) (2023). Accessed: 2025-09-28

2. Keep SNAP-authorized stores active in 2024 with type = *Supermarket*; drop entries with missing coordinates or duplicates, and clip to the city boundary.

3. Pair each origin to its nearest supermarket using a K-nearest-neighbors search ($K = 1$).

4. For each origin–destination pair, call Google Directions (transit mode) to obtain an itinerary and arrival time, and Google Routes Distance Matrix (transit) for travel duration and distance, using a fixed local departure time.

5. Fallbacks: if Directions returns *OK* but no transit leg, re-query Routes in walking mode; if that fails, estimate walking time from straight-line distance at 4.8 km/h.

C Equitable-access regulation history

Atlanta, GA: The City of Atlanta adopted Ordinance 18-O-1322 in 2018 to regulate shareable dockless mobility devices before updating the permit framework through Ordinance 20-O-1174 in March 2020. The initial phased deployment began July 1, 2020, with operators Bird, Helbiz, Spin, and VeoRide; the 2021–2022 annual permit cycle (June–December 2021) included Bird, Spin, and Lime. The Atlanta Department of Transportation (ATLDOT) designated three Equity Zones based on metrics including households without vehicle access and other demographic indicators, and required each operator to deploy at least 2% of its daily fleet in each zone. Compliance was audited monthly, with permit revocation as the enforcement mechanism. Permit holders were required to submit an Equity Plan as part of their application and to offer payment and access options that do not require a credit card or smartphone.[3]

Baltimore, MD: The Baltimore City Department of Transportation (BCDOT) launched a dockless vehicle pilot in 2018 before the Baltimore City Council authorized a formal permit program in 2019. The regulations initially designated 20 Equity Zones in historically underserved neighborhoods, selected based on proximity to transit and bike lanes, percentage of households without a motor vehicle or in poverty, population density, and distance to amenities. After one year, eight zones were eliminated based on utilization data, leaving 12 active Equity Zones for the June–August 2021 permit period—each typically encompassing two to three city blocks. Every morning between 5:00 and 9:00 a.m., each permitted operator was required to place at least three vehicles in each Equity Zone, with permit revocation as the enforcement mechanism. All permit holders were required to offer a low-income access plan for SNAP/WIC-enrolled individuals, cash payment options, and text-to-unlock plans for riders without smartphones.[4]

[3] City of Atlanta. "Ordinance 18-O-1322: Amendment to Chapter 150 (Traffic and Vehicles) Creating Article X — Shareable Dockless Mobility Devices" (link) (2018). Accessed: 2025-10-15; City of Atlanta. "Ordinance 20-O-1174: Updated Permit Structure for Shareable Dockless Mobility Devices" (link) (2020). Accessed: 2025-10-15

[4] Baltimore City Department of Transportation. "Dockless Vehicles for Hire Regulations (Title 14, Subtitle 02)" (link) (2021). Accessed: 2025-10-15; Baltimore City Department of Transportation. "Dockless Vehicles — Shared

Chicago, IL: The Chicago Department of Transportation (CDOT) piloted shared e-scooters in 2019 (June–October) and 2020 (August–December) before transitioning to a permanent licensing program in 2022. The 2020 pilot introduced an Equity Priority Area divided into 20 zones, mandating that at least 50% of each operator's fleet be located in the Equity Priority Area with a minimum of 2.5% in every zone. The 2022 permanent program retained the 50% Equity Priority Area requirement but redefined it into 10 sub-areas, each requiring a 4% minimum. All programs required operators to offer low-income pricing, non-smartphone access programs, and adaptive scooters on request.[5]

Denver, CO: The Denver Department of Transportation and Infrastructure (DOTI) ran a dockless mobility vehicle pilot from August 2018 to May 2021 before transitioning to licensing agreements with Lyft and Lime in May 2021. DOTI developed an Equity Index combining demographic, socioeconomic, and transportation-access characteristics—including race and ethnicity, income, vehicle ownership, and transit ridership—to identify 25 Opportunity Areas across the city. Under the permanent program, operators are required to deploy a minimum of 30% of their fleet daily across these Opportunity Areas. Compliance is audited daily through the Ride Report platform, and enforcement includes fines, fleet-size reductions, and permit revocation. Operators must offer reduced pricing for riders qualifying for local, state, or federal assistance programs, as well as a cash payment option.[6]

Washington, DC: The District Department of Transportation (DDOT) regulated shared scooters through pilot phases (2018–2019) before transitioning to annual permits in 2020. DDOT designated 18 Equity Emphasis Area subareas across the District. The 2020 permit mandated a minimum of 400 devices jointly across these areas between 05:00–07:00 and required at least 20 vehicles in each ward. Compliance was audited daily at 6:00 a.m., with permit revocation as the enforcement mechanism. Subsequent permits (2021–2022) adjusted fleet sizes and ward-level requirements while maintaining equity deployment mandates. All permits required operators to implement a Low-Income Customer Plan offering free unlimited 30-minute trips, provide non-smartphone access options, conduct outreach in underserved neighborhoods, and supply adaptive scooters on request.[7]

Scooters and Bicycles Frequently Asked Questions" (link) (2021). Accessed: 2025-10-15

[5]Chicago Department of Transportation. "E-Scooter Pilot Evaluation 2019" (link) (2020). Accessed: 2025-10-06; Chicago Department of Transportation. "2020 E-Scooter Pilot Program Terms" (link) (2020). Accessed: 2025-10-06; Chicago Department of Transportation. "Chicago Scooter License Rules and Regulations" (link) (2022). Accessed: 2025-10-06

[6]Denver Department of Transportation and Infrastructure. "Dockless Mobility Vehicle Permit Pilot Program Final Report" (link) (2021). Accessed: 2025-10-15

[7]District Department of Transportation. "Dockless Demonstration Phase I Evaluation" (link) (2019). Accessed: 2025-10-06; District Department of Transportation. "Dockless Terms and Conditions Phase II – Scooters" (link) (2019). Accessed: 2025-10-06; District Department of Transportation. "Shared Dockless 2020 Terms and Conditions – Scooter" (link) (2020). Accessed: 2025-10-06; District Department of Transportation. "Shared Dockless

Kansas City, MO: The City of Kansas City launched a yearlong scooter and e-bike pilot program in May 2019 with Bird, Spin, and RideKC before transitioning to annual operating agreements. Each company signed an operating agreement that included equity and inclusion provisions. The program designated six equity subareas in underserved neighborhoods and required operators to jointly deploy at least 10% of their fleet across these areas, with deployment restrictions as the enforcement mechanism. No formal audit cadence was specified in the public operating agreements. Operators were required to offer low-income pricing plans and non-smartphone access options.[8]

Los Angeles, CA: The Los Angeles Department of Transportation (LADOT) piloted shared scooters in 2019–2020 before adopting permanent On-Demand Mobility permits. Operators deploying in Special Operations Zones were required to stage at least 20% of their total fleet in Equity-Focus and Mobility Development Districts, with a minimum of 5% in Equity-Focus Districts. The permit program paired these deployment requirements with a variable per-trip fee structure that charged $0.00 in Equity-Focus Districts. Operators were also required to offer low-income pricing programs, non-smartphone access, and multilingual customer support.[9]

Louisville, KY: Louisville Metro Public Works established a dockless vehicle policy in October 2018 following the arrival of shared e-scooters and updated the policy in December 2020. The city permits up to four operators, each beginning with a six-month probationary period before receiving annual permits. The operational zone is bounded by the Watterson Expressway (I-264), with extensions south to Iroquois Park and west to Shawnee Park. Within this area, the city established nine distribution zones for geographic coverage, of which two are designated equity subareas in historically underserved neighborhoods. Operators were required to jointly deploy at least 20% of their fleet across these two equity subareas (June–August 2021), with fines and device removal as enforcement mechanisms. No formal audit cadence was specified. Operators were required to submit monthly data reports documenting compliance with deployment requirements.[10]

Minneapolis, MN: The City of Minneapolis piloted shared e-scooters in 2018 (August–November) and 2019 (May–November) before transitioning to annual license agreements. The 2019 pilot required at least 30% of each operator's fleet in Areas of Concentrated Poverty. In March 2023, the

2021 Terms and Conditions – Bike" (link) (2021). Accessed: 2025-10-06; District Department of Transportation. "Dockless Scooter Terms and Conditions" (link) (2022). Accessed: 2025-10-06

[8]City of Kansas City, Missouri. "Scooter and eBike Pilot Program Launched" (link) (2019). Accessed: 2025-10-15

[9]Los Angeles Department of Transportation. "Dockless On-Demand Personal Mobility: Year-One Report" (link) (2019). Accessed: 2025-10-06; Los Angeles Department of Transportation. "Dockless Mobility Annual Permit Rules" (link). Accessed: 2025-10-06

[10]Louisville Metro Public Works. "City Announces Policy for Dockless Vehicle Transportation" (link) (2018). Accessed: 2025-10-15; Louisville Metro Public Works. "Dockless Scooter Policy (December 2020 Update)" (link) (2020). Accessed: 2025-10-15

City codified Equity Distribution Areas with zone-level minimums: Zone 1 (2%), Zone 2 (15%), and Zone 3 (21%). For compliance monitoring, these zones were subdivided into six subzones (A–F) with minimum shares of 2%, 7%, 8%, 7%, 7%, and 7%, respectively. License agreements also required adaptive scooters on request, low-income pricing, multilingual education campaigns, and non-smartphone access.[11]

Portland, OR: The Portland Bureau of Transportation (PBOT) piloted shared e-scooters from July to November 2018 before transitioning to a permit program in 2019. The 2018 pilot required permittees to deploy at least 100 scooters daily (or 20% of their fleet, whichever was less) in the East Neighborhoods. The 2019 rules mandated that at least 15% of the fleet be deployed daily in the East Neighborhoods; the 2021 permit restored the threshold to at least 20% of the fleet deployed daily in the East Neighborhoods. Compliance was audited daily, with fines, suspension, and revocation as enforcement mechanisms. Permits also required low-income pricing, user education, multilingual outreach, and non-smartphone access options.[12]

Sacramento, CA: The Sacramento City Council adopted an updated shared rideable ordinance on April 2, 2019, authorizing citywide operating permits for shared e-scooters and bike-share. The ordinance designated nine disadvantaged-community subareas and required operators to jointly deploy at least 20% of their fleet across these areas. Compliance was audited daily, with fleet-size reductions and permit suspension as enforcement mechanisms. Operators were required to offer reduced-cost plans for low-income riders and non-smartphone access options.[13]

San Francisco, CA: The San Francisco Municipal Transportation Agency (SFMTA) launched its powered scooter share pilot in November 2019 before transitioning to annual permits. The 2020 distribution guidelines required operators to maintain 50% fleet coverage for 75% of the day (06:00–22:00) across designated areas and set minimum thresholds in Key Neighborhoods. The 2021–2022 permit retained the 75/75 coverage requirement while shifting focus to systemwide equity compliance. Both pilots and permits required operators to offer low-income pricing, non-smartphone access options, adaptive devices on request, and multilingual outreach programs.[14]

Seattle, WA: The Seattle Department of Transportation (SDOT) launched a free-floating

[11] City of Minneapolis. "Shared Scooter Presentation" (link) (2018). Transportation & Public Works Committee presentation, Accessed: 2025-10-06; City of Minneapolis. "Minneapolis Scooter Pilot Summary 2018–2019 Update" (link) (2019). Accessed: 2025-10-06; City of Minneapolis. "Scooter Data" (link). Accessed: 2025-10-06

[12] Portland Bureau of Transportation. "Scooter Share Permit Application" (link) (2018). Accessed: 2025-10-06; Portland Bureau of Transportation. "Administrative Rule TRN-15.01: New Mobility – Shared Electric Scooters" (link) (2021). Accessed: 2025-10-06

[13] City of Sacramento. "City Council Adopts New Ordinance on Shared Electric Scooters and Bikes" (link) (2019). Accessed: 2025-10-15

[14] San Francisco Municipal Transportation Agency. "Powered Scooter Share Policy Directive" (link) (2019). Accessed: 2025-10-06; San Francisco Municipal Transportation Agency. "Powered Scooter Share Distribution Guidelines" (link) (2020). Accessed: 2025-10-06; San Francisco Municipal Transportation Agency. "2021–2022 Scooter Permit Terms and Conditions and Appendices" (link) (2022). Accessed: 2025-10-06

scooter share pilot in fall 2020, initially permitting Lime, LINK, and Wheels; the study period covers August–December 2021. SDOT designated four equity subareas—encompassing neighborhoods including Lake City, Bitter Lake, and Haller Lake in North Seattle; Seward Park, South Beacon Hill, Rainier Valley, and Rainier Beach in South Seattle; Delridge and High Point in West Seattle; and the Central District, Georgetown, and South Park—selected based on higher proportions of communities of color, immigrants, refugees, people with low incomes, and limited-English-proficient individuals. Operators were required to deploy at least 10% of their fleet across these areas daily, with fleet-size reductions and permit revocation as enforcement mechanisms. No formal audit cadence was specified. Subsequent permit versions (Permit Requirements 5.0) increased this threshold to 15%. All operators were required to provide reduced-rate plans for low-income residents (capped at $1.50 per hour for riders enrolled in government income-qualified assistance programs), non-smartphone and non-credit-card access options, and seated scooters for riders with mobility impairments.[15]

[15]Seattle Department of Transportation. "SDOT Invites Three Scooter Share Companies to Seek Permits After a Thorough Selection Process" (link) (2020). Accessed: 2025-10-15; Seattle Department of Transportation. "Free-Floating Shared Micromobility Program: Permit Requirements 5.0" (link). Includes equity deployment minimums (15% in Equity Focus Neighborhoods), compliance and data-reporting requirements

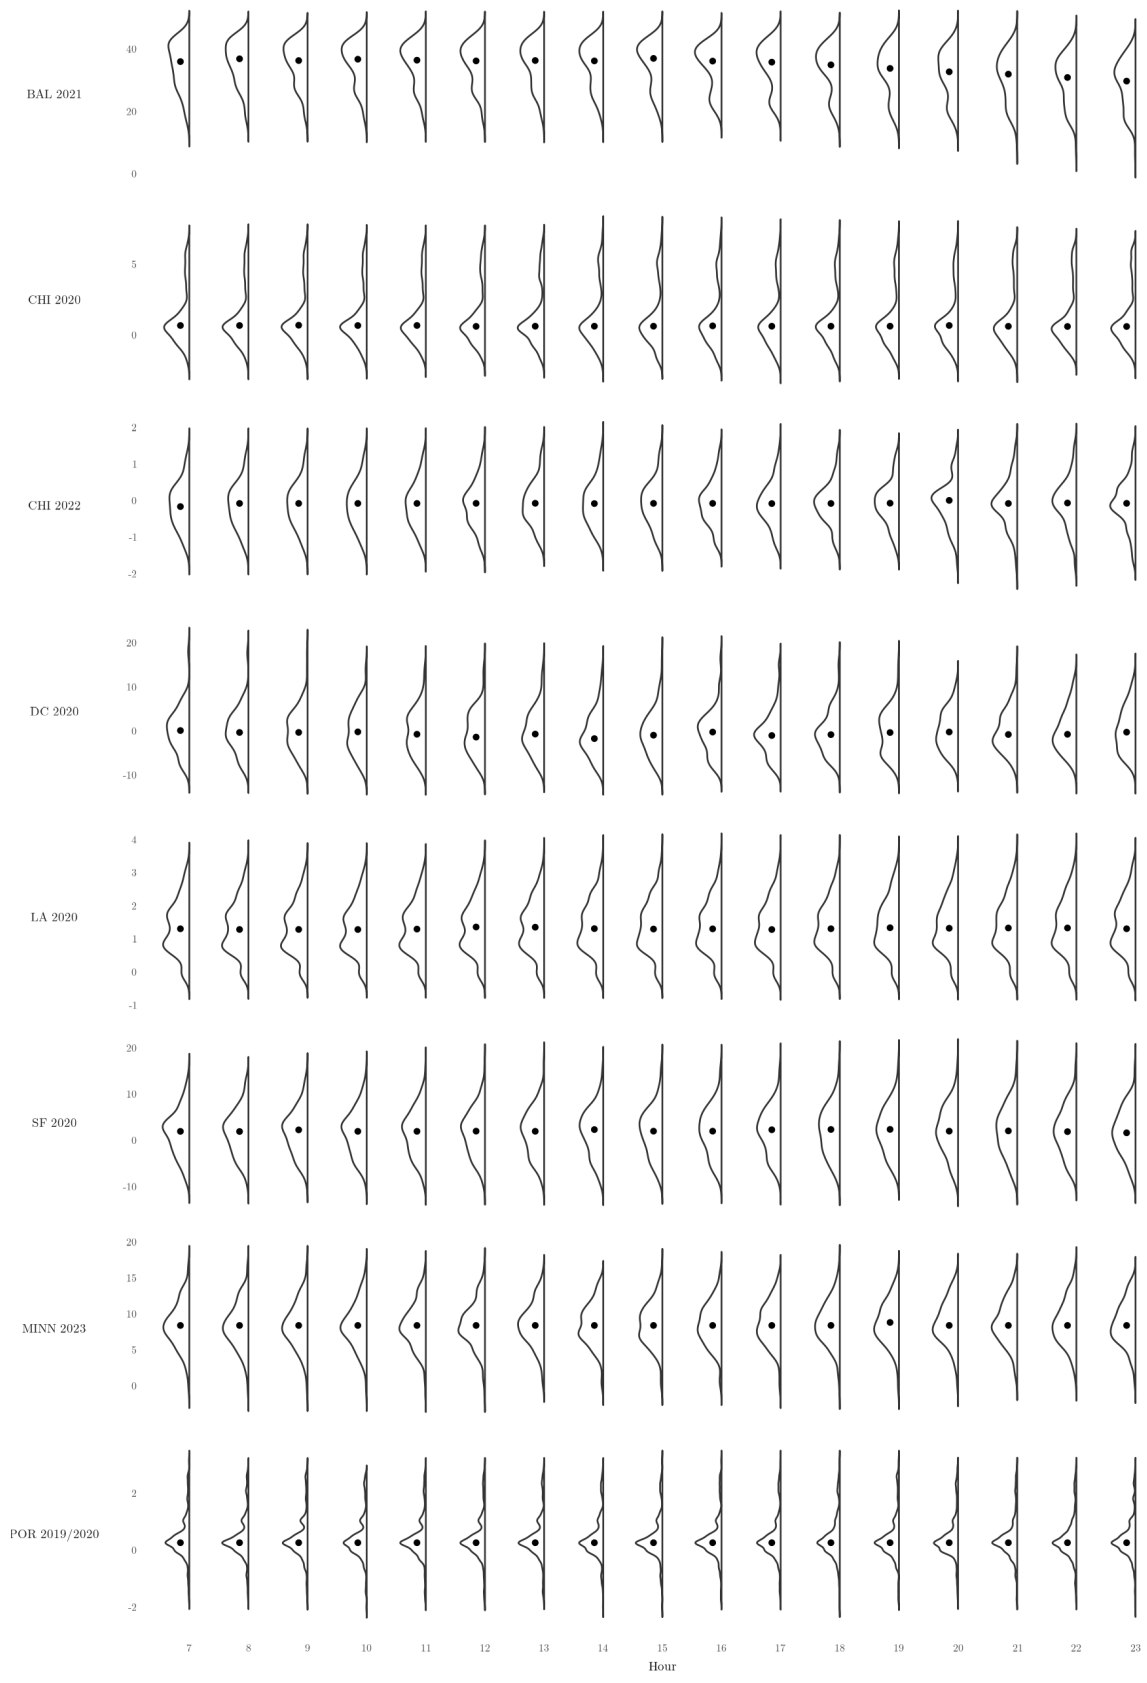


Figure A.1: Hourly differences in the number of available scooters between $\mathcal{A}|_\epsilon$ and $\mathcal{A}'|_\epsilon$ for each program, calculated for every hour from 7:00 a.m. to 11:00 p.m.

D Essential service locations

We identified essential service locations using the ArcGIS Business Analyst "Points of Interest" tool, which provides geocoded facility data (latitude and longitude) for each city–pilot. Searches were conducted separately by service category, classified as follows:

- **Health:** General medical and surgical hospitals (including emergency facilities), surgical centers, medical offices, and pharmacies.
- **Education:** Colleges and universities, including constituent departments such as dental, engineering, and law schools.
- **Transportation:** Local and suburban transit stops and stations.

Program	Total number of rides	Total number of essential rides
BAL 2021	294,127	81,422
CHI 2020	133,891	61,335
CHI 2022	66,648	29,638
DC 2020	202,944	76,379
MSP 2023	79,334	27,018
LA 2021	156,856	37,997
PDX 2019	147,663	89,059
PDX 2021	80,274	34,451
SF 2021	119,051	69,951

These facilities were geocoded and stored as point layers for each city–program. Scooter rides were classified as *essential* if either the origin or the destination lay within 150 m of a health or education facility, or within 20 m of a transit stop. Distances were computed as great-circle distances (Haversine formula) based on latitude–longitude coordinates in WGS84.

E Randomly shifting mandate boundaries

To construct placebo boundaries, we randomly re-positioned each pair of adjacent polygons—the areas immediately inside and outside the mandate boundary, $\mathcal{A}|_\epsilon$ and $\mathcal{A}'|_\epsilon$—independently within each city. For each pair, we calculated the centroid of their union and applied a random rotation and translation, moving both polygons together to preserve their geometry and adjacency while randomizing their location within the city. Transformations were retained only if the shifted polygons remained entirely within the municipal boundary. We repeated this procedure independently for each pair, generating 500 placebo boundary sets per city for inference.

F Placebo distribution validation

We verify that placebo distributions are centered at zero by testing whether the median of each placebo distribution equals zero using the Wilcoxon signed-rank test. Under the null hypothesis

of no systematic bias, placebo statistics should be symmetrically distributed around zero.

Program	Median	95% CI	p-value
ATL 2021	0.19	[−0.13, 0.52]	0.244
BAL 2021	0.29	[−0.08, 0.66]	0.122
CHI 2020	0.14	[−0.72, 1.02]	0.741
DEN 2021	0.50	[−0.12, 1.13]	0.116
KC 2021	−0.15	[−1.06, 0.52]	0.659
LA 2021	−0.30	[−2.09, 1.42]	0.723
LOU 2021	−0.07	[−0.22, 0.02]	0.152
PDX 2021	−0.41	[−1.61, 0.76]	0.485
SAC 2021	0.06	[−0.37, 0.45]	0.800
SEA 2021	−3.97	[−8.41, 0.10]	0.056
SF 2021	0.61	[−1.23, 2.44]	0.505

Table A.1: Wilcoxon signed-rank test for placebo median = 0 (availability analysis, 11 cities). All 95% confidence intervals include zero.

Table A.1 shows results for the 11-city availability analysis. All programs have $p > 0.05$, indicating we cannot reject the null hypothesis that placebo medians equal zero. Seattle shows a marginally significant deviation ($p = 0.056$).

G Statistical power analysis

We evaluate statistical power by comparing the outer-ring baseline ($\mathcal{A}'|_\epsilon$) to benchmarks in central neighborhoods centered on the median tract. For each city and outcome, tracts within the central neighborhood are ranked by the median value of the metric, and the benchmark is computed as the average across three tracts: the tract at the exact median and its immediate neighbors in the ranking (one above, one below). Daily differences are winsorized at the 1st and 99th percentiles to reduce sensitivity to outliers.

Central neighborhood boundaries follow the same definitions used throughout the analysis. In all cities, we use the administrative downtown area as the central neighborhood, except in **Chicago**, where scooter operations are prohibited within the central business district. In that case, we use nearby neighborhoods with high scooter activity—Lake View, Lincoln Park.

Program	Point estimate	p-value	Adj. p-value
ATL 2021	+32.76	0.005	0.008
BAL 2021	+39.23	0.002	0.003
CHI 2020	+55.00	0.002	0.006
DEN 2021	+83.04	0.004	0.008
KC 2021	+234.67	0.003	0.003
LA 2021	+154.92	0.002	0.003
LOU 2021	+14.02	0.002	0.006
PDX 2021	+174.26	0.002	0.003
SAC 2021	+80.98	0.005	0.006
SEA 2021	+122.68	0.002	0.003
SF 2021	+54.67	0.008	0.008

Table A.2: Power analysis point estimates with unadjusted and adjusted p-values. The p-value column reports standard (unadjusted) randomization inference p-values for each program individually. The Adj. p-value column reports p-values adjusted for multiple-hypothesis testing using the stepdown procedure described in Section 3.

H SafeGraph covariate categories

For the covariate balance analysis described in Section 4 of the main text, we use six categories of points of interest (POIs) from the SafeGraph Places dataset. Each category corresponds to a six-digit North American Industry Classification System (NAICS) code. For each program city, we count the number of POIs in each category that fall within the inner and outer bands adjacent to the mandate boundary and compute the cross-boundary difference. Table A.3 lists the six categories, their NAICS codes, and descriptions.

NAICS	Category	Description
611110	Elementary and Secondary Schools	Public and private K–12 schools, including elementary, middle, and high schools.
722410	Drinking Places (Alcoholic Beverages)	Bars, taverns, nightclubs, and similar establishments primarily serving alcoholic beverages.
445110	Supermarkets and Other Grocery Retailers (except Convenience Retailers)	Full-line grocery stores and supermarkets; excludes convenience stores.
621498	All Other Outpatient Care Centers	Outpatient care centers not classified elsewhere, including sleep disorder clinics, dialysis centers, and other specialty outpatient facilities.
457110	Gasoline Stations with Convenience Stores	Fuel retailers that also operate convenience stores on the same premises.
456110	Pharmacies and Drug Retailers	Retail pharmacies, including chain drugstores and independent pharmacies that dispense prescription medications.

Table A.3: SafeGraph POI categories used in the covariate balance analysis. Each category corresponds to a 6-digit North American Industry Classification System (NAICS) code.

These six categories were selected to capture a diverse set of neighborhood amenities—spanning education, food access, healthcare, fuel, and social infrastructure—that are plausibly unaffected by scooter deployment mandates. Under the identification assumption, the density of these amenities should be similar on both sides of the mandate boundary, since the boundary cuts through neighborhoods rather than separating fundamentally different areas. Finding balance in these covariates supports the validity of our spatial regression discontinuity design.

Program	Outpatient Care Centers			Drinking Places			Elementary/Secondary Schools		
	Diff	Adj. p-value	p-value	Diff	Adj. p-value	p-value	Diff	Adj. p-value	p-value
ATL 2021	0.007	1.000	0.889	−0.668	1.000	0.337	−0.904	1.000	0.172
BAL 2021	0.006	1.000	0.960	1.078	1.000	0.282	0.292	1.000	0.802
CHI 2020	−0.047	1.000	0.824	0.385	1.000	0.519	0.447	1.000	0.539
DC 2020	1.278	1.000	0.158	0.205	1.000	0.904	−1.171	1.000	0.739
DEN 2021				−0.021	1.000	0.962	−0.259	1.000	0.629
KC 2021				−0.195	1.000	0.545	−0.061	1.000	0.924
LA 2021	−0.072	1.000	0.402	0.599	1.000	0.096	0.054	1.000	0.840
LOU 2021	0.574	1.000	0.358	−0.026	1.000	0.931	−1.675	0.894	0.039
MSP 2023	0.630	1.000	0.060	0.420	1.000	0.651	1.889	1.000	0.286
PDX 2021				1.001	1.000	0.257	0.230	1.000	0.713
SAC 2021	−0.058	1.000	0.642	0.123	1.000	0.623	0.118	1.000	0.816
SEA 2021	−0.175	1.000	0.415	−1.111	1.000	0.200	0.722	1.000	0.305
SF 2021	0.466	1.000	0.471	0.978	1.000	0.719	−2.278	1.000	0.240
	Gas Stations			Pharmacies			Supermarkets		
	Diff	Adj. p-value	p-value	Diff	Adj. p-value	p-value	Diff	Adj. p-value	p-value
ATL 2021	0.558	1.000	0.112	−0.161	1.000	0.603	−0.230	1.000	0.523
BAL 2021	0.545	1.000	0.299	3.162	0.014	0.002	2.847	0.162	0.002
CHI 2020	−0.192	1.000	0.585	−0.057	1.000	0.852	0.338	1.000	0.447
DC 2020	0.455	1.000	0.853	−3.330	0.950	0.059	−1.607	1.000	0.351
DEN 2021	0.627	0.986	0.052	−0.132	1.000	0.578	0.287	1.000	0.373
KC 2021	0.261	1.000	0.386	0.000	1.000	0.998	−0.390	1.000	0.091
LA 2021	0.138	1.000	0.439	0.244	1.000	0.257	0.240	1.000	0.295
LOU 2021	0.956	0.966	0.046	0.765	1.000	0.152	−0.218	1.000	0.661
MSP 2023	0.210	1.000	0.779	0.420	1.000	0.486	1.050	1.000	0.198
PDX 2021	−0.413	1.000	0.337	0.129	1.000	0.800	0.625	1.000	0.263
SAC 2021	−0.033	1.000	0.880	−0.345	1.000	0.184	0.029	1.000	0.906
SEA 2021	0.317	1.000	0.465	0.542	1.000	0.252	0.017	1.000	0.974
SF 2021	−0.682	1.000	0.338	−0.212	1.000	0.855	0.974	1.000	0.493

Table A.4: Differences in nearby amenities. The differences in the number of selected amenities—outpatient care centers, drinking places, schools, gas stations, pharmacies, and supermarkets—located within the inner ring relative to the outer ring for each program. For each amenity, we report the estimated difference, the Westfall–Young adjusted p-value, and the unadjusted p-value.

I Availability analysis results

Program	**Point estimate**	**p-value**	**Adj. p-value**	**Max-t p-value**
ATL 2021	−1.72	0.586	1.00	1.00
BAL 2021	+38.38	0.000	0.00	0.00
CHI 2020	−27.19	0.104	0.59	1.00
DEN 2021	+70.30	0.000	0.01	0.01
KC 2021	−3.65	0.598	1.00	1.00
LA 2021	−19.34	0.376	0.99	0.92
LOU 2021	−4.01	0.114	0.75	0.82
PDX 2021	+3.11	0.718	1.00	1.00
SAC 2021	+7.73	0.202	0.74	0.82
SEA 2021	−23.48	0.556	1.00	1.00
SF 2021	+8.91	0.098	0.58	1.00

Table A.5: Availability point estimates with unadjusted and adjusted p-values. The max-t p-value column reports inference using the maximum absolute t-statistic across programs; conclusions are identical to the FWER-adjusted p-values.

J Nested boundaries

To test whether policy effects vary with distance from the mandate boundary, we divided each mandate area into a sequence of concentric interior layers. Layers were created by buffering inward from the boundary in fixed-width increments of ϵ. Each step produced an annulus defined by the difference between two successive buffers, resulting in a set of non-overlapping interior layers that together cover the full mandate area.

For inference, each interior and corresponding exterior layer were treated as replacements for $\mathcal{A}|_\epsilon$ and $\mathcal{A}'|_\epsilon$ in equation (5). To generate placebo distributions, we randomized both the placement and orientation of the buffered mandate area, applying random translations and rotations within the city boundary. Transformations were retained only if the shifted area remained entirely within the pilot boundary. Each accepted draw was then intersected with the nested buffers to produce placebo versions of the interior layers.

We repeated this procedure 500 times per city and, stratifying pilots by the number of feasible layers, computed family-wise p-values by comparing observed statistics to the corresponding placebo distributions, adjusting for multiple outcomes and variation in layer depth across cities.

Pilot	**Depth**	**Est.**	***p***	Adj. ***p***	**Pilot**	**Depth**	**Est.**	***p***	Adj. ***p***
ATL 2021	1	-0.71	0.36	0.93	LOU 2021	1	-0.07	0.64	0.73
	2	0.90	0.32	0.93		2	0.00	0.91	0.91
	3	3.05	0.11	0.78		3	-0.02	0.61	0.81
CHI 2020	1	4.12	0.06	0.61		4	0.00	1.00	1.00
	2	-0.92	0.44	0.89	PDX 2021	1	-0.18	0.81	0.81
	3	2.29	0.15	0.73		2	-1.28	0.31	0.54
	4	-2.17	0.21	0.75		3	0.62	0.34	0.64
	5	1.00	0.54	0.89		4	3.16	0.04	0.26
	6	3.19	0.25	0.71		5	-2.67	0.09	0.29
KC 2021	1	-0.14	0.37	0.93	SAC 2021	1	-0.94	0.45	0.93
	2	-0.52	0.18	0.93		2	0.61	0.44	0.93
	3	0.03	0.39	0.93		3	0.83	0.34	0.93
LA 2021	1	0.43	0.66	0.93	SEA 2021	1	-1.71	0.42	0.53
	2	2.27	0.36	0.82		2	0.55	0.65	0.65
	3	-0.07	0.73	0.93	SF 2021	1	2.09	0.76	0.76

Table A.6: Point estimates and p-values by program and depth.

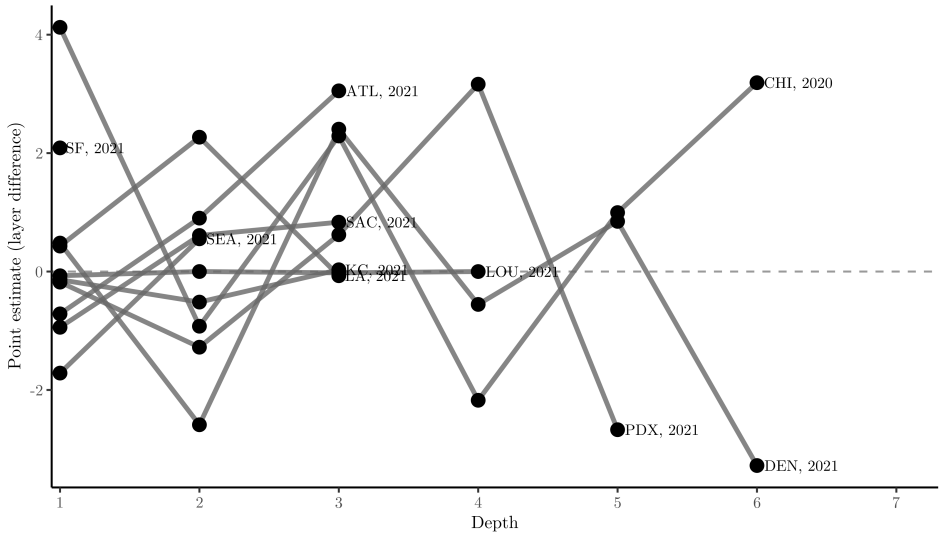


Figure A.2: No systematic pattern emerges as we move deeper into mandate areas.

K Compliance

Program	**Threshold**	**Median daily average availability**
ATL 2021	2%, 2%, 2%	5.79%, 4.14%, 3.39%
BAL 2021	3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3	7.00, 5.31, 5.61, 10.94, 5.68, 6.25, 7.08, 4.03, 12.23, 9.00, 3.91, 6.30
CHI 2020	2.5%, 2.5%	2.88%, 2.78%, 2.72%, 2.86%, 2.76%, 2.78%, 2.76%, 2.86%, 2.69%, 2.69%, 2.89%, 2.83%, 2.75%, 2.74%, 2.75%, 2.74%, 2.77%, 2.80%, 2.81%, 2.75%
DEN 2021	30%	34.76%
KC 2021	10%	4.63%
LA 2021	20%	21.03%
LOU 2021	20%	2.93%
PDX 2021	20%	17.08%
SAC 2021	20%	23.32%
SEA 2021	10%	15.91%
SF 2021	34, 94, 86, 12	69.69, 81.71, 53.77, 7.17

Table A.7: Each program's mandate threshold and median daily average availability across subareas.