

UC Berkeley

CUDARE Working Papers

Title

Uber versus Trains? Worldwide Evidence from Transit Expansions

Permalink

<https://escholarship.org/uc/item/29v1x3gc>

Authors

Gonzalez-Navarro, Marco

Hall, Jonathan D.

Wheeler, Harrison

et al.

Publication Date

2026-09-29

Uber versus Trains?

Worldwide Evidence from Transit Expansions*

Marco Gonzalez-Navarro[†]
Harrison Wheeler[‡]

Jonathan D. Hall[§]
Rik Williams[¶]

September 2026

Abstract

We study how ride-hailing and public transit interact – an open question with important implications for transportation policy and urban form. We develop a model of mode choice that features three key mechanisms: (i) *substitution*, whereby riders switch from Uber to transit near stations; (ii) *last-mile complementarity*, whereby Uber is used to access transit; and (iii) a dynamic choice-set effect, whereby outbound mode choice constrains return options, giving Uber *option value*. Using proprietary Uber trip data around 650 rail station openings worldwide, we estimate how ride-hailing responds when public transit improves. We document three patterns consistent with these mechanisms: Uber usage rises sharply within 100 meters of new stations (last-mile complementarity), declines locally among nearby residents and workers (substitution), and falls modestly overall for these users as local reductions are partially offset by increased Uber usage elsewhere (option value). Estimating the model on these moments, we find that all three channels are quantitatively important, with substitution the largest individual channel but last-mile complementarity and option value together more than offsetting it, so that Uber and transit are net complements. Policy implications include: (a) a ride-hailing tax induces users to switch to private vehicles rather than transit; (b) a subsidy for last-mile ride-hailing trips is more effective at increasing transit use when it targets the per-kilometer cost rather than the fixed cost; and (c) increases in transit use are best induced by a congestion charge that subsidizes transit fares rather than by policies relying on ride-hailing taxes and/or last-mile subsidies.

*We are grateful to Jonathan V. Hall, Mariya Shappo, and Felipe Vial for facilitating data access from Uber. We thank Chandra Krishnamurthy, Enrico Moretti, Lin Tian, and Wenxuan Xu for helpful feedback, as well as seminar participants at the 2023 UEA North America Conference, the 2024 LA Urban Conference, the 2026 AREUEA-ASSA conference, and the 2026 UEA European Conference. This work has received funding from the Social Sciences and Humanities Research Council of Canada [grant number 435-2019-0507], the European Union’s Horizon 2020 research and innovation programme [grant number 101022491], the NET Institute, and an interagency agreement between the US Department of Transportation and the National Science Foundation [award number 1559013].

[†]University of California Berkeley, marcog@berkeley.edu

[‡]Rotman School of Management, University of Toronto, harry.wheeler@rotman.utoronto.ca

[§]University of Alabama, jonathan.hall@ua.edu

[¶]Independent researcher, rjw297@gmail.com

1 Introduction

The global expansion of ride-hailing services has sparked a critical debate over their impact on cities. Policymakers and urban planners are grappling with how best to respond, with some countries, including Denmark, Hungary, and Bulgaria, even banning them outright. Making good policy decisions requires understanding how ride-hailing interacts with existing modes, especially public transit. These interactions have first-order welfare effects because each mode's price differs from its social marginal cost. Automobile-based modes, including ride-hailing, are too cheap because of their unpriced externalities such as congestion, pollution, noise, and crash risk,¹ while public transportation is too expensive because of economies of scale and pricing constraints.² Understanding the impact of ride-hailing matters beyond today's trips: new transportation technologies have repeatedly reshaped cities. Steam railways enabled London's expansion and highways hollowed out U.S. city centers.³ This question will only grow in importance as autonomous vehicles make on-demand, car-based travel far cheaper and more widespread than ride-hailing is today.

Predicting the impact of any transportation policy that affects ride-hailing requires understanding not just the aggregate relationship between ride-hailing and other modes, but also the mechanisms through which they interact. A tax on ride-hailing, for example, pushes some riders from ride-hailing to transit, as intended, but it also raises the cost of using ride-hailing to reach a station or as a backup for when transit fails. Whether the tax raises or lowers transit ridership depends on the relative strength of these mechanisms.

Separating these mechanisms requires observing how travelers respond when one mode becomes more attractive. A new train station is such a change. It is a discrete improvement in transit, on a known date and in a known place. Station openings also matter in their own right. Whether a new station pulls travelers out of their cars or only serves those who would have ridden anyway determines its effect on congestion and emissions, and how ride-hailing responds offers evidence on that question.⁴ This paper's contribution is to investigate how ride-hailing usage responds when a new train station opens, through which mechanisms, and what the answer implies for transportation policy.

To answer these questions, we develop a tractable dynamic model of travel mode choice, assemble proprietary data on Uber trips in 55 cities across 35 countries and on the 650 train station openings in those cities, and introduce a new empirical strategy. In our model, travelers make multiple trips during the day, choosing among personal automobile, Uber, and public transit. If they take transit, they choose how to get to the station: either by walking or by Uber. A key feature

¹See, respectively, [Currie and Walker \(2011\)](#) and [Yang et al. \(2020\)](#); and [Anderson \(2020\)](#); [Moretti and Wheeler \(2025\)](#); and [Ashenfelter and Greenstone \(2004\)](#).

²Fares nonetheless fall short of average cost, so agencies run deficits, and limits on subsidies keep fares from falling to marginal social cost. See [Parry and Small \(2009\)](#) and [Basso and Silva \(2014\)](#).

³See [Heblich et al. \(2020\)](#) and [Baum-Snow \(2007\)](#).

⁴The impact of a new train station on ride-hailing also bears on station design, from pickup and drop-off areas to parking requirements, and on the size of a station's catchment area; we return to these in Section 6.

is that the outbound mode choice constrains the traveler’s later options: whoever drives to work must drive home, and whoever leaves the car at home must use Uber or transit for every later trip. The model highlights three mechanisms through which transit improvements affect ride-hailing: substitution (travelers switch from Uber to transit), last-mile complementarity (travelers use Uber to reach transit), and option value (travelers leave the car at home to preserve the option of using the improved transit, and then rely on Uber for the trips transit serves poorly).

Our model predicts that the effect of a new station on travel behavior differs across trips and travelers, both near and far from the station. Accordingly, we estimate the effect of a new train station on ride-hailing for three quadrants: trips near the station by all travelers, trips near the station by travelers who live or work nearby, and citywide trips by those same travelers.

Our identification strategy exploits the sharp timing of rail transit station openings. Because rail expansions are planned more than a decade in advance ([Gendron-Carrier et al., 2022](#)), typically before Uber existed, their opening provides plausibly exogenous variation in transit access. In each quadrant, we compare the change in Uber trips after a station opens with the change for the same kind of trips about a kilometer farther from the station, whether that distance is where the trip starts or ends or where the traveler lives or works. In the first two quadrants, we estimate the effect separately by 100-meter concentric distance bands, tracing out how the response varies with distance. Comparing trips near the station with those only slightly farther away flexibly controls for local trends in Uber ridership.

We find evidence for all three mechanisms. Among all travelers, trips within 100 m of a new station rise by 48%, consistent with last-mile complementarity, while trips 200–700 m away fall by about 2%, consistent with substitution among travelers within walking distance. Restricting attention to travelers who live or work near the station removes most last-mile trips. Their trips near the station fall by 4–8% within 500 m of the station, confirming that substitution away from Uber is the central response among these travelers. Following those same travelers across the city, we find that their total Uber use falls by only 2%, because their trips elsewhere in the city rise by 6%. This is consistent with option value: travelers who leave the car at home to use the improved transit rely on Uber for the trips transit serves poorly. The pattern is similar across the 55 cities in our sample, with most cities close to the pooled estimates.

Several tests support a causal interpretation of these estimates. Our dynamic difference-in-differences design relies on parallel trends, and event studies show no differential pre-trends. Because our estimates are relative to the comparison group, we also test for effects there; we find that effects decay toward zero with distance, and a design using only the timing of openings finds no detectable change in the comparison group. The estimates are robust to placebo opening dates, alternative specifications, geographies, and trip definitions, and we find no evidence that Uber changed its pricing when stations opened.

To quantify the mechanisms and evaluate policy counterfactuals, we estimate the model’s parameters using moments from the reduced-form evidence. Because mode shares vary with distance to transit, changes in Uber use reveal substitution across modes: near stations, Uber

primarily competes with transit; farther away, it primarily competes with cars. This allows us to infer mode-share changes for all modes using evidence on Uber trips alone.

We first use the model to quantify the importance of the three mechanisms. For travelers in the station’s catchment area, substitution reduces Uber ridership by 1.55%, while last-mile complementarity and option value raise it by 0.92% and 1.18%, respectively. Substitution is the largest single channel, but the two complementarities together outweigh it, so on net Uber and transit are complements.

We also use the estimated model to evaluate policy counterfactuals, each of which operates differently across the substitution, last-mile, and option-value channels. Ride-hailing taxes reduce transit ridership, but uniform ride-hailing subsidies do so as well, as travelers increasingly substitute from transit to Uber. By contrast, subsidies targeted specifically to last-mile access increase transit ridership by more than 5.5%. A policymaker hoping to increase transit use in a fiscally neutral way would fund both fixed and marginal-cost last-mile subsidies with a congestion tax, while not levying a ride-hailing tax. However, if direct transit subsidies are also available, they are a more effective tool than last-mile subsidies. Banning ride-hailing outright, as some countries have, lowers transit ridership by about 10%, because most former Uber riders switch to driving rather than to transit.

This paper makes four contributions. First, we provide the first evidence on how new train stations impact ride-hailing.⁵ This builds on existing work that asks the reverse question: how ride-hailing affects transit. In a meta-analysis, [Cairncross et al. \(2025\)](#) shows that estimates across this literature disagree in sign and magnitude, and that every existing study uses data from a single country, 69% of them the United States.⁶ Our evidence comes from 55 cities in 35 countries, and the response is similar across them, so the mechanisms we document are not specific to one country. Our estimated model also answers the reverse question: the ban counterfactual implies that ride-hailing and transit are net complements.

Second, we develop a dynamic model of mode choice that captures substitution, last-mile complementarity, and option value together. This builds on [Agrawal and Zhao \(2023\)](#), who model last-mile complementarity in a monocentric city to study optimal ride-hailing taxation, but in their model travelers choose a mode for each trip separately, so option value cannot arise.

Third, we provide the first estimates of the relative magnitudes of substitution, last-mile complementarity, and option value. Designing good policy requires these magnitudes; for example, a tax on ride-hailing raises transit ridership if substitution dominates, and lowers it if the two complementarities do.

⁵In contemporaneous work, [Bosker et al. \(2025\)](#) document last-mile use of motorcycle ride-hailing following the inauguration of Jakarta’s first subway line. Our results also add to the literature on how public transit shapes urban economic geography ([Billings, 2011](#), [Gonzalez-Navarro and Turner, 2018](#), [Gendron-Carrier et al., 2022](#), [Gupta et al., 2022](#), [Zárate, 2024](#), [Tsivanidis, 2026](#)).

⁶Complementary to our work is [Christensen et al. \(2026\)](#), who study how travel patterns change with discounts to ride-hailing services in Cairo, Egypt. Our paper is also related to a broader literature on the economic effects of ride-hailing, including its effects on traffic safety ([Greenwood and Wattal, 2017](#), [Teltser et al., 2021](#), [Barrios et al., 2023](#), [Barreto et al., 2021](#), [Anderson and Davis, 2026](#)), congestion ([Tarduno, 2021](#)), labor markets ([Angrist et al., 2021](#), [Hall and Krueger, 2018](#)), and housing costs ([Gorback, 2026](#)).

Fourth, we show that transportation policy must account for all three mechanisms, not substitution alone. This builds on existing studies of ride-hailing taxes and other transportation policies.⁷ For example, a ride-hailing tax, which under substitution alone could only raise transit ridership, lowers it in our estimates by making last-mile access more costly and weakening option value, while congestion pricing raises it because it targets cars directly.

The remainder of this paper is structured as follows. Section 2 presents the model, Section 3 the data, Section 4 the empirical strategy, Section 5 the reduced-form results, and Section 6 the model estimation and counterfactuals. Section 7 concludes.


2 Model

In this section, we develop a dynamic discrete-choice model of travel mode choice. We keep the model analytically tractable so that it clearly illustrates how substitution, last-mile complementarity, and option value affect choices at different distances from a transit stop and when transit service improves. Our model’s predictions guide our reduced-form empirical analysis. At the same time, the model is flexible enough that we can estimate its parameters, and use it to quantify the importance of each mechanism and evaluate counterfactual transportation policies.

2.1 Model setup

Consider a single public transit route with two stops and a risk-neutral representative traveler who must travel from an origin d kilometers from the first stop to a destination near the second. Later, she will make the return trip, however, with probability p_s , she will first need to make another trip. Figure 1 depicts how trips are linked between origin and destination in our model.

Figure 1: Model geography and trips



Notes: This figure depicts the model geography and the trips taken by the representative traveler.

⁷See, respectively, [Agrawal and Zhao \(2023\)](#), [Tarduno \(2025\)](#), and [Leccese \(2025\)](#); and [Basso and Silva \(2014\)](#), [Barwick et al. \(2024\)](#) and [Almagro et al. \(2024\)](#).

For each trip, the traveler chooses which transportation modes to use, with her outbound mode choice constraining her options for later trips. She has three options for her primary mode: personal automobile, Uber, and public transit. If she uses public transit, she chooses an access mode to travel the d kilometers to the transit stop: walking or Uber. If she drives her personal automobile for the outbound trip, she must also use the automobile for the return journey, but she may use any mode for the additional trip during the day. If she does not drive for the outbound trip, she cannot use a personal automobile for the remainder of the day's trips.

We define the following parameters. Let the generalized in-vehicle base utility of mode m be denoted by α_m for the outbound and return journeys, and β_m for the additional trip during the day. Using access/egress mode a has fixed cost f_a and marginal cost c_a per kilometer of distance d . We assume walking has no fixed cost, but has a higher marginal cost than taking Uber to the transit stop ($f_{\text{walk}} = 0$; $c_{\text{walk}} > c_{\text{Uber}}$).

The utility of a one-way trip from the origin to the destination depends on the distance to the transit stop (d), the primary transportation mode chosen ($m_\ell \in \{\text{car}, \text{transit}, \text{Uber}\}$ for period $\ell \in \{1, 2, 3\}$), the access mode chosen if using transit ($a_\ell \in \{\text{walk}, \text{Uber}\}$ for $\ell \in \{1, 3\}$), and a random shock over mode preferences. For notational purposes, let $a_\ell = \emptyset$ whenever $m_\ell \neq \text{transit}$.

We model each mode's shock as the sum of a persistent component and a transitory component. The persistent component, η_m , is realized at the start of the day and captures shocks to a traveler's preferences across travel modes that persist across both outbound and return trips. For example, a traveler may have a unique preference for driving, or today's destination may be unusually well suited to traveling by public transit. The η_m are drawn from a mean-zero bivariate normal distribution, with η_{car} normalized to zero.⁸

$$\begin{pmatrix} \eta_{\text{Uber}} \\ \eta_{\text{transit}} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\text{Uber}}^2 & \phi \sigma_{\text{Uber}} \sigma_{\text{transit}} \\ \phi \sigma_{\text{Uber}} \sigma_{\text{transit}} & \sigma_{\text{transit}}^2 \end{pmatrix} \right), \quad \eta_{\text{car}} \equiv 0.$$

The parameters σ_{Uber} and σ_{transit} govern the dispersion of persistent shocks. The parameter $\phi \in (-1, 1)$ governs the correlation between these shocks, capturing whether persistent preferences over Uber and transit (relative to personal car) tend to move together or in opposite directions.

The transitory components, ϵ_{ma1} , ϵ_{m2} , and ϵ_{ma3} , are realized at the start of each period and capture shocks to a traveler's preferences that only affect the current trip. For example, the next bus may be arriving shortly, or it may be raining, making the walk to the transit stop less appealing. The transitory shocks are independent and identically distributed random shocks with a type I extreme value distribution.

With these definitions, the utility of the outbound trip ($\ell = 1$) and the return trip ($\ell = 3$) are given by

$$U_\ell(m, a; d) = \alpha_m - \mathbb{1}_{m=\text{transit}}(f_a + d \cdot c_a) + \eta_m + \epsilon_{ma\ell}, \quad \ell \in \{1, 3\}. \quad (1)$$

⁸When we try to fit the parameters of our model to the data, a more relevant interpretation of the η_m is persistent preferences across the larger population of travelers. Also, we do not explicitly model the fact that travelers' true destinations often lie beyond the public transit terminus. This omission is absorbed by η .

With probability p_s , the traveler needs to make the intermediate trip during the day. The utility of the intermediate trip during the day only depends on the mode chosen and a transitory shock, ϵ_{m2} , which is drawn from the same distribution as the ϵ_{mal} , and is given by

$$U_2(m_2) = \beta_{m_2} + \epsilon_{m2}. \quad (2)$$

Two features of the model are worth reiterating. First, the traveler has all modes available for the intermediate trip only if they chose car for the outbound trip. Otherwise, they can only choose between public transit and Uber for the intermediate trip. Second, the traveler *must* return via car if they drove for the outbound trip.

The traveler chooses a sequence of modes to maximize expected utility. We derive the expected utility from each period's mode choice using backward induction, and recover closed-form logit choice probabilities. Finally, we calculate unconditional probabilities by integrating over the persistent shocks η_m . The details of these calculations are in Appendix C.

2.2 Theoretical implications from improving public transit

We use the model to assess how opening a new train station affects Uber ridership. We model the new train station as an improvement in the base utility of public transit, relative to car (α_{transit}). The change in α_{transit} bundles all the improvements that come with a new train station: the speed and frequency advantage of rail over the pre-existing bus service, complementary adjustments to feeder routes and schedules, and any easing of road congestion.⁹

The model highlights three mechanisms through which transit improvements affect the demand for ride-hailing: *substitution* (riders switch from Uber to transit), *last-mile complementarity* (riders use Uber to access the new train station), and *option value*. Option value is a dynamic choice-set effect, and occurs because the traveler's outbound mode choice constrains her options later in the day: if she drives to work, she must drive home; if she doesn't drive to work, then she cannot drive for any other trips. Because the traveler is forward-looking, these constraints influence her outbound decision. As *dynamic choice-set effect* names this object precisely but reads awkwardly, we retain it through the formal results below, then continue with *option value* for the remainder of the paper.

We can interpret the dynamic choice-set effect in two ways, each emphasizing a different temporal perspective. Viewed from the first period, it is an option-value effect: the forward-looking traveler may now choose Uber over driving for her outbound trip precisely because doing so preserves the option of taking transit home.¹⁰ Viewed from later in the day, it is a mode lock-in effect: the outbound choice changed the set of available modes. When improved transit induces a traveler to switch from driving to transit for her outbound trip, she is no longer locked into driving home. Uber is now available for the return journey, increasing Uber ridership. These

⁹Improvements in road congestion affect car and Uber similarly, leaving α_{Uber} unchanged.

¹⁰A closely related interpretation is an insurance effect: the traveler is more willing to take transit outbound knowing that if transit is delayed or disrupted, she can take Uber home instead.

are not distinct mechanisms but alternative descriptions of the same underlying dynamic; the outbound mode choice constrains the traveler's later options.

We now formally decompose the impact of a transit improvement on total Uber use into these three channels and sign each. Substitution and the dynamic choice-set effect also increase *transit* use; we classify the resulting Uber access and egress trips as last-mile complementarity. Because the channels compete, the net effect of a transit improvement on Uber ridership is ambiguous.

Proposition 1 (Three-channel decomposition) *For each period $t \in \{1, 2, 3\}$, the derivative of total Uber ridership with respect to α_{transit} , conditional on (d, η) , decomposes as*

$$\frac{\partial P_t(\text{Uber}_{\text{total}}|d, \eta)}{\partial \alpha_{\text{transit}}} = \underbrace{S_t(d, \eta)}_{\substack{\text{Direct} \\ \text{substitution} \\ \leq 0}} + \underbrace{D_t(d, \eta)}_{\substack{\text{Dynamic} \\ \text{choice-set effect} \\ \geq 0}} + \underbrace{L_t(d, \eta)}_{\substack{\text{Last-mile} \\ \text{complementarity} \\ \geq 0}},$$

For $t \in \{1, 3\}$, there exist parameter values for which the derivative is strictly negative, and others for which it is strictly positive. For $t = 2$, $S_2 \equiv L_2 \equiv 0$, and so the derivative reduces to $D_2 \geq 0$.

The model also maps marginal mode switches directly to the three channels. In doing so, we treat each combination of transit and an access mode as a distinct choice, separating “transit-by-Uber” from “transit-by-walk”.

Corollary 1 (Mapping mode switches to mechanisms) *Conditional on (d, η) , the marginal response of mode probabilities to an increase in α_{transit} maps to the three channels of Proposition 1 by mode pair as follows.*

1. From Uber to transit-by-walk: direct substitution only.
2. From car to transit-by-Uber: last-mile complementarity only.
3. From Uber to transit-by-Uber: direct substitution and last-mile complementarity.
4. From car to Uber: dynamic choice-set effect only.

In the intermediate trip (period 2), there is no Uber access leg, so clauses 2 and 3 do not arise. Because the station opening does not directly change intermediate-trip flow utilities, clause 1 does not arise either; only the car-to-Uber margin in clause 4 remains. Appendix C.2 contains the full proof of Proposition 1 and Corollary 1.

The following proof sketch shows how the three channels arise for the outbound trip. Suppressing the conditioning on (d, η) , the probability that a traveler uses Uber is $P_1(\text{Uber}) + P_1(\text{transit-by-Uber})$. The first term, $\partial P_1(\text{Uber})/\partial \alpha_{\text{transit}}$, operates through the Uber-vs-transit and Uber-vs-car margins. Let $V_1(m)$ denote the value of choosing primary mode m for the

outbound trip, including both period-1 utility and continuation value. Then

$$\begin{aligned} \frac{\partial P_1(\text{Uber})}{\partial \alpha_{\text{transit}}} &= P_1(\text{Uber}) P_1(\text{transit}) \cdot \frac{\partial(V_1(\text{Uber}) - V_1(\text{transit}))}{\partial \alpha_{\text{transit}}} \\ &\quad + P_1(\text{Uber}) P_1(\text{car}) \cdot \frac{\partial(V_1(\text{Uber}) - V_1(\text{car}))}{\partial \alpha_{\text{transit}}}. \end{aligned} \quad (3)$$

The Uber-vs-transit margin carries only direct substitution. Choosing Uber or transit for the outbound trip leaves the traveler with the same choice sets in periods 2 and 3. A transit improvement therefore changes the relative values of Uber and transit only through period-1 utility, where transit gets a unit increase relative to Uber:

$$\frac{\partial(V_1(\text{Uber}) - V_1(\text{transit}))}{\partial \alpha_{\text{transit}}} = -1.$$

The Uber-vs-car margin carries only the dynamic choice-set effect. Neither mode's period-1 utility depends directly on α_{transit} , so the relative value changes only through continuation values:

$$\frac{\partial(V_1(\text{Uber}) - V_1(\text{car}))}{\partial \alpha_{\text{transit}}} = 1 \cdot \left[P_3(\text{transit} \mid m_1 \neq \text{car}) - \overbrace{P_3(\text{transit} \mid m_1 = \text{car})}^{=0} \right] \equiv \Phi(d, \eta).$$

The bracketed difference is the increase in later transit use when the traveler cannot choose car because she did not drive outbound, which is simply the probability of returning by transit conditional on not driving in period 1. The term is non-negative, so $\Phi \geq 0$.

We next consider last-mile complementarity captured by $\partial P_1(\text{transit-by-Uber}) / \partial \alpha_{\text{transit}}$. To evaluate this derivative, write

$$P_1(\text{transit-by-Uber}) = P_1(a_1 = \text{Uber} \mid m_1 = \text{transit}) \cdot P_1(\text{transit}),$$

where the first factor is the probability of choosing Uber as the access mode conditional on choosing transit and the second is the probability of choosing transit. As the access-mode share is unaffected by α_{transit} , the response depends on the change in transit use:

$$\begin{aligned} \frac{\partial P_1(\text{transit})}{\partial \alpha_{\text{transit}}} &= P_1(\text{transit}) P_1(\text{Uber}) \cdot \frac{\partial(V_1(\text{transit}) - V_1(\text{Uber}))}{\partial \alpha_{\text{transit}}} \\ &\quad + P_1(\text{transit}) P_1(\text{car}) \cdot \frac{\partial(V_1(\text{transit}) - V_1(\text{car}))}{\partial \alpha_{\text{transit}}} \\ &= P_1(\text{transit}) P_1(\text{Uber}) + P_1(\text{transit}) P_1(\text{car}) (1 + \Phi). \end{aligned}$$

The first term is transit's gain from Uber through direct substitution. The second is transit's gain from car, with the 1 due to the direct increase in transit's period-1 utility and the Φ due to the dynamic choice-set effect. We note that *transit usage* increases due to substitution and dynamic choice-set effects, and the increase in transit usage in turn increases last-mile Uber use.

For the return trip, the same three channels operate, although the dynamic choice-set effect is backward-looking: travelers who did not drive outbound cannot drive home. For the intermediate trip, there is no access mode or substitution, and only the dynamic choice-set effect remains.

Despite its simplicity, the model can generate qualitatively distinct spatial patterns in Uber ridership following a station opening under different parameter values. Figure 2 presents three versions of the model, chosen to illustrate the range of empirical patterns it can rationalize.¹¹ Each column corresponds to a different parameterization. The first row reports average outbound and return trip mode shares under baseline preferences. The second row shows how those mode shares respond to a 0.5-unit increase in α_{transit} , which captures an improvement in public transit. The third row plots the log change in Uber trips induced by the same transit improvement. All figures report values for locations within 1200 meters of the station on the x-axis. Parameter values for each simulation are described in Section C.5.

In the first parameterization, low fixed and marginal Uber access costs strengthen the last-mile channel, generating an increase in Uber use at all distances. In the second, high fixed access costs weaken the last-mile channel, while greater dispersion in persistent Uber preferences strengthens substitution, producing declines at all distances. The third combines low access costs with high preference dispersion, generating a sharp increase near the station and declines at intermediate distances. This final case generates the same qualitative pattern we observe empirically in Section 5: a sharp increase in Uber use near the station alongside declines at intermediate distances.

2.3 Relating the model to data on Uber trips

The model yields clear predictions for how Uber ridership should respond to the opening of a new transit station as a function of the distance d between the station and trip origin.¹² To bring these predictions into our setting, we must therefore recover d from the data. For the outbound and return trips, the trip itself reveals this distance by originating or terminating near the station. Because the intermediate trip can occur anywhere in the city, it reveals nothing about d ; we must instead recover it from the traveler. We identify travelers whose trips typically start or end near the station, which in practice captures people who likely live or work in the station’s catchment area. Accordingly, we classify our empirical margins by where *trips* and where *travelers* are located:

Table 1: Classification of empirical margins by trip and traveler proximity

Trips/People	Near station	Anyone
Near station	II	I
Everywhere	III	IV

¹¹We assign first- and last-mile trips to the station itself (located at $d = 0$) but allow for measurement error in the data. Specifically, we assume that the origins and destinations of first- and last-mile trips follow an exponential distribution centered at the station, so that the density of such trips decays with distance.

¹²To map model-generated probabilities to our Uber trips data, we aggregate the model’s choice probabilities across origins at varying distances from each station, weighting by the geographic area at each distance band. Appendix C.4 describes this mapping in detail.

- I:** *Trips near station by all travelers.* These trips correspond to the model’s outbound and return trips. Because access and egress legs begin or end at the station itself, this margin captures all last-mile trips, making it the most informative about last-mile complementarity.
- II:** *Trips near station by travelers near station.* By restricting attention to travelers residing or working within the station’s catchment area, this sample isolates responses among local users. This helps minimize the influence of last-mile complementarity, leaving substitution and option value.
- III:** *All trips by travelers near station.* Following the same local travelers across the entire city adds the intermediate trip, which can occur anywhere and which Proposition 1 shows responds through option value alone.
- IV:** *All trips by all travelers.* This margin corresponds to aggregate, citywide effects; because our identification strategy is station-level, we do not measure it directly.

Together, these margins cover all three of the model’s trips and all three of its channels. In Section 6, we estimate the model by fitting it to the causal evidence across them.

3 Data

To investigate the effect of new transit stations on Uber ridership we require data describing new transit station locations and dates of opening, as well as panel data on Uber ridership and riders near these stations. We use data on Uber ridership constructed from Uber’s trip database. A unique feature of these data is that they allow us to observe interactions between Uber and transit that transportation data typically cannot capture: highly disaggregated transit data are unavailable in most cities, and even station-level ridership data reveal neither how riders access the station nor how far they travel to reach it. Our data on new transit stations are the result of primary data collection. We describe these datasets and their construction below.

3.1 Uber ridership

We use Uber’s trip database to calculate the number of trips starting or stopping within a given distance band from the transit station in a month, for example, all trips within 100–200 meters (m). We do this for 12 different 100 m bands from 0–100 m up to 1100–1200 m. To avoid double-counting, each Uber pickup or drop-off is allocated to the station to which it is closest. This means the distance bands are not always perfect circles. We use this dataset for our analysis of station-adjacent trips among all travelers (Quadrant I of Table 1).

We also make use of a fixed area hexagon in robustness exercises as well as to identify individuals with a point-of-interest (POI) near the station. The hexagons have an average edge length of 174 m and an average area of 0.105 km². We define a person’s POI as the hexagon containing the most trip origins and destinations for the user in the 12 to 7 months prior to the

station opening, subject to the person making at least three trips to or from the hexagon. The set of travelers with a POI whose centroid is within 1200 m of a new station makes up our *POI sample*. We use this dataset for the nearby-traveler analyses (Quadrants II and III of Table 1). Both hexagons and concentric ring spatial units are depicted for train openings in Toronto in Figure A1.¹³ We emphasize that throughout our analysis, as per our data usage agreement, we have access to data aggregated from the individual trip and rider databases, but not the underlying microdata itself.

Our Uber trip data spans January 2012 to December 2018, but note that Uber was expanding in this time period so that not every city has trip data from 2012. This provides us with between 15 and 79 months of Uber trip data for each transit expansion, with a mean and median of 61 months. We exclude cities where Uber availability began, ended, or paused within six months of the rail expansion. Under the data use agreement, New York City and China are excluded and geography-month observations with fewer than five trips are set to zero to protect privacy.

3.2 Transit stations

To collect data on new transit stations, we start with a list of cities Uber operates in worldwide, and for each city, find all rail transit stations that opened after Uber entered the city. We obtain data on subway openings through 2017 from [Gendron-Carrier et al. \(2022\)](#), extend this dataset to include light rail and commuter rail, and update it through 2018 using online sources such as www.urbanrail.net, www.wikipedia.org, and news sites. We limit attention to rail stations as these are better documented than bus stops. For each station, we record opening date, latitude and longitude, station name, whether it is the terminal station, city, and country. We define the exact latitude and longitude of each station using Google Maps. We also find the locations of the pre-existing stations that had been the terminal stations before the transit expansion. We restrict our sample to transit that does not interact with traffic. This excludes tram and streetcar station openings. We are left with 650 station openings across 55 cities and 35 countries.

Figure 3 plots the geographic distribution of cities used in our analysis. The size of each circle is proportional to the number of new stations that opened between Uber’s entry date and 2018. The map reveals that transit openings in our sample are heavily concentrated in India and Southeast Asia, Europe, and North America. Table A1 reports the cities in our sample that expanded their rail transit between the date Uber entered and 2018. Figure A2 shows the time series of rail expansion events and number of stations opened, showing there is a notable drop in openings in January and February, but they are otherwise fairly uniform over time of year.¹⁴

¹³These hexagons are defined using the H3 hexagonal hierarchical geospatial indexing system, designed by Uber. The advantage of hexagons is that they have a uniform area, unlike the concentric rings we use in our main analysis. For more information, see <https://h3geo.org/>. We have data on all hexagons whose centroids are within 1,200 m of a transit station.

¹⁴Figure A3 plots the distribution of days and months in our sample, and shows no discernible patterns outside of the drop in transit openings during the winter.

3.3 Additional variables

We use additional data to explore heterogeneity in our treatment effects. We define the center of each city using [United Nations, Department of Economic and Social Affairs, Population Division \(2018\)](#). To find the footprint of each city, we use the Global Human Settlement Layer ([Schiavina et al., 2023](#)) which defines urban clusters “via a logic of cell clusters population size, population and built-up area densities as defined by the stage I of the Degree of Urbanisation.” This is measured for 1 km x 1 km cells. For the purpose of finding average values, we define the footprint of a city as the urban cells within 25 km of the city center. We use the city center and footprints to determine how close a new train station is to the city center. We measure “closeness” in two ways. First, we measure the straight-line distance, and second, we measure the travel time using public transit.¹⁵

3.4 Summary statistics

Our analysis requires us to observe Uber ridership at a transit station for some time before and after the transit station opens. Although each regression uses the full available history of Uber ridership around each station, we focus our estimates on a 13-month event window: six months before the station opening, the opening month, and six months after. We exclude transit stations for which this full window is not observed.

Table 2 reports summary statistics for the two main Uber samples used in the analysis. Panel A describes station-adjacent trips over the 13-month analysis window (Quadrant I). The first five rows report mean and median values for the number of pickups, the number of drop-offs, total kilometers traveled, average trip length, and average fare per trip. The next row reports the share of distance-bin-by-month observations with zero Uber trips. The remaining rows report the share of trips that occurred during weekday mornings, weekday afternoons, Friday and Saturday nights, and weekend days. We report these summary statistics across the full station-by-distance-bin sample in Column (1), and for selected distance bins in Columns (2) through (5).

On average, the station-adjacent data contain 1,663 Uber pickups and 1,656 drop-offs in each 100 m distance-bin-by-month cell. Uber trips are 11.20 kilometers long on average, with an average fare of \$7.84. Within 100 m of a station opening, 25% of observations have zero Uber trips; the corresponding share in the 1100–1200 m comparison band is 15%. The median values show that the data are severely right-skewed toward larger cities and denser locations. The existence of zeros and outliers motivates our use of the Poisson Pseudo-Maximum Likelihood estimator, which we discuss in further detail in Section 4.

Panel B reports summary statistics for the POI sample used in the citywide analysis (Quadrants II and III). These observations are defined by station-month cells and groups of riders with a fixed point of interest near the station. We report citywide trips, local trips, non-local trips, the number of riders, average trip length, average fare per trip, and the share of observations with zero trips.

¹⁵In calculating this, we remove walking time from the last transit stop and the city center.

We define local trips as those that originate or terminate within the POI distance band, whereas non-local trips are those that occur farther than the outer edge of the distance band.¹⁶ The table reports these statistics for the full POI sample and separately for riders with POIs located 0–500 m, 500–800 m, and 800–1200 m from the station. On average, riders in each POI sample group take 5,441 citywide trips per station-month cell, of which 2,558 are local and 2,404 are non-local. As in the station-adjacent trip data, the POI sample is right-skewed, with mean trip counts substantially larger than medians.

4 Methodology

Section 2.3 introduced three empirical settings in Table 1, each providing evidence for the mechanisms in the model. We begin by outlining our empirical strategy for station-adjacent Uber trips among all travelers (Quadrant I). We then describe how we adapt this design for travelers with a fixed location near the station, examining both their station-adjacent trips (Quadrant II) and their citywide Uber usage (Quadrant III).

Our identification strategy uses a dynamic difference-in-differences approach, leveraging high-frequency Uber trip data and the precise timing of station openings. The first difference compares Uber ridership before and after a station opens, within a given distance band around the station. The second difference controls for broader trends in Uber usage over time – an important adjustment, since Uber ridership was growing rapidly throughout our sample period. To flexibly account for these local time trends, we use trips occurring at distances between 1100–1200 meters from the station as a reference group.

The nearby-traveler analyses restrict attention to individuals with a point of interest (POI) – such as home or work – near the station. For station-adjacent trips (Quadrant II), this restriction removes most last-mile trips, sharpening our focus on substitution and option value. For citywide trips (Quadrant III), we instead track Uber use across the city among the same population, allowing us to examine option value in isolation. In both cases, we use individuals with POIs farther from the station as the reference group.

We first define the core research design. Let y_{dit} denote the number of station-adjacent trips beginning or ending in distance band d around transit station i during month t . Denote the month that station i opens as t'_i and define the time since the station opened (“relative time”) as $\tau_{it} = t - t'_i$, noting that τ_{it} is negative before the station opens. In the month a station opens, $\tau_{it} = 0$; for stations that open at the start of the month, nearly the entire month is treated, while for stations that open at the end of the month, nearly the entire month is untreated.

¹⁶We prefer this measure of local and non-local trips because it is conceptually intuitive. However, while these measures make up the vast majority of citywide trips, they do not sum to the citywide total: each POI group’s local and non-local measures exclude trips that are closer to the station than the group’s own distance band. For example, for the 500–800 m POI group, trips originating or terminating within 0–500 m of the station are counted in neither the local nor the non-local measure; the same applies to the 800–1200 m group with trips within 0–800 m. Because these excluded trips are less than 9% of citywide trips, and because we do not explicitly model them, we exclude them from our local and non-local measures; however, for completeness we do include them when we consider the total citywide effect.

We model the effect of a new transit station opening on Uber ridership in distance band d^* using observations from both distance band d^* and the farthest distance band $\bar{d} = 1200$ m. The dynamic difference-in-differences specification is as follows:

$$y_{dit} = \gamma_{it} + \delta_{di} + \sum_{j \in \{-6, -5, \dots, 6\} \setminus \{-2\}} \alpha_{d^*j} \times \mathbb{1}_{\tau_{it}=j} \times \mathbb{1}_{d=d^*} + (\beta_{1,d^*} \times \mathbb{1}_{\tau_{it} < -6} + \beta_{2,d^*} \times \mathbb{1}_{\tau_{it} > 6}) \mathbb{1}_{d=d^*} + \epsilon_{dit},$$

$$\forall d \in \{d^*, \bar{d}\}, \quad (4)$$

where $\mathbb{1}_{\tau_{it}=j}$ is an indicator function for whether the given observation occurs j months after the station opens, $\mathbb{1}_{d=d^*}$ is an indicator function for whether the given observation is at distance d^* , γ_{it} is a station-time fixed effect, and δ_{di} is a station-distance fixed effect. The index j runs from -6 to 6, excluding -2, so the second month prior to the transit station openings is the reference category.¹⁷

The coefficients of interest are the α_{d^*j} , which in our estimation can be interpreted as log changes in Uber trips j months after a new transit station opens.¹⁸ For $j \geq 0$, the α_{d^*j} are the treatment effects we seek to measure at distance d^* in each month after the opening. For $j < 0$, the α_{d^*j} allow us to test whether there are pre-trends in Uber trips before stations open. The coefficients β_{1,d^*} and β_{2,d^*} are event study coefficients for before and after our treatment window. This specification leverages up to 79 months of data per station, while isolating variation around the opening date.

We estimate the following model via Poisson Pseudo-Maximum Likelihood (PPML) (Silva and Tenreiro, 2006, 2011). The PPML estimator has several advantages for our setting. First, it allows for a natural interpretation of our treatment effects as percentage changes. This is an attractive feature for our setting since farther away distance bands are mechanically larger in area, and thus contain more trips. Second, it handles zeros and outliers in the dependent variable well. Many 100 m distance band by month observations will contain zero trips. As evidenced in Table 2, 25% and 15% of observations have zero trips in the 0–100 m and 1100–1200 m distance bands, respectively. Lastly, it handles high-dimensional fixed effects well relative to other non-linear models. Our main specification is fully saturated with station by distance and station by calendar time fixed effects. Throughout, we cluster our standard errors at the station level.

We often wish to aggregate the treatment effect to a single coefficient, which is helpful when reporting regression results across multiple specifications in a single table or figure. We will want to report these coefficients for each distance band, and so we rely on the following

¹⁷We use the second month prior to the opening as the reference category to test for anticipation effects.

¹⁸For effects smaller than about 0.10 log points, we report $100 \times \alpha_{d^*j}$ as an approximate percentage change. Larger effects are converted exactly as $100 \times [\exp(\alpha_{d^*j}) - 1]$.

difference-in-difference specification:

$$\begin{aligned}
y_{dit} = & \gamma_{it} + \delta_{di} + \sum_{d^* \neq 1200 \text{ m}} \alpha_{d^*} \times \mathbb{1}_{\tau_{it} \in \{1, \dots, 6\}} \times \mathbb{1}_{d=d^*} \\
& + \sum_{d^* \neq 1200 \text{ m}} (\beta_{1,d^*} \times \mathbb{1}_{\tau_{it} < -6} + \beta_{2,d^*} \times \mathbb{1}_{\tau_{it} > 6} + \beta_{3,d^*} \times \mathbb{1}_{\tau_{it} \in \{-2, -1, 0\}}) \mathbb{1}_{d=d^*} + \epsilon_{dit},
\end{aligned}
\quad \forall d, \quad (5)$$

In this regression, the omitted category is the third through sixth months before the station opened and the coefficients of interest are the α_{d^*} . The α_{d^*} capture the average percentage effect of a station opening on Uber trips at distance d^* relative to $\bar{d} = 1200$ m, for months $\tau_{it} \in \{1, \dots, 6\}$ relative to months $\tau_{it} \in \{-6, -5, -4, -3\}$. We separately control for the effect two months prior through the month of the station opening. This specification mitigates concerns over the partial treatment of month $\tau_{it} = 0$ and soft openings which may have had an impact on Uber usage prior to the station’s official open date.

We have two core identifying assumptions. First, we assume that local governments do not time the opening of new transit stations to coincide with a sharp break in Uber ridership. Indeed, [Gendron-Carrier et al. \(2022\)](#) find that it typically takes 11 years between when the plan is approved for a new subway and the opening date. This means that the vast majority, and maybe even all, of the transit openings in our data were approved before Uber existed. Moreover, we have chosen a narrow study window – six months before and after the station opening – for our analysis. Any effects we observe are most likely driven by the sudden opening of public transit, rather than other changes to public policy. Second, we make the standard difference-in-differences assumptions that Uber ridership at all distances would have moved in parallel in the absence of a new transit station opening. A strength of our approach is the granularity of our data, which uses nearby distances to provide a more plausible counterfactual. We will present evidence that Uber ridership in the 1100–1200 m band moved in parallel to closer distances prior to the station opening.¹⁹

Given our station openings occur over the period 2013–2018, an additional concern might be whether our setting is subject to the usual issues of difference-in-differences estimators in the presence of staggered treatment timing and treatment effect heterogeneity ([Goodman-Bacon, 2021](#), [Callaway and Sant’Anna, 2021](#), [Sun and Abraham, 2021](#), [Borusyak et al., 2024](#)). Our main specification is fully saturated with station by distance and station by month fixed effects. Consequently, the α_{dj} will be a variance weighted-average of each station’s individual difference-

¹⁹A potential concern is that transit stations are announced well in advance, potentially inducing households to relocate to the area prior to opening. Such sorting could affect Uber usage before service begins or alter the composition of travelers near the station. We believe this channel is unlikely to be important. Residential relocation involves substantial costs and long planning horizons, making it implausible that households would systematically adjust ride-hailing behavior for extended periods before a station opens – especially given that many openings are delayed. More directly, our event-study estimates show no pre-trends in Uber usage in the months leading up to openings, even when we extend the sample to 18 months prior. The results are also robust to restricting the sample to stations that are neither terminal nor centrally located, where residential sorting is less likely.

in-differences estimate at event time j and will not suffer from contaminated control units.²⁰

The specification above uses Uber trips farther from the station as a control group for station-adjacent trips affected by its opening. For the nearby-traveler analyses (Quadrants II and III), we instead compare riders according to the location of their POI relative to the new station. We bin riders into three groups: 0–500 m, 500–800 m, and 800–1200 m from the station, and use riders with a POI 800–1200 m away as the control group for riders in the 0–500 m and 500–800 m groups. We pool bins because of the smaller amount of trips in this sample. The specification is the same as above, except observations y_{dit} refer to the number of trips for a rider with a POI in distance group d from station i during month t . Our identifying assumption is now that riders who live or work 800–1200 m away from the station would have behaved similarly to riders who live or work closer to the station in the absence of the station opening.

While our main outcome of interest is Uber trips, we also consider the effects on pickups and drop-offs separately. We also study trips that originate or end at certain times of day - for example, weekday morning, weekday afternoon, or the weekend. We also consider other margins of Uber usage, like total distance traveled and total fares. We use average trip distance as an outcome to consider if trips get longer or shorter after the station opening. These alternative outcomes are considered in Section 5.1.

5 The Effect of a New Station on Uber Ridership

This section presents our core empirical results. We organize the analysis around the three empirical settings introduced in Table 1. We first examine station-adjacent Uber trips among all travelers (Quadrant I), which are particularly informative about first- and last-mile behavior. We then restrict attention to travelers with a fixed POI near the station, examining both their station-adjacent trips (Quadrant II), which shed light on substitution, and their citywide and non-local Uber usage (Quadrant III), which speaks most directly to option value. These estimates also provide the moments used to discipline the model in Section 6.

5.1 All trips near station, all travelers (Quadrant I)

We begin by estimating how overall Uber trips respond to the opening of a new transit station, focusing on variation in distance from the station. The top panel of Figure 4 plots the estimated effect on trips within 0–100 meters of a station, relative to trips within 1100–1200 meters. These estimates correspond to the $\alpha_{0-100\text{ m},j}$ in Equation 4. We find a large and statistically significant increase in ridership of 42% within one month of the opening (0.35 log points). In the two months prior to opening, there were an average of 419 trips within 100 m of a station, so the estimated effect corresponds to roughly 176 additional trips. Six months after opening, trips had increased

²⁰An alternative way to say this is for each station, each distance is subject to the same, contemporaneous treatment timing - including our control group 1100–1200 m away (Baker et al., 2022). This is equivalent to a “stacked” staggered adoption specification (Cengiz et al., 2019).

by 0.45 log points, and we can rule out increases smaller than 0.33 log points at the 95% level. This effect is persistent and shows no signs of decay over the six-month window.

The smaller increase in month 0 likely reflects partial exposure: stations do not always open on the first day of the month. In addition, travelers may take time to adjust their travel behavior. We also observe a modest effect one month prior to opening, consistent with “soft” openings or early access at some stations. Overall, we find a swift and large rise in Uber trips within 100 m of a new transit station.

The immediacy of the ridership response also helps rule out alternative explanations. For example, it is unlikely that new restaurants or retail developments – whose impacts typically unfold gradually – could account for such a sharp jump. While businesses may attempt to coordinate openings with new stations, doing so precisely is difficult. The high frequency of our data allows us to isolate the effects of sudden infrastructure shocks from slower-moving neighborhood change. Consistent with our argument, empirical evidence from [Kovvuri and Sharma \(2026\)](#) shows limited firm-entry responses during the first six months following Delhi Metro station openings.

The figure also provides evidence in support of our identification strategy. There are no signs of pre-trends in Uber trips from six to three months before the opening, and the large increase in ridership only occurs once the station opens. Our estimates are precise: we can reject pre-treatment differences between the 0–100 m and 1100–1200 m groups greater than 5.8% at the 95% level.

To assess substitution responses, we next examine effects slightly farther from the station, where the model suggests substitution responses may dominate. The lower panel of Figure 4 presents the event study for trips 300–400 m from the station, where we do not expect last-mile trip effects to occur. Unlike the sharp increases observed in the immediate vicinity, we observe a moderate but statistically significant decline in Uber usage beginning shortly after the station opens. This pattern is consistent with substitution away from Uber toward improved public transit access, particularly for travelers within comfortable walking distance of the station.

Figure 5 summarizes the average effect during months 1–6 of a new transit station on Uber ridership changes with distance. The reported coefficients are those from the model in Equation 5. On average, a new station opening increases Uber trips by 48% or 0.39 log points (s.e. 0.045) within 100 m. Farther out, we find that a new transit station opening increases Uber ridership at 100–200 m by 0.05 log points (a 5% effect, significant at the 95% level). For distances 200–700 m, we find evidence that new rail openings *decrease* Uber trips (though insignificantly at 400–500 m and 600–700 m). We find significant effects on the order of a 2% decline in trips at these distances. Thus, while Uber users show strong evidence of complementarities at the station, we do find evidence consistent with substitution away from Uber at intermediate distances. This is a point we take up in more detail in the next section.

Beyond these distances there is no detectable effect. Moreover, the magnitude of the estimates converges toward zero. For example, we can rule out an effect of the station opening on Uber

trips at 1000–1100 m of greater than 2.1%. We take this as additional evidence in favor of our identification assumption: that 1100–1200 m is sufficiently far from the station as to be unaffected by its opening. Figure A4 plots the coefficients for the full dynamic difference-in-differences specification for each distance. Importantly, the plots consistently show no pre-trends in the months leading up to the station opening across each distance band.

Together, these findings indicate that new station openings trigger a discrete shift in Uber demand. Trips that begin or end at the station rise sharply, consistent with last-mile complementarities, while trips at intermediate distances fall modestly, consistent with substitution toward transit.

Stability of the control group 1100–1200 m. Our finding of no effect beyond 800 m, relative to the 1100–1200 m band, suggests that transit–Uber interactions at these distances are minimal. An alternative possibility is that effects are constant across 800–1200 m and thus net out in our design. Indeed, if riders are willing to walk long distances to access public transit, then even those far from a station may substitute away from Uber. However, in North American transportation planning, a 1/2-mile (approximately 800 m) radius is widely used as the default catchment area for rail transit stations (Guerra et al., 2012). Empirical studies show that ridership tends to decline sharply beyond this range (van Soest et al., 2020, Daniels and Mulley, 2013). Our control group lies well outside this threshold, but we formally assess its validity below.

To test this, we estimate an event study – essentially a pre–post comparison – for locations 1100–1200 m from the new station. Finding no change in Uber usage at this distance would support the identifying assumption underlying our main dynamic difference-in-differences strategy.

Our event study test only uses observations from the 1100–1200 m band, treating not-yet-treated observations as controls. We follow the estimator of de Chaisemartin and D’Haultfœuille (2026) and include city-by-month fixed effects to absorb secular trends. We use a linear model with the inverse hyperbolic sine of trips (in hundreds) as the outcome.²¹ Full specification details are provided in Appendix Section D.1.

Results are shown in Figure 6. Consistent with our identifying assumption, we find no significant changes in Uber trips before or after the opening. The results of this identifying assumption test reinforce the validity of the 1100–1200 m band as a comparison group. Figures A5 and A6 confirm similar results when using a basic two-way fixed effects model estimated by a PPML regression and an IHS-transformed linear model, respectively.

Effects across other ridership margins. Our main outcome thus far has been the total number of Uber trips that either originated or terminated near the station. We now explore whether our

²¹To our knowledge, a PPML-style estimator along the lines of de Chaisemartin and D’Haultfœuille (2026) has not been developed. Our compromise is to use linear regression with a transformed outcome so that we can interpret treatment effects as log changes. We normalize by 100 so that estimated treatment effects near to the station are comparable to those from our PPML specifications. However, estimated effects using inverse hyperbolic sine-transformed outcomes are not scale-invariant (Chen and Roth, 2024).

results hold across alternative ride metrics. Figure A7 shows that the effects are nearly identical when estimating impacts on pickups and drop-offs separately. We also find consistent results when examining total fares paid and kilometers traveled, as shown in Figure A8.

Given our broad sample of 55 cities across four continents, we next test whether the average effects mask regional heterogeneity. We estimate city-specific treatment effects within 0–100 m and 300–400 m of a station and plot the results in funnel graphs (Figure 7). Estimates are centered around the aggregate treatment effect, with most cities – especially those with larger Uber markets – falling near the aggregate estimate. While a handful of smaller cities show outlier responses,²² the core pattern is remarkably consistent across the global sample. These findings reinforce the external validity of our main findings and highlight their applicability across the variety of contexts in our sample.

Last-mile behavior. If the increase in Uber ridership near stations reflects last-mile use, we would expect Uber trips in those areas to be shorter on average. To test this, we compute the average kilometers per trip across all station-month observations (excluding those with zero trips). Estimates from our main specification are shown in Figure 8.

Consistent with last-mile behavior, we find that average trip length declines sharply near the station. Within 100 meters, trips are 0.17 log points (16%) shorter, corresponding to a reduction of 1.8 kilometers. At 100–200 m, trips are 5% shorter. In contrast, trip length increases slightly at intermediate distances – for example, by 0.02 log points (2%) at 400–500 m. This pattern is consistent with substitution away from previous last-mile trips: if travelers previously took Uber to access a farther transit station, but now walk to the new one nearby, the remaining trips at those distances will appear longer on average.

We next examine whether these patterns are stronger for stations where last-mile behavior is likely more important. Terminal stations – those at the end of a transit line – are natural candidates. Figure A9 reports treatment effects by distance band for terminal and non-terminal stations. Though noisier, estimates for terminal stations are substantially larger. We find an 82% increase in Uber ridership within 100 m (0.60 log points), more than double the 38% increase (0.32 log points) observed near non-terminal stations. At 100–200 m, terminal stations induce an additional 0.30 log point increase in trips relative to non-terminal ones.

Stations far from the city core are also locations where we might expect last-mile behavior to be strongest. We therefore examine whether ridership gains are larger for more peripheral stations by interacting each of our distance-specific effects with measures of station peripherality. Figure A10 shows that effects within 100 m are significantly larger for more peripheral stations, both when measured by travel time and by kilometers from the central business district (significant at the 5% and 10% levels, respectively).

²²Opposite-signed effects are concentrated in very few cities, including Panama City and Baku.

5.2 Trips near station by travelers near station (Quadrant II)

When analyzing all users, Uber usage increases sharply within 100 meters of a new transit station and declines modestly at intermediate distances. We now turn to a subsample analysis of the subset of riders with a pre-existing connection to a specific location near the station.

This section focuses on a subset of Uber trips taken by users whose regular activity (either home, work, or habitual travel) is near the station – what we refer to as the POI sample. These riders have a fixed relationship to the treated area, allowing us to hold user composition constant and observe how their Uber usage responds to improved transit access. Because these riders start or end their trips near the station, they are unlikely to use Uber to get to or from it, so this sample largely removes last-mile complementarity; the response that remains is due to substitution and option value.

We begin by examining how Uber usage changes across distance bands for users with a fixed point of interest (POI) near the station. For each 100 meter distance band, we estimate a separate difference-in-differences regression comparing trips in that band against a control group of trips farther from the station. Riders are grouped into three POI distance bins based on where their pre-opening point of interest falls relative to the station: 0–500 m, 500–800 m, and 800–1200 m. For the POI 0–500 m and POI 500–800 m groups, the control group is trips in the 800–1200 m distance band within the POI 800–1200 m group. To assess if the effect decays toward zero with distance, we also consider the response of Uber trips within the POI 800–1200 m group by using the 1100–1200 m band within the same group as the control. We plot the average treatment effect – one to six months after opening – for each treated distance band, with separate markers and colors for each POI group. We exclude the 0–100 m distance band since the resolution of our POI hexes, which are approximately 174 meters in width, means riders cannot be precisely assigned to this innermost band.²³

These estimates are presented in Figure 9. Following the opening, we find a significant and sustained decline in Uber usage near the station. Uber trips fall by 7.1% at 100–200 m, 8.2% at 200–300 m, 5.9% at 300–400 m, and 4.0% at 400–500 m. The effects taper off beyond this range, lessening to a 1.2% and 1.4% decline for 500–600 m and 600–700 m from the station, respectively. By 700–800 m, the effect is indistinguishable from zero. The analysis using just the 800–1200 m POI group finds precise null effects for distances 800–1100 m as well. Because this exercise mostly removes last-mile trips, the reductions in Uber usage indicate that substitution toward transit outweighs the offsetting option-value response. This contrasts sharply with the increase in station-adjacent trips in the all-traveler sample, which primarily reflects first- and last-mile trips to and from the station.

Event-time-specific estimates for the average effects in Figure 9 are presented in Figure A11. Across distances, we find balanced trends in Uber usage prior to the station opening.

²³Even with these adjustments, 150 stations (depending on the specification) are omitted from the regression due to sparsity of the data and our saturated fixed effects.

Heterogeneity by time/day of week. Table 3 examines whether the effect varies by time of day and day of week. Each column reports the average treatment effect over months 1–6 following station opening for a different time block, comparing near-station trips (100–500 m) for the POI 0–500 m group against trips in the 800–1200 m band for the POI 800–1200 m group. The effect is negative and highly significant across all time blocks, indicating that the substitution away from Uber near the station is not concentrated in any particular part of the week. The largest effects are found during weekday morning commutes (-7.6%, column 2) and weekday afternoons (-7.3%, column 3), consistent with train stations being most effective substitutes for Uber during peak commuting hours when transit frequency is highest. Both weekday effects are statistically distinguishable from Friday and Saturday nights (-4.9%, column 4) and the rest of the week (-5.6%, column 6), but not from weekend days (-7.0%, column 5). The relatively large weekend day effect likely reflects the role of trains as substitutes for leisure trips during daytime hours, when transit service remains frequent. By contrast, the smaller effects on Friday and Saturday nights are consistent with Uber trips for nighttime activity when late-night transit options are more limited. Overall, the pattern suggests that substitution is strongest when transit service is most reliable and frequent, though it remains significant and economically meaningful across all time periods.

5.3 All trips by travelers near station (Quadrant III)

We next ask whether improved transit access changes Uber usage elsewhere in the city for these same travelers. Using the same POI sample from the previous section, we examine the total number of Uber trips taken by each user across the city – what we refer to as *citywide* Uber ridership. To better understand the nature of these changes, we disaggregate trips into those that are *local* (originating or ending within the rider’s own POI distance band) and *non-local* (trips that occur farther than the outermost edge of the POI distance band).²⁴ We assign each user to one of three groups based on the distance between their primary POI and the new station: 0–500 meters, 500–800 meters, and 800–1200 meters. Following the strategy used in the previous section, we treat users whose POI is 800–1200 meters from the station as the control group.

Table 4 presents estimates of the average effect over months one through six after the station opens. The first three columns report results for users with POIs within 0–500 m of the station. Column (1) shows that citywide Uber usage declines significantly by 2.1% following the station opening. Disaggregating this effect, Column (2) shows a 6.1% decline in local trips – consistent with the substitution effects documented in the previous section – while Column (3) reveals a partially offsetting 6.2% increase in non-local trips, significant at the 5% level. The next three columns focus on users with POIs located 500–800 m from the station. Here, citywide Uber usage is essentially unchanged (Column (4), -0.1%, insignificant), with a modest 1.3% decline in local trips (Column (5), significant at the 5% level) and a 3.4% increase in non-local trips (Column (6),

²⁴We define local trips as the sum of pickups and drop-offs within the rider’s POI distance band. We define non-local trips as trips that occur farther than the outermost edge of the POI distance band — equivalently, citywide pickups minus all pickups and drop-offs occurring anywhere from the station out to the POI band’s outer edge. Citywide trips are measured as citywide pickups, one per trip, ensuring no double counting.

significant at the 5% level).

Figure A12 presents the corresponding event-study plots. For citywide and local trips, trends are balanced across all pre-opening periods in both POI samples, supporting the validity of the research design. For non-local trips, the event studies are also generally balanced. The only exception is an isolated negative coefficient for the 0–500 m POI group six months before station opening, which is significant at the 5% level. We assess the sensitivity of our results to this pre-trend by using a trimmed five-month symmetric event window. Our estimates are slightly attenuated but remain significant at the 5% level. To be conservative, we use these estimates in the model estimation in Section 6.

These findings provide causal evidence of the option-value complementarity predicted by the model. The increase in non-local trips mirrors the dynamics described in Proposition 1: some travelers may choose not to drive in the morning in order to preserve improved transit options for the return leg, causing them to rely more heavily on Uber for other trips that arise. At the same time, the decline in local Uber trips is strongest for travelers within 0–500 m of the station, consistent with stronger direct substitution where transit access costs are lowest. Columns (1) and (4) are the net effect of these offsetting channels for travelers with a POI near the station;²⁵ Section 6 uses the estimated model to separate the channels and to compute their net effect for all travelers within 2.5 km of the station.

5.4 Robustness

We conduct a series of robustness checks for our empirical results, focusing primarily on Quadrant I, where statistical power is greatest, the last-mile response is most visible, and the resulting moments are central to the model-based analysis in Section 6.

Placebo tests. We assign each station a fake opening date randomly drawn from the set of dates with sufficient pre- and post-opening data and re-estimate our main specification 100 times. Figure A13 compares the distribution of placebo coefficients with the true estimates (blue circles), showing that the true near-station increases and intermediate-distance declines lie well outside the distribution generated by randomly assigned opening dates.

Uber fare structure. If Uber were strategically repricing near new stations – for example, by lowering fares to compete with transit or through surge pricing induced by the demand shock near the station – this could confound our interpretation of the ridership patterns as demand-side responses. To test this possibility, we estimate our main specification using fare per trip as the

²⁵A trip to the new station may replace an Uber trip the traveler took before, whether a direct trip or an access trip to a different station. The citywide estimates include every trip these travelers take, so any such replaced trip is netted against the trip that replaced it. Travelers with no POI within 1.2 km are outside this sample. For them the station is a waypoint rather than an end of the journey, so the trip they replaced enters the Quadrant I estimates only if one of its ends lies within 1.2 km of the station. The model accounts for these trips because in the model last-mile trips originate throughout the catchment area, not only within 1.2 km of the station.

outcome and interact the treatment effect with kilometers per trip to separately identify changes in the fixed and marginal components of fares. These estimates are depicted in Figure A14. We find no meaningful change in either component, providing no evidence that pricing adjustments drive our results.

Additional checks. Results are also robust to a battery of further specification and sample changes detailed in the Appendix. Alternative sets of fixed effects (Tables A2 and A3) and event-study estimates using not-yet-treated periods as controls (Figures A15 and A16) leave both the magnitude of the effects and the absence of pre-trends unchanged. Using adjacent 100-meter bands as controls instead of the 1100–1200 m band (Figure A17) yields estimates that closely track our main results in shape, magnitude, and significance. Replacing concentric rings with fixed-area hexagonal bins (Figure A18) produces muted but qualitatively similar patterns, as expected given spatial averaging within bins. Extending the event window to 18 months (Figure A19) reveals no pre-trends and shows that the near-station increase persists and grows over time.

POI-sample results. The decline in citywide Uber trips among the POI 0–500 m group is robust to alternative fixed-effect specifications (Table A4). Placebo tests for the Quadrant III estimates (Figure A20) show that, for the 0–500 m POI sample, the actual citywide, local, and non-local estimates all lie outside the placebo distribution; for the 500–800 m sample, the non-local estimate does as well. City-specific treatment effects (Figure A21) show our estimates for larger Uber markets cluster near the pooled effects. This indicates that the citywide, local, and non-local results are not driven by a small number of outlier cities.

6 Model Estimation and Counterfactuals

While our reduced-form findings are consistent with Uber’s role in substituting for, complementing via last-mile access, and providing option value to public transit, they do not isolate the relative contributions of specific mechanisms, nor can they evaluate how transit behavior might change under alternative policy environments – such as the effect of Uber exiting the market or subsidizing last-mile trips. To that end, we estimate the parameters of the model introduced in Section 2 to match key reduced-form moments from Section 5. Our goal is to organize the reduced-form evidence, quantify mechanisms, aggregate effects, and conduct counterfactuals which cannot be done using the findings in Section 5 alone.

6.1 Approach and moments matched

We estimate the model by matching its predictions to the causal estimates from Section 5 using the method of simulated moments. The goal is not to recover a fully structural model of travel demand, but to select parameter values consistent with our empirical patterns and use them to quantify mechanisms and evaluate counterfactuals.

As outlined in Section 2, the parameter vector θ includes baseline mode preferences for primary and intermediate trips, persistent preference heterogeneity, Uber access costs (with walking fixed costs normalized to zero). To this set, we include one additional parameter capturing how the station opening shifts transit utility for primary trips: $\Delta_{\text{transit}}^{\alpha}$. We also estimate two spatial parameters: an exponential kernel rate λ governing the spatial dispersion of first- and last-mile trips around the station, and a population density gradient κ capturing the decline in potential last-mile travelers with distance. In total, θ contains 13 parameters. Details of the simulation and spatial assignment of first- and last-mile trips are provided in Appendix E.

We match 24 reduced-form moments drawn from the three empirical quadrants, plus two calibration targets, to discipline these 13 parameters. We estimate θ by minimizing the weighted distance between data and model-implied moments, with an additional penalty for deviations from the calibration targets; the objective function and standard error formula are reported in Appendix E. The moments and the parameters they primarily discipline are as follows.

Quadrant I (13 moments): station-adjacent Uber usage. These moments include the estimated change in Uber trips across distance bands near the station, along with changes in average trip length within 0–100 m and 100–200 m. They are primarily informative about first- and last-mile access costs ($f_{\text{Uber}}, c_{\text{Uber}}, c_{\text{walk}}$), the magnitude of the transit quality improvement ($\Delta_{\text{transit}}^{\alpha}$), and the spatial parameters (κ, λ). These moments, in combination with the spatial patterns elsewhere, also inform the distribution of preferences ($\sigma_{\text{transit}}, \sigma_{\text{Uber}}, \phi$).

Quadrant II (7 moments): station-adjacent Uber usage for POI users. These moments are the average treatment effects on station-adjacent Uber usage for POI travelers (0–500 m and 500–800 m from the station) across distance bins from 100–200 m through 700–800 m. By conditioning on a stable set of travelers, these moments largely remove the last-mile behavior present in the all-user sample, leaving substitution and option value channels.

Quadrant III (4 moments): citywide and non-local Uber usage for POI users. These moments are the estimated changes in citywide and non-local Uber usage for the 0–500 m and 500–800 m POI samples. They relate to the option value mechanism and primarily discipline the intermediate-trip utilities β_{Uber} and β_{transit} .

Car and transit shares. We observe only Uber trips, yet the estimated model also reports changes in car and transit use. Two features of the model make this inference possible. First, distance to the station affects only the cost of reaching transit: driving and taking Uber directly cost the same wherever the traveler lives, but reaching transit costs more the farther away she lives. When the station opening reduces Uber trips by more at 300 m than at 700 m, the model attributes the difference to transit’s larger share at 300 m. The spatial pattern of Uber’s response therefore traces out how transit’s share declines with distance. Second, each of the three channels moves

travelers between a different pair of modes: substitution between Uber and transit, last-mile access between car and transit-by-Uber, and option value between car and Uber. Within each preference type, the size of each flow is proportional to the product of the two shares involved. Because the three quadrants isolate these channels, and we observe each across many distance bands, the reduced-form estimates identify how mode shares vary across space. The moments do not, however, pin down the level of the shares. Because the moments are differences in ridership, a large transit improvement acting on small shares produces moments similar to those of a small improvement acting on large shares.

Calibration targets. We therefore supplement the difference moments with two additional targets. The first anchors the level of the mode shares: we target a transit mode share of 40% within 800 m of stations after opening, the share observed near stations in dense, transit-oriented cities.²⁶ Together with the spatial and channel-specific variation in the moments, this target separates the level of the mode shares from the size of the transit improvement, and hence pins down the base utilities α_{transit} and α_{Uber} . The second target is the pre-opening ratio of Uber trips per square kilometer at 1000–1100 m to that at 1100–1200 m, which is 1.07 in the data. Before the station opens, distance to it does not affect travel costs, so this ratio reflects only differences in population density. It therefore pins down the density gradient κ separately from the marginal access cost c_{Uber} , two parameters that otherwise produce similar spatial patterns in first- and last-mile trips.

6.2 Estimation results

Before reporting parameter estimates, we describe how we convert utility units to willingness-to-pay (WTP). Let ρ denote the marginal utility of income, which governs the mapping between utility units and dollar-valued WTP. Our identification exploits transit expansions rather than price variation, so we cannot separately identify mode utilities from ρ without additional structure. Moreover, we found that Uber’s pricing structure – fixed costs and cost per kilometer – does not change in response to station openings (Figure A14). We therefore calibrate ρ using an external Uber demand elasticity of $\varepsilon^D = -0.6$ from [Cohen et al. \(2016\)](#) and [Castillo \(2025\)](#). Appendix Proposition 2 derives the mapping from the elasticity target to ρ in our mode-choice model. This yields $\hat{\rho} = 0.47$, so that one utility unit corresponds to \$2.13 in WTP.

Improvement in transit quality. Parameter estimates and standard errors are reported in Table 5. The station opening corresponds to an improvement in outbound/return utility from public transit of $\Delta_{\text{transit}}^{\alpha} = 0.60$ (s.e. 0.12), corresponding to a \$1.28 improvement in WTP per primary trip – roughly 16% of the average Uber fare in our sample. Public transit is a more attractive mode

²⁶The [Center for Transit-Oriented Development \(2024\)](#) data show transit mode shares within 1/2 mile of stations range from 15% in Los Angeles to 50% in New York. We target a value near the New York benchmark because its density and transit quality more closely resemble the Asian and European cities in our sample.

than car for those that live near the station by \$3.38, with 38% of that value coming from the new station itself. At the station itself ($d = 0$), the opening increases the probability of choosing transit from 60% to 67%, with gains coming both from driving and direct Uber trips. The increase in transit use peaks at 8.7 pp at 450 m from the station, while the decline in car use peaks at 7.6 pp at 550 m from the station.

Access costs and last-mile behavior. The fixed cost of using Uber for station access is estimated at $f_{\text{Uber}} = 5.47$ (s.e. 0.15), corresponding to approximately \$11.65 in WTP. Marginal access costs imply each additional 100 meters imposes disutility equivalent to \$0.95 by walking and \$0.26 by Uber. While last-mile trips remain rare at any given location – accounting for 1.4% of total trips at peak – they aggregate across the catchment area to produce the substantial 48% increase in station-adjacent Uber ridership documented in Section 5.1. The estimated spatial dispersion parameter λ implies an average measurement error in last-mile trip geo-locations of 45 m. The estimated population density gradient $\kappa = 0.55$ (s.e. 0.16) implies that the pool of potential last-mile travelers nearly halves every kilometer from the station, reflecting the declining share of residents for whom the new station is the closest transit option.²⁷

Mode preferences and heterogeneity. The estimated baseline utilities imply that, adjacent to the station, transit is preferred to car use, which in turn is preferred to direct Uber, for both commute and intermediate trips. The estimated standard deviations of persistent preference shocks are $\sigma_{\text{Uber}} = 2.83$ (s.e. 0.18) and $\sigma_{\text{transit}} = 2.18$ (s.e. 0.55), reflecting substantial heterogeneity in preferences over both modes relative to private car use. This is consistent with Uber serving a mix of travelers for whom it is highly convenient and those for whom it is rarely used, and similarly with public transit serving travelers ranging from those with strong preferences for it to those who find it costly or inconvenient. The estimated correlation between Uber and transit preferences is strongly negative ($\phi = -0.74$, s.e. 0.06), indicating that travelers who prefer ride-hailing over driving tend to dislike transit, and vice versa. This suggests that Uber and transit draw from different segments of the non-driving population, which indicates that policies discouraging ride-hailing will tend to shift travel toward private vehicles rather than transit.

Model fit. Table A5 compares model-implied moments to their empirical counterparts. Importantly, the model predicts negligible changes in Uber ridership within our control groups: a 0.002 log point change at 1100–1200 m for Quadrant I, and 0.012 for local POI trips at 800–1200 m for Quadrant II. These predictions support our identifying assumption that treatment effects are negligible sufficiently far from the station. For Quadrant III, the model predicts log point increases of 0.023 and 0.013 for non-local and citywide trips among the 800–1200 m POI sample –

²⁷The estimated costs imply travelers are willing to walk up to 1688 m before Uber becomes the preferred access mode. Beyond 800 m, transit mode share falls below 25%, consistent with standard 1/2-mile catchment areas (Guerra et al., 2012, Daniels and Mulley, 2013). By 1200 m the model predicts less than 1% change in Uber ridership, providing theoretical justification for our control group choice.

suggesting if anything that our reduced-form estimates of the non-local trip increase are biased downward. Turning to the matched moments, the model performs well overall, with an RMSE of 0.015 log points. It slightly understates the decline in trips for the POI sample near the station, and overstates the overall decline in citywide Uber trips. On the latter point, since the citywide effect combines the model’s local decline and non-local increase in Uber usage, this over-prediction stems from the model placing more weight on local trips than the data support. This is unsurprising: we cannot distinguish non-local trips that are genuinely chained to a commute from those that occur for unrelated reasons elsewhere in the city.

6.3 Decomposition of mechanisms

Proposition 1 provides a framework for decomposing the change in Uber ridership induced by an improvement in transit quality. We apply it to the opening of a new station, aggregating across traveler and trip types within a 2.5 km potential catchment area and weighting each distance band by population.²⁸ We focus on this population because these travelers are most likely to be affected by the new station. The 2.5 km threshold is also consistent with the radius used in Berlin and Munich pilots for last-mile Uber subsidies. Figure 10 shows how the transit improvement propagates through three mode-choice margins: Uber versus transit, transit versus car, and car versus Uber. These margins generate the three mechanisms in Proposition 1: direct substitution, last-mile complementarity, and option value.

Direct substitution is the largest individual channel, reducing aggregate Uber ridership by 1.55% as travelers switch from Uber to improved transit. Last-mile complementarity offsets 0.92 percentage points of this decline, as some of the additional transit trips are accessed or egressed using Uber. The option-value channel adds another 1.18% to Uber usage, as improved transit makes it more attractive to avoid driving and preserve access to Uber and transit later in the day. Thus, although substitution dominates either complementary channel individually, last-mile and option value together more than offset it.

The underlying mode flows also show how the channels interact. Overall, transit usage rises by 12.16%, reflecting both direct substitution toward transit from Uber and car usage, as well as transit use generated by option value. Most of these additional transit trips are accessed by walking, but some use Uber, generating the 0.92% last-mile effect. Importantly, more than half of this last-mile effect comes from transit trips that are themselves induced by option value.

Car usage also falls by 4.69% as travelers switch toward transit. At the same time, some travelers switch from car to Uber because Uber preserves the option of using improved transit later in the day, generating the 1.18% option-value increase in Uber usage. Combining the three channels, a station opening increases aggregate Uber ridership by 0.54% in a 2.5 km radius around

²⁸To obtain an exact decomposition of Uber usage, we integrate each term in the derivative in Proposition 1 over $\Delta_{\text{transit}}^{\alpha}$ from 0 to $\hat{\Delta}_{\text{transit}}^{\alpha} = 0.60$. Tracking the mode switches underlying each channel also gives the associated changes in car and transit use. The decomposition is exact for all modes. Appendix C.3 derives the flows underlying each channel, and Appendix E describes how we compute them from the estimated model.

a new station. Improved public transit generates substantial substitution away from Uber, but the two complementary channels together dominate on net.

Our model handles the possibility that trips may have been diverted by the new station. Consider a traveler 2 km from the station who replaces a direct Uber trip with an Uber access trip. This trip appears in the direct substitution and the last-mile channels. Of the aggregated 0.92% last-mile effect, 0.11 percentage points are trips that were already Uber trips before the station opened; the remaining 0.81 percentage points are access trips by travelers who previously drove.

6.4 Counterfactuals

Policymakers often seek to increase public transit use to reduce driving externalities.²⁹ The effectiveness of these policies depends on how transit interacts with other modes, including ride-hailing and personal vehicles. We use the estimated model to evaluate counterfactual policies that target different margins of travel behavior. Because we do not explicitly value reductions in congestion and pollution, the welfare and fiscal measures described below should be interpreted as the direct costs required to induce higher transit ridership.

Ride-hailing bans. Section 6.3 showed that a new train station increases Uber ridership in our setting. We probe this complementarity further by examining how transit ridership changes when ride-hailing is removed.³⁰

Table 6 reports model-predicted changes in mode shares, welfare, and fiscal costs across counterfactual scenarios. Panel A considers a ride-hailing ban directly. Removing Uber causes former users to substitute primarily toward cars rather than transit, and aggregate transit ridership falls by 9.87%. This drop is even larger (10.93%) without a train station. Taken together, these counterfactuals show that last-mile access and Uber’s option value are critical: removing Uber reduces transit use in this setting, consistent with Uber and transit being complements on net.

Subsidizing last-mile trips. Panel B considers subsidies to first/last-mile access, a commonly proposed tool for expanding station catchment areas and one that transit agencies have adopted (Hall, 2026). To compare similarly sized policies, we hold the government’s outlay fixed: we compare a flat subsidy of \$2.50 per last-mile trip to a per-km subsidy of \$1.68/km (65% of marginal last-mile access costs), chosen so that both policies carry the same net fiscal cost of \$0.14 per traveler.

At the same cost, the per-km subsidy is more effective at increasing transit ridership. The per-km subsidy increases overall transit use by 6.18% as compared to 5.69% for the flat subsidy.

²⁹These policies are typically motivated by the goal of reducing automobile externalities, including congestion (Currie and Walker, 2011, Yang et al., 2020), emissions (Anderson, 2020, Gendron-Carrier et al., 2022), noise (Moretti and Wheeler, 2025), and safety (Ashenfelter and Greenstone, 2004). For a survey, see Parry et al. (2007); Kashner and Ross (2025) provides recent evidence quantifying these externalities.

³⁰We treat Uber as representative of the broader ride-hailing market. In settings with close substitutes, the effects of Uber exit would be attenuated; in our sample, however, Uber was typically the dominant ride-hailing platform.

The difference reflects the margins each policy targets: fixed subsidies partly shift existing transit users from walking to Uber access, whereas marginal-cost subsidies induce new transit use, especially among travelers farther from stations where access costs are high.

Taxes. Panel C compares ride-hailing and congestion taxes, both of which are commonly used to discourage driving and increase transit use. Because Uber also provides last-mile access, however, taxing ride-hailing can work against the objective of increasing transit use. A \$2.00 ride-hailing tax reduces Uber usage for primary trips by 10.89 pp, but most marginal users switch to cars rather than transit, consistent with the estimated negative correlation in Uber and transit preferences. Transit-by-walk rises while transit-by-Uber falls. Intermediate-trip transit use also falls, as travelers who switch toward car commuting become less likely to use transit for subsequent trips. Overall, the ride-hailing tax reduces aggregate transit ridership by 1.60%. The negative sign reflects the complementarities. Under pure substitution, a ride-hailing tax could only shift riders from Uber toward transit or cars and so could not lower transit ridership; the decline therefore reflects the last-mile and option-value channels, which more than offset the switching into transit.

To model congestion pricing, we interpret the traveler as commuting from the urban periphery to a job in the dense CBD and impose a cordon charge on outbound and return trips by car or direct Uber. We assume last-mile access trips do not enter the cordon and that intermediate trips occur within it, so neither is taxed. To compare the two policies, we choose the congestion charge, \$0.56 per cordon crossing, to generate the same fiscal impact as the \$2.00 ride-hailing tax. At this charge, aggregate transit ridership increases by 5.50%. Congestion pricing is more effective in this setting because it directly raises the cost of the principal fallback option to transit – driving – without simultaneously raising the cost of last-mile access.

Figure A22 shows that these differences extend beyond the particular tax levels considered in Panel C. Increasing the congestion charge monotonically increases transit ridership, and does so at an increasing rate as the non-transit fallback option becomes progressively less attractive; conversely, subsidizing car and direct ride-hailing trips into the cordon reduces transit use. The response to a uniform ride-hailing tax or subsidy is very different: aside from a negligible increase ($<0.10\%$) for subsidies of \$0.10-\$0.90 per trip, moving the price of ride-hailing in either direction reduces transit ridership. Taxes discourage last-mile access and induce substitution toward cars, while large enough subsidies make Uber more attractive for both commute and intermediate trips, causing direct substitution from transit to Uber that eventually dominates.

Taken together, Panels B and C highlight the importance of targeting policy to the margin on which ride-hailing complements transit. The first/last-mile subsidies increase aggregate transit ridership by more than 5.50% at a net fiscal cost of only \$0.14 per traveler. By contrast, a uniform tax on ride-hailing reduces transit ridership, while the congestion charge increases transit ridership by the same amount despite collecting approximately \$0.99 per traveler.

Fiscally neutral policy mix. Finally, we solve for a fiscally neutral policy that increases aggregate transit use by 5 pp while minimizing welfare losses, financing last-mile subsidies with ride-hailing and congestion taxes. Panel D reports the resulting policy mixes.

When only ride-hailing taxes are available as a revenue source, the model selects a marginal last-mile subsidy of \$1.93 per km, a fixed subsidy of \$3.35, and a \$1.28 ride-hailing tax. Achieving the transit-use target in this way is costly, generating a 8.1% welfare loss (\$0.52 per traveler). Allowing congestion pricing substantially improves the policy mix: the model selects a \$0.35 congestion tax, no ride-hailing tax, and smaller marginal and fixed last-mile subsidies of \$1.60 per km and \$2.02, respectively. The welfare cost falls to 5.6% (\$0.36 per traveler). Thus, congestion pricing dominates ride-hailing taxation as a revenue instrument because it discourages driving without simultaneously raising the cost of last-mile transit access.

Table A6 shows that the resulting welfare losses are largest for travelers farthest from the station, who have the least scope to adjust behavior, and smallest near the station, where many travelers already use transit. When direct transit subsidies are also available, the policy mix shifts markedly: a \$1.00 transit subsidy replaces last-mile subsidies entirely, paired with a \$0.40 congestion tax and a negligible \$0.03 ride-hailing tax. The welfare cost falls further to 1.2% (\$0.08 per traveler). Direct transit subsidies are more effective because they increase the attractiveness of transit for both commute and intermediate trips, whereas last-mile subsidies operate primarily through commute access and become increasingly costly at longer distances.

6.5 Policy implications

Our results suggest that the policy question is not whether ride-hailing and public transit are complements or substitutes in general, but how and where they interact. The same transit improvement generates substitution away from Uber through improved direct service, and complementarity with Uber through last-mile access and option value; the net effect depends on the relative strength of these mechanisms across trip types and distance from the station. This helps reconcile the wide range of existing estimates of Uber's effect on transit ridership ([Hall et al., 2018](#), [Diao et al., 2021](#)), and it speaks directly to the debate that opened this paper: several countries, including Denmark, Hungary, and Bulgaria, have banned ride-hailing outright. Our ban counterfactual (Panel A) suggests this policy can backfire where last-mile access and option value are important – in our setting, removing Uber reduces rather than increases transit ridership. The same logic applies to transit agencies weighing restrictions on ride-hailing to protect farebox revenue, since these restrictions can have the opposite effect.

A first implication is for transit agencies planning station openings or service expansions. Changes in ride-hailing demand around new stations are not solely evidence of competition with transit. Increases in Uber use near stations reflect greater use of Uber as an access mode to transit, while improved transit quality reduces direct Uber use for riders who live or work nearby. Station-area planning should account for both margins simultaneously – in practice, this points to the importance of curb management, pickup/drop-off design, and enforcement near

stations. It also highlights that new transit access, intended to reduce local congestion, may have the unintended consequence of *increasing* it.

A second implication concerns the design of transit–ride-hailing partnerships and subsidies. Berlin and Munich piloted partnerships with Uber in 2021 offering a flat fee for trips beginning or ending at U-Bahn stations, later expanded to additional stations. Our counterfactuals show that subsidies of equal fiscal cost can have markedly different effects on transit ridership depending on delivery: subsidies that reduce the marginal cost of Uber access, akin to the Berlin/Munich flat fee, generate larger increases in transit use than equal-valued fixed-cost subsidies, because they target the last-mile margin directly rather than reallocating access from walking to Uber among existing transit users. At the same time, marginal-cost subsidies become fiscally expensive as they induce more distant travelers to switch, providing a rationale for distance caps on eligible trips – such as the 2.5 km limit in the Berlin/Munich pilot. When policymakers can also subsidize transit directly, however, last-mile subsidies are no longer the most efficient tool, since direct subsidies raise the attractiveness of transit for both commute and intermediate trips simultaneously. More broadly, our results complement [Agrawal and Zhao \(2023\)](#) by showing that whether ride-hailing complements public transit depends on the policy instruments that shape last-mile access.

A third implication concerns who bears the cost of transit-promoting policy. The distribution of welfare losses under our fiscally neutral policy mix (Table A6) is uneven across distance from the station: travelers farthest from the station bear the largest welfare costs, with the least scope to adjust behavior, while losses are smallest near the station, where many travelers already use transit. Pairing congestion pricing with last-mile subsidies mitigates losses for some travelers with limited walking access to transit, particularly at intermediate distances, although those farthest from the station continue to bear the largest welfare costs even under the fiscally neutral policy mix. To the extent that income varies systematically with distance from the station, these findings are relevant to broader equity concerns surrounding congestion pricing ([Cook et al., 2025](#)).

More generally, our framework implies that policies targeting ride-hailing or driving inevitably affect multiple behavioral margins at once: it is difficult to tax substitution away from transit without also changing the last-mile access and option-value channels that work in the opposite direction ([Tarduno, 2025](#), [Leccese, 2025](#)). Ride-hailing taxes may be particularly poorly targeted in peripheral catchment areas or in settings with unreliable transit, precisely where last-mile access and option value matter most; congestion pricing performs better because it targets cars directly, though the relative strength of these mechanisms – and thus the case for either instrument – should depend on local transit quality, density, and car ownership.

7 Conclusion

The question of whether ride-hailing complements or substitutes for public transit has generated substantial debate and a wide range of empirical estimates. Our results suggest that the disagreement is not only a matter of measurement or context: substitution and complementarity

operate simultaneously, on different margins, for different travelers, and at different distances from transit infrastructure. Aggregate estimates often conflate these forces, and the sign of the net effect depends on which margins are observed and which policies shape these interactions.

In this paper, we provide new causal evidence on how ride-hailing responds to improvements in public transit using Uber data around 650 rail station openings across 55 cities in 35 countries. Our identification strategy exploits the sharp timing of station openings – planned years in advance, typically before Uber existed – to compare changes in ride-hailing across concentric distance bands. We document three patterns: a 48% increase in Uber trips within 100 m of new stations, consistent with first- and last-mile complementarity; among travelers with a persistent connection to station areas, substitution away from Uber use near the station by 6%; and, among these same travelers, increased Uber use elsewhere in the city, consistent with option value.

We then develop and estimate a model of mode choice that formalizes these mechanisms and allows us to quantify their relative importance and evaluate counterfactuals. We find that all three channels are quantitatively important: although substitution dominates either complementary channel individually, last-mile and option value together more than offset it. More than half of the last-mile effect arises from transit trips induced by dynamic considerations, highlighting the interaction between the two sources of complementarity. On net, Uber and transit are complements in our setting: new stations increase Uber use and removing Uber reduces transit use.

The policy counterfactuals further show that instruments designed to increase transit use can have sharply different effects depending on which margins they target. Ride-hailing taxes reduce transit ridership by weakening last-mile access and option value, while congestion pricing is more effective because it targets cars directly without undermining these complementarities. More broadly, our results suggest that the key policy question is not whether ride-hailing competes with transit, but how the two modes interact across space and across margins of behavior.

Several limitations remain. Our estimates capture short-run behavioral responses within six months of station openings and may not reflect longer-run adjustments in residential location, car ownership, or travel patterns. Those longer-run changes could also include supply-side responses – such as changes in ride-hailing pricing or transit frequency – that could amplify or dampen the mechanisms we identify. And while our international sample provides unusually broad external validity, the estimated parameters reflect an average across contexts that vary substantially in transit quality, urban density, and car ownership.

These limitations aside, the same three mechanisms we document – substitution, access-mode complementarity, and the option value of flexible mode-switching – extend to other technologies that alter the cost of feeder trips or the flexibility of return journeys, from micromobility to autonomous vehicles. Understanding these interactions, rather than treating any new mode as a simple substitute or complement to transit, is essential for designing transportation infrastructure and policy that works alongside the modes travelers actually use.

Bibliography

- Agrawal, David R. and Weihua Zhao (2023) "Taxing Uber," *Journal of Public Economics*, 221, 104862.
- Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel Hickok, and Tobias Salz (2024) "Optimal Urban Transportation Policy: Evidence from Chicago," NBER Working Paper 32185, National Bureau of Economic Research.
- Anderson, Michael L. (2020) "As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality," *Journal of the European Economic Association*, 18 (4), 1886–1927.
- Anderson, Michael L. and Lucas W. Davis (2026) "Uber and Traffic Fatalities," *Review of Economics and Statistics*, 108 (2), 525–532.
- Angrist, Joshua D., Sydnee Caldwell, and Jonathan V. Hall (2021) "Uber versus Taxi: A Driver's Eye View," *American Economic Journal: Applied Economics*, 13 (3), 272–308.
- Ashenfelter, Orley and Michael Greenstone (2004) "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy*, 112 (S1), S226–S267.
- Baker, Andrew C., David F. Larcker, and Charles C.Y. Wang (2022) "How Much Should We Trust Staggered Difference-In-Differences Estimates?" *Journal of Financial Economics*, 144 (2), 370–395.
- Barreto, Yuri, Raul Silveira Neto, and Luís Carazza (2021) "Uber and Traffic Safety: Evidence from Brazilian Cities," *Journal of Urban Economics*, 123, 103347.
- Barrios, John Manuel, Yael V. Hochberg, and Hanyi Yi (2023) "The Cost of Convenience: Ridehailing and Traffic Fatalities," *Journal of Operations Management*, 69 (5), 823–855.
- Barwick, Panle Jia, Shanjun Li, Andrew Waxman, Jing Wu, and Tianli Xia (2024) "Efficiency and Equity Impacts of Urban Transportation Policies with Equilibrium Sorting," *American Economic Review*, 114 (10), 3161–3205.
- Basso, Leonardo J. and Hugo E. Silva (2014) "Efficiency and Substitutability of Transit Subsidies and Other Urban Transport Policies," *American Economic Journal: Economic Policy*, 6 (4), 1–33.
- Baum-Snow, Nathaniel (2007) "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics*, 122 (2), 775–805.
- Billings, Stephen B. (2011) "Estimating the Value of a New Transit Option," *Regional Science and Urban Economics*, 41 (6), 525–536.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024) "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 91 (6), 3253–3285.
- Bosker, Maarten, Mark Roberts, Sailesh Tiwari, and Putu S. Wibisana (2025) "Going the First/Last-Mile: Ride-Hailing Services and Mass Rapid Transport," *CEPR Discussion Paper 20344*.
- Cairncross, John, Jonathan D. Hall, and Craig Palsson (2025) "VancUber: The Long-Run Effect of Ride-Hailing on Public Transportation, Congestion, and Traffic Fatalities," *Canadian Journal of Economics*, 58 (3), 892–939.
- Callaway, Brantly and Pedro H. C. Sant'Anna (2021) "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225 (2), 200–230.
- Castillo, Juan Camilo (2025) "Who Benefits from Surge Pricing?" *Econometrica*, 93 (5), 1811–1854.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, 134 (3), 1405–1454.

- Center for Transit-Oriented Development (2024) "Transit-Oriented Development Database," <https://toddata.cnt.org/>, Accessed: January 23, 2025.
- de Chaisemartin, Clément and Xavier D'Haultfœuille (2026) "Difference-in-Differences Estimators of Intertemporal Treatment Effects," *Review of Economics and Statistics*, 108 (4), 863–880.
- Chen, Jiafeng and Jonathan Roth (2024) "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 139 (2), 891–936.
- Christensen, Peter, Gustavo Nino, and Adam Osman (2026) "The Demand for Mobility: Evidence From an Experiment With Uber Riders," *Journal of Political Economy* (forthcoming).
- Cohen, Peter, Robert W. Hahn, Jonathan V. Hall, Steven D. Levitt, and Robert D. Metcalfe (2016) "Using Big Data to Estimate Consumer Surplus: The Case of Uber," *NBER Working Paper* 22627.
- Cook, Cody, Aboudy Kreidieh, Shoshana Vasserman, Hunt Allcott, Neha Arora, Freek van Sambeek, Andrew Tomkins, and Eray Turkel (2025) "The Short-Run Effects of Congestion Pricing in New York City," *NBER Working Paper* 33584.
- Currie, Janet and Reed Walker (2011) "Traffic Congestion and Infant Health: Evidence from E-ZPass," *American Economic Journal: Applied Economics*, 3 (1), 65–90.
- Daniels, Rhonda and Corinne Mulley (2013) "Explaining Walking Distance to Public Transport: The Dominance of Public Transport Supply," *Journal of Transport and Land Use*, 6 (2), 5–20.
- Diao, Mi, Hui Kong, and Jinhua Zhao (2021) "Impacts of Transportation Network Companies on Urban Mobility," *Nature Sustainability*, 4 (6), 494–500.
- Gendron-Carrier, Nicolas, Marco Gonzalez-Navarro, Stefano Polloni, and Matthew A. Turner (2022) "Subways and Urban Air Pollution," *American Economic Journal: Applied Economics*, 14 (1), 164–196.
- Gonzalez-Navarro, Marco and Matthew A. Turner (2018) "Subways and Urban Growth: Evidence from Earth," *Journal of Urban Economics*, 108, 85–106.
- Goodman-Bacon, Andrew (2021) "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 225 (2), 254–277.
- Gorback, Caitlin (2026) "Ridesharing and the Redistribution of Economic Activity," *Journal of Urban Economics* (forthcoming).
- Greenwood, Brad N. and Sunil Wattal (2017) "Show Me the Way to Go Home: An Empirical Investigation of Ride-Sharing and Alcohol Related Motor Vehicle Fatalities," *MIS Quarterly*, 41 (1), 163–187.
- Guerra, Erick, Robert Cervero, and Daniel Tischler (2012) "Half-Mile Circle: Does It Best Represent Transit Station Catchments?" *Transportation Research Record*, 2276 (1), 101–109.
- Gupta, Arpit, Stijn Van Nieuwerburgh, and Constantine Kontokosta (2022) "Take the Q Train: Value Capture of Public Infrastructure Projects," *Journal of Urban Economics*, 129, 103422.
- Hall, Jonathan D. (2026) "Ride-hailing and Urban Transportation: Evidence and Policy," *Regional Science and Urban Economics*.
- Hall, Jonathan D., Craig Palsson, and Joseph Price (2018) "Is Uber a Substitute or Complement for Public Transit?" *Journal of Urban Economics*, 108 (1), 36–50.

- Hall, Jonathan V. and Alan B. Krueger (2018) "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," *ILR Review*, 71 (3), 705–732.
- Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm (2020) "The Making of the Modern Metropolis: Evidence from London," *Quarterly Journal of Economics*, 135 (4), 2059–2133.
- Kashner, Zane and Brad Ross (2025) "Measuring and Mitigating Traffic Externalities," Working Paper.
- Kovvuri, Akhila and Karmini Sharma (2026) "Moving Opportunity Closer: How Public Transit Transforms Firm Composition and Employment," Working Paper.
- Leccese, Mario (2025) "The Effects of Targeted Platform Taxes: Evidence from the Ride-Sharing and Taxi Industry," Working Paper.
- Moretti, Enrico and Harrison Wheeler (2025) "The Traffic Noise Externality: Costs, Incidence and Policy Implications," NBER Working Paper 34298.
- Parry, Ian W. H. and Kenneth A. Small (2009) "Should Urban Transit Subsidies Be Reduced?" *American Economic Review*, 99 (3), 700–724.
- Parry, Ian W. H., Margaret Walls, and Winston Harrington (2007) "Automobile Externalities and Policies," *Journal of Economic Literature*, 45 (2), 373–399.
- Schiavina, Marcello, Michele Melchiorri, and Martino Pesaresi (2023) "GHS-SMOD R2023A - GHS settlement layers, application of the Degree of Urbanisation methodology (stage I) to GHS-POP R2023A and GHS-BUILT-S R2023A, multitemporal (1975-2030)," doi: 10.2905/JRC.SPX1FDR.
- Silva, Joao M. C. Santos and Silvana Tenreiro (2006) "The Log of Gravity," *Review of Economics and Statistics*, 88 (4), 641–658.
- (2011) "Further Simulation Evidence on the Performance of the Poisson Pseudo-Maximum Likelihood Estimator," *Economics Letters*, 112 (2), 220–222.
- van Soest, Dennis, Miles R. Tight, and Christopher D. F. Rogers (2020) "Exploring the Distances People Walk to Access Public Transport," *Transport Reviews*, 40 (2), 160–182.
- Sun, Liyang and Sarah Abraham (2021) "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225 (2), 175–199.
- Tarduno, Matthew (2021) "The Congestion Costs of Uber and Lyft," *Journal of Urban Economics*, 122, 103318.
- (2025) "Elasticities and Tax Incidence in Urban Ridesharing Markets: Evidence from Chicago," *Journal of Urban Economics*, 150, 103809.
- Teltser, Keith, Conor Lennon, and Jacob Burgdorf (2021) "Do Ridesharing Services Increase Alcohol Consumption?" *Journal of Health Economics*, 77, 102451.
- Train, Kenneth E. (2009) *Discrete Choice Methods with Simulation*, 2nd edition, Cambridge: Cambridge University Press.
- Tsivanidis, Nick (2026) "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio," *American Economic Review*, 116 (2), 418–463.
- United Nations, Department of Economic and Social Affairs, Population Division (2018) "World Urbanization Prospects: The 2018 Revision, Online Edition. File 22: Annual Population of Urban Agglomerations with 300,000 Inhabitants or More in 2018, by Country, 1950-2035 (thousands)," <https://population.un.org/wup/>, Accessed: 2025-01-28.

Yang, J., A.-O. Purevjav, and S. Li (2020) "The Marginal Cost of Traffic Congestion and Road Pricing: Evidence from a Natural Experiment in Beijing," *American Economic Journal: Economic Policy*, 12 (1), 418–453.

Zárate, Román David (2024) "Spatial Misallocation, Informality and Transit Improvements: Evidence from Mexico City," *World Bank Policy Research Working Paper* 9990.

Tables

Table 2: Uber trips descriptive statistics

(a) Quadrant I - Near-station trips, all riders

	Summary Statistics				
	<i>Full: $\in [-6, 6]$</i>	<i>0-100 m</i>	<i>3-400 m</i>	<i>7-800 m</i>	<i>11-1200 m</i>
Uber trips	mean [median]	mean [med.]	mean [med.]	mean [med.]	mean [med.]
Pickups	1,663 [188]	419 [48]	1,530 [181]	1,862 [255]	2,110 [227]
Dropoffs	1,656 [188]	424 [47]	1,583 [188]	1,828 [247]	2,091 [227]
Kilometers	29,173 [3,986]	7,722 [951]	28,648 [3,882]	31,683 [5,275]	35,760 [5,003]
Km/trip	11.20 [10.05]	10.24 [9.10]	11.32 [10.12]	11.33 [10.22]	11.32 [10.12]
Fare/trip (USD)	7.84 [5.48]	6.98 [4.77]	7.92 [5.48]	7.96 [5.57]	7.89 [5.60]
Zero Trips Shr	0.16	0.25	0.14	0.14	0.15
Weekday AM Shr	0.09 [0.09]	0.08 [0.08]	0.09 [0.09]	0.09 [0.09]	0.09 [0.09]
Weekday PM Shr	0.11 [0.13]	0.11 [0.12]	0.12 [0.13]	0.12 [0.13]	0.11 [0.13]
Fri-Sat Night Shr	0.08 [0.07]	0.06 [0.06]	0.08 [0.07]	0.09 [0.07]	0.08 [0.08]
Weekend Day Shr	0.11 [0.11]	0.10 [0.11]	0.11 [0.11]	0.11 [0.12]	0.11 [0.12]
N	99,720	8,310	8,310	8,310	8,310

Notes: This table contains summary statistics on Uber trips in our sample. All data are from the study window of six months before and after the station opening. Each column contains means and medians (in brackets) for Uber trip outcomes. Column (1) is for our full sample through 1200 meters away from the station. Columns (2) through (5) disaggregate those means by four 100 meter distance bands. Rows (1) through (4) contain average counts for the number of pickups, drop-offs, kilometers travelled, and average trip length. Row (5) contains the share of month-distance band cells that had zero trips in them. This includes true zeros as well as those that had fewer than 5 trips, and so were censored. Rows (6) through (9) contain the share of trips during different parts of the week. Row (10) contains observation counts.

(b) Quadrants II & III - Local, non-local, and citywide trips, POI riders

	Summary Statistics			
	<i>Full</i>	<i>POI 0-500 m</i>	<i>POI 500-800 m</i>	<i>POI 800-1200 m</i>
Uber trips	mean [median]	mean [med.]	mean [med.]	mean [med.]
Citywide trips	5,441 [661]	3,229 [361]	4,566 [596]	8,528 [1,205]
Local trips	2,558 [248]	1,930 [174]	2,316 [246]	3,428 [358]
Non-local trips	2,404 [299]	1,224 [157]	1,814 [243]	4,174 [639]
Riders (citywide)	2,562 [344]	1,543 [191]	2,145 [315]	3,998 [603]
Km/trip (citywide)	10.96 [10.40]	10.88 [10.41]	10.97 [10.35]	11.03 [10.44]
Fare/trip (USD)	7.70 [5.61]	7.74 [5.58]	7.69 [5.61]	7.68 [5.63]
Zero Trips Shr	0.13	0.16	0.13	0.11
N	22,131	7377	7377	7377

Notes: This table contains summary statistics on citywide Uber trips for the POI sample. All data are from the study window of six months before and after the station opening. Each column contains means and medians (in brackets) for Uber trip outcomes. Column (1) is for the full POI sample. Columns (2) through (4) disaggregate by the distance between the rider's point of interest (POI) and the station: 0–500 m, 500–800 m, and 800–1200 m. Rows (1) through (3) contain average counts for citywide trips, local trips, and non-local trips, as described in the main text. Row (4) contains the average number of unique riders. Row (5) contains average kilometers per citywide trip. Row (6) contains average fare per citywide trip in USD. Row (7) contains the share of month-POI distance bin cells that had zero citywide trips. This includes true zeros as well as those that had fewer than 5 trips and so were censored. Row (8) contains observation counts.

Table 3: Quadrant II - Effects of rail station opening on station-adjacent Uber trips in POI sample by time/day of week

	Log Uber trips					
	All (1)	Weekday AM (2)	Weekday PM (3)	Fri/Sat night (4)	Weekend day (5)	Rest of week (6)
Avg. effect (1–6 mths)	-0.061*** (0.007)	-0.076*** (0.011)	-0.073*** (0.009)	-0.049*** (0.010)	-0.070*** (0.008)	-0.056*** (0.007)
Observations	182,439	129,044	132,634	126,570	127,095	170,621
Station Openings	521	466	464	455	454	508
Station x Dist FE	✓	✓	✓	✓	✓	✓
Station x Date FE	✓	✓	✓	✓	✓	✓
Event Time FE	✓	✓	✓	✓	✓	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table contains average treatment effect estimates for local Uber trips by time of day and day of week in the POI sample. Each column corresponds to a different time block. The outcome variable is near-station trips within 100–500 meters of the station for riders whose POI is within 500 meters of the station. The control group is trips in the 800–1200 meter distance band for riders whose POI is between 800 and 1200 meters from the station. The reported coefficient is the average effect over months 1–6 following station opening, relative to the pre-opening period. All specifications include station-by-POI, station-by-date, and event-time fixed effects. Standard errors are clustered at the station level.

Table 4: Quadrant III - Effects of rail station opening on citywide, local and non-local Uber trips in POI sample

	Log Uber trips					
	POI (0–500 m)			POI (500–800 m)		
	Citywide (1)	Local (2)	Non-local (3)	Citywide (4)	Local (5)	Non-local (6)
Avg. effect (1–6 mths)	-0.021*** (0.004)	-0.061*** (0.007)	0.062** (0.027)	-0.0014 (0.002)	-0.013** (0.005)	0.034** (0.016)
Observations	54,492	46,524	54,248	55,508	46,978	54,964
ATE 1–5 mths (%)	-2.18***	-6.03***	5.30**	-0.16	-1.15**	3.02**
Station Openings	529	494	528	540	502	536
Station x POI FE	✓	✓	✓	✓	✓	✓
Station x Date FE	✓	✓	✓	✓	✓	✓
Event Time FE	✓	✓	✓	✓	✓	✓

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table contains coefficient estimates for the average effect of a new train station opening on citywide, local, and non-local POI sample Uber trips during the first six months after station opening, relative to POI sample Uber trips 800–1200 meters away. Coefficients are relative to six to three months prior to the station opening. All specifications include station by distance, station by date and event-time fixed effects. All errors are clustered at the station level.

Table 5: Model estimates

Parameter		Estimate	Std. Error
α_{transit} :	Transit base utility for primary trip	0.99***	(0.14)
α_{Uber} :	Uber base utility for primary trip	-1.25***	(0.31)
β_{transit} :	Transit base utility for intermediate trip	0.26**	(0.10)
β_{Uber} :	Uber base utility for intermediate trip	-1.97***	(0.11)
$\Delta_{\text{transit}}^{\alpha}$:	Change in transit utility for primary trip	0.60***	(0.12)
f_{Uber} :	Uber cost of access, fixed component	5.47***	(0.15)
c_{walk} :	Walking cost of access, distance component	4.45***	(0.30)
c_{Uber} :	Uber cost of access, distance component	1.21***	(0.36)
σ_{transit} :	Transit preference dispersion (SD)	2.18***	(0.55)
σ_{Uber} :	Uber preference dispersion (SD)	2.83***	(0.18)
ϕ :	Uber and transit preference correlation	-0.74***	(0.06)
λ :	Rate of spatial dispersion of last-mile trips around station	22.21***	(0.57)
κ :	Density of potential last-mile travelers around station	0.55***	(0.16)
Number of moments		24	
Number of parameters		13	
Degrees of freedom		11	
RMSE		0.0146	

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table contains coefficient estimates and standard errors from the model estimated in Section 6. The α and β utility parameters are relative to private vehicle utility. The probability a last-mile trip originates d kilometers from the station is given by $\rho(d) = \frac{\lambda e^{-\lambda d}}{1 - e^{-\lambda D_{\max}}}$, where $D_{\max}$ is set to 5 km. The fraction of potential last-mile travelers d kilometers from a station is given by $\exp(-\kappa \cdot d)$. The probability of an intermediate trip p_s is set to 0.5. Errors are computed using numerical gradients with $N = 650$ station openings.


Table 6: Counterfactuals

Scenario	Commute (pp)			Intermediate (pp)			Transit, welfare, & fiscal				
	Tr. walk	Tr. Uber	Dir. Uber	Car	Transit	Uber	Car	Δ Tr. (%)	ΔW (%)	ΔW (\$)	Fiscal (\$)
<i>Panel A: Ride-hailing ban</i>											
With a train station	+0.31	-0.86	-34.48	+35.03	-4.53	-3.76	+8.29	-9.87	-46.92	-3.05	—
Without a train station	+0.05	-0.52	-34.65	+35.11	-4.64	-3.70	+8.34	-10.93	-48.53	-3.04	—
<i>Panel B: Last-mile subsidies (flat \$2.50 subsidy vs. a per-km subsidy with the same net fiscal cost)</i>											
Flat: \$2.50 flat subsidy	-0.38	+1.86	+0.03	-1.51	+0.28	+0.03	-0.31	+5.69	+1.25	+0.08	-0.14
Per-km: 65% of per-km charge	-0.16	+1.77	+0.07	-1.67	+0.31	+0.03	-0.34	+6.18	+1.17	+0.08	-0.14
<i>Panel C: Taxes (\$2.00 ride-hailing tax vs. a congestion charge with the same net fiscal cost)</i>											
Ride-hailing tax (\$2.00)	+0.31	-0.54	-10.89	+11.12	-0.46	-2.31	+2.76	-1.60	-18.84	-1.22	+0.99
Congestion charge (\$0.56)	+1.24	+0.20	-0.10	-1.34	+0.25	+0.03	-0.28	+5.50	-15.29	-0.99	+0.99
<i>Panel D: Fiscally-neutral policy mix that increases transit ridership by 5 percentage points</i>											
RH tax + LM subsidies $s_c^{LM} = 1.93, s_f^{LM} = 3.35, t^{RH} = 1.28$	-0.57	+6.29	-7.04	+1.32	+1.06	-1.63	+0.57	+21.90	-8.06	-0.52	-0.00
Cong. charge + RH + LM $s_c^{LM} = 1.60, s_f^{LM} = 2.02, t^{RH} = 0, t^{cong.} = 0.35$	+0.02	+5.69	+0.05	-5.77	+1.07	+0.11	-1.18	+21.90	-5.59	-0.36	+0.00
C + RH + LM + transit subsidy $s^{Tr} = 1.00, s_c^{LM} = s_f^{LM} = 0, t^{RH} = 0.03, t^{cong.} = 0.40$	+3.42	+0.60	+0.01	-4.03	+4.46	-1.14	-3.32	+21.90	-1.22	-0.08	-0.00

Notes: This table contains counterfactual estimates under different scenarios using the model estimates in Table 5. Each row corresponds to a different scenario, compared to a baseline in which that scenario did not occur. For example, row 1 shows changes in traveler behavior in response to a ride-hailing ban. Columns (2) through (5) calculate the percentage point change in different mode shares averaged across outbound and return legs. Columns (6) through (8) calculate percentage point changes in intermediate trip mode shares. Column (9) captures the average trip percentage increase in transit ridership. Columns (10) and (11) calculates the welfare change in percentage and WTP terms respectively, relative to the baseline. Column (12) calculates the fiscal cost of a policy change per traveler, summed across all trip types. All columns are averaged over distances through 2.5 kilometers from the station using the κ -population weights.


Figures

Figure 2: Model simulations of the response to an improvement in public transit




Notes: This panel contains figures from three separate model simulations. Each model simulation corresponds to a column of figures. Row 1 contains average initial mode shares for outbound and return trips plotted by distance from the transit station. Row 2 contains changes in average outbound and return trip mode shares by distance from the transit station. Row 3 contains log changes in the number of trips at each distance from the transit station. First- and last-mile trips are assumed to be located at $d = 0$ but have some measurement error, following an exponential distribution with rate parameter $\lambda = 5$. Details of the model and calculations are included in the text and Appendix C.

Figure 3: Locations of transit expansions in our data




Notes: This map plots the locations of new transit openings in our data. The size of each circle is proportional to the number of new stations in our sample for each city.

Figure 4: Quadrant I - Effects of rail station opening on Uber trips over time




Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips 0–100 m away (Top panel) and 300–400 m (Bottom panel), relative to 1100–1200 m away. The coefficients are plotted for each month between six months prior and six months after an expansion. The PPML estimation follows a dynamic difference-in-differences model of station openings on Uber trips that either originated or terminated at a given distance band from the opening. The model contains station by distance band and station by time fixed effects; additionally, fixed effects for distance band by more than 6 months before or after an expansion are included. All errors are clustered at the station level.

Figure 5: Quadrant I - Effects of rail station opening on Uber trips by distance band (average effect months 1-6)




Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating equation (5). These coefficients estimate the effect of a new train station opening on Uber trips at varying distances from the train station during months 1 through 6, relative to the farthest distance band 1100–1200 m, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station.

Figure 6: Quadrant I - Event study estimates of rail station opening on Uber trips 1100–1200 m from station




Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips 1100–1200 m away. The coefficients are plotted for each month between six months prior and six months after an expansion. The regression is an event study model of station openings on the inverse hyperbolic sine of Uber trips (in 100s) that either originated or terminated at a given distance band from the opening. The model contains station and city by time fixed effects; additionally, fixed effects for more than 6 months before or after an expansion are included. The model is estimated using [de Chaisemartin and D'Haultfœuille \(2026\)](#). All errors are clustered at the station level.

Figure 7: Quadrant I - Heterogeneous effects of rail station opening on Uber trips by city, 0–100 m and 300–400 m from the station




Notes: This funnel graph shows city-specific treatment effects based on equation (5) where the coefficient α is allowed to vary by city. Panel A contains estimates for 0–100 m whereas Panel B contains estimates for 300–400 m from the station. The x-axis shows coefficient estimates, the y-axis shows standard errors. The region in white contains estimates that are not statistically different from zero. The light, medium, and dark gray regions contain estimates that are statistically different from zero at 10%, 5%, and 1%, respectively. The size of each circle is proportional to baseline Uber ridership. The large red triangle indicates the average effect reported in Figure 5.

Figure 8: Quadrant I - Effect of rail station opening on kilometers per Uber trip by distance from station




Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating equation (5) for the effect of a new train station opening on kilometers travelled per Uber trip at varying distances from the train station. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station.

Figure 9: Quadrant II - Effect of rail station opening on Uber trips by distance and POI group



Notes: This figure plots the average treatment effect estimates (hollow symbols) and 95% confidence intervals (bars) for the effect of a new train station opening on Uber trips at varying distances from the station, separately by rider POI distance group. Each point corresponds to a separate difference-in-differences regression comparing trips in a given 100 meter distance band against a control group of trips farther from the station. Blue circles correspond to riders whose POI is within 500 meters of the station, red diamonds to riders whose POI is between 500 and 800 meters, and green triangles to riders whose POI is between 800 and 1200 meters. For the POI 0–500 m and POI 500–800 m groups, the control group is trips in the 800–1200 meter distance band within the POI 800–1200 m group. For the POI 800–1200 m group, the control group is the 1100–1200 meter distance band within the same group. Each distance is measured over a 100 meter band ending at the given distance; for example, the coefficient at 400 meters reports the change in trips between 300 and 400 meters from the station. The average treatment effect is computed over months 1–6 following station opening. All specifications include station-by-distance and station-by-date fixed effects, as well as common event-time fixed effects. Standard errors are clustered at the station level.

Figure 10: Decomposing the effect of transit improvements on Uber usage



Notes: The figure traces the effects of improved transit on Uber usage through direct substitution, option value, and last-mile complementarity. We focus on travelers located within 2.5 kilometers of a station. All values are percentage changes relative to the primary mode's baseline share. Within 2.5 kilometers, these shares are 30% for Uber, 50% for automobile, and 20% for transit. The "transit-by-walk" box depicts the percent increase in total transit use due to transit-by-walk for primary trips. Arrows are labeled with usage changes only when multiple arrows enter the same box. The bottom row reports each channel's contribution to Uber usage and their sum, which is exactly equal to the model-implied change in total Uber usage. Dashed lines and lower contrast indicate branches that do not affect Uber usage. We use Proposition 1 and Appendix C.3 to perform the decomposition. Appendix E describes the computation.

Online Appendix

Contents

A	Supplemental tables	2
A.1	Descriptives	2
A.2	Robustness	3
A.3	Model estimates and counterfactuals	6
B	Supplemental figures	8
B.1	Descriptives	8
B.2	Quadrant I Results	10
B.3	Quadrant II and III results	16
B.4	Robustness	18
B.5	Counterfactuals	26
C	Model appendix	27
C.1	Finding choice probabilities	27
C.1.1	Value function for outbound journey (Period 1)	27
C.1.2	Conditional choice probabilities	28
C.1.3	Unconditional choice probabilities	29
C.2	Theoretical results	30
C.3	Mode-by-mechanism decomposition behind Figure 10	38
C.4	Mapping the model to the data	44
C.5	Model simulations	45
D	Empirics appendix	46
D.1	Event study for far distances	46
D.2	Adjacent distance bins methodology	46
E	Model estimation and counterfactuals appendix	48

A Supplemental tables

A.1 Descriptives

Table A1: City-level rail service openings

City	# Exp.	# Stn.	Light	Heavy	Subway	First Yr	Last Yr	Uber Yr.
Amsterdam	1	6	0	0	6	2018	2018	2012
Baku	1	1	0	0	1	2016	2016	2015
Bangalore	5	34	0	0	34	2014	2017	2013
Bangkok	4	18	0	0	18	2016	2017	2014
Boston	1	1	0	0	1	2014	2014	2012
Bratislava	1	1	1	0	0	2016	2016	2015
Bucharest	1	1	0	0	1	2017	2017	2015
Buenos Aires	1	1	0	0	1	2018	2018	2016
Charlotte	1	11	11	0	0	2018	2018	2013
Chennai	5	29	0	0	29	2015	2018	2014
Chicago	2	2	0	0	2	2015	2017	2012
Cleveland	1	1	0	0	1	2015	2015	2014
Dallas	3	6	6	0	0	2014	2016	2012
Delhi	13	65	0	0	65	2014	2018	2013
Denver	2	15	15	0	0	2013	2017	2012
Dubai	5	18	11	0	7	2011	2017	2013
Dublin	1	4	4	0	0	2017	2017	2014
Dusseldorf	1	5	5	0	0	2016	2016	2014
Edmonton	1	2	2	0	0	2015	2015	2014
Fortaleza	3	7	6	0	1	2017	2018	2016
Frankfurt	1	9	0	0	9	2016	2016	2014
Hong Kong	4	10	0	0	10	2014	2016	2014
Hyderabad	2	39	0	0	39	2017	2018	2014
Istanbul	6	17	0	0	17	2014	2017	2014
Jaipur	1	9	0	0	9	2015	2015	2014
Kochi	3	16	0	0	16	2010	2017	2014
Kuala Lumpur	5	50	22	0	28	2015	2017	2013
Lisbon	1	1	0	0	1	2016	2016	2014
Los Angeles	2	10	10	0	0	2016	2016	2012
Lucknow	1	7	0	0	7	2017	2017	2016
Milan	7	10	0	0	10	2014	2015	2013
Moscow	18	41	0	0	41	2010	2021	2013
New Jersey	1	1	0	1	0	2016	2016	2013
Novosibirsk	1	2	2	0	0	2016	2016	2015
Oslo	1	1	0	0	1	2016	2016	2014
Panama City	2	2	0	0	2	2015	2015	2014
Paris	1	1	0	0	1	2013	2013	2012
Portland	1	3	3	0	0	2015	2015	2014
Prague	1	4	0	0	4	2015	2015	2014
Rio de Janeiro	1	5	0	0	5	2016	2016	2014
Rome	3	22	0	0	22	2014	2015	2013
Sacramento	1	3	3	0	0	2015	2015	2013
Salvador	4	11	0	0	11	2016	2018	2016
San Francisco	5	15	0	11	4	2014	2018	2012
Santiago	1	10	0	0	10	2017	2017	2013
Santo Domingo	1	4	0	0	4	2018	2018	2015
Sao Paulo	9	15	0	0	11	2017	2018	2014
Seattle	2	3	3	0	0	2016	2016	2012
Singapore	9	39	6	0	33	2013	2017	2013
Stockholm	1	1	0	1	0	2017	2017	2013
Taipei	3	21	0	0	21	2014	2017	2013
Toronto	2	10	0	4	6	2015	2017	2012
Vienna	1	5	0	0	5	2017	2017	2014
Warsaw	1	7	0	0	7	2015	2015	2014
Washington D.C.	1	5	0	0	5	2014	2014	2012

Notes: This table contains summary statistics on the 650 stations in our sample. Column (1) contains the 55 cities in our sample. Columns (2) through (6) count the number of expansions, the number of stations associated with those expansions, and whether they were for light rail, heavy rail, or subways, respectively. Column (7) and (8) denote the first and last year we have Uber data for that city. Column (9) contains the year that Uber entered the market.

A.2 Robustness

Table A2: Quadrant I - Effects of rail station opening on Uber trips 0–100 m from station with alternative fixed effects

	(1)	(2)	(3)	(4)
	Uber trips	Uber trips	Uber trips	Uber trips
-6 months after	-0.00966 (0.0262)	0.00671 (0.0260)	0.00676 (0.0259)	0.0399 (0.0304)
-5 months after	-0.00417 (0.0223)	0.00506 (0.0212)	-0.00327 (0.0234)	0.0239 (0.0263)
-4 months after	0.0177 (0.0211)	0.00980 (0.0181)	0.00756 (0.0194)	0.0229 (0.0205)
-3 months after	0.0125 (0.0157)	0.00807 (0.0141)	0.0124 (0.0160)	0.0201 (0.0169)
-1 months after	0.0324 (0.0233)	0.0365* (0.0188)	0.0482** (0.0225)	0.0478** (0.0211)
Month of expansion	0.141*** (0.0285)	0.161*** (0.0266)	0.180*** (0.0292)	0.177*** (0.0296)
1 months after	0.328*** (0.0415)	0.343*** (0.0417)	0.350*** (0.0428)	0.342*** (0.0429)
2 months after	0.336*** (0.0437)	0.366*** (0.0440)	0.366*** (0.0432)	0.350*** (0.0432)
3 months after	0.353*** (0.0459)	0.379*** (0.0470)	0.372*** (0.0456)	0.351*** (0.0460)
4 months after	0.391*** (0.0494)	0.410*** (0.0488)	0.396*** (0.0454)	0.375*** (0.0459)
5 months after	0.407*** (0.0605)	0.427*** (0.0584)	0.406*** (0.0556)	0.386*** (0.0560)
6 months after	0.436*** (0.0629)	0.469*** (0.0624)	0.451*** (0.0603)	0.423*** (0.0610)
Observations	72,618	70,998	58,298	58,298
Station Openings	635	634	608	608
Station x Dist FE	✓	✓	✓	✓
Cal. and Evt. Time FE	✓	✓	✓	✓
City x Date FE		✓		
Station x Date FE			✓	✓
Dist x Date FE				✓

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table contains coefficient estimates and standard errors from equation (4) for trips 0–100 meters from the station under alternative fixed effects. These coefficients estimate the effect of a new train station opening on Uber trips 0–100 meters away in every month leading up to and after the station opening relative to Uber trips 1100–1200 meters away. Coefficients are relative to two months prior to the station opening. All specifications include station and calendar and event-time fixed effects. Column (2) adds city by calendar time fixed effects. Column (3) adds station by date fixed effects, and is our main specification. Column (4) includes distance by date fixed effects. All errors are clustered at the station level.

Table A3: Quadrant I - Effects of rail station opening on Uber trips 300–400 m from station with alternative fixed effects

	(1) Uber trips	(2) Uber trips	(3) Uber trips	(4) Uber trips
-6 months after	-0.0117 (0.0137)	0.000533 (0.0107)	0.00584 (0.0113)	-7.43e-06 (0.0114)
-5 months after	0.00198 (0.0118)	0.00527 (0.00843)	0.00755 (0.00873)	0.00394 (0.00892)
-4 months after	0.00400 (0.0119)	0.00563 (0.00935)	0.00864 (0.00979)	0.00817 (0.00998)
-3 months after	0.00527 (0.00813)	0.00476 (0.00627)	0.00603 (0.00658)	0.00609 (0.00650)
-1 months after	-0.00716 (0.0120)	-0.00964 (0.00854)	-0.00998 (0.00674)	-0.00825 (0.00690)
Month of expansion	-0.0297*** (0.0110)	-0.0251** (0.0104)	-0.0163* (0.00918)	-0.0134 (0.00927)
1 months after	-0.0472*** (0.0130)	-0.0431*** (0.0120)	-0.0255** (0.0110)	-0.0208** (0.0106)
2 months after	-0.0562*** (0.0152)	-0.0447*** (0.0139)	-0.0299** (0.0130)	-0.0235* (0.0125)
3 months after	-0.0533*** (0.0148)	-0.0437*** (0.0140)	-0.0340** (0.0135)	-0.0267** (0.0132)
4 months after	-0.0444** (0.0180)	-0.0385** (0.0157)	-0.0302** (0.0140)	-0.0220 (0.0152)
5 months after	-0.0654*** (0.0229)	-0.0627*** (0.0167)	-0.0632*** (0.0154)	-0.0546*** (0.0159)
6 months after	-0.0732*** (0.0234)	-0.0556*** (0.0150)	-0.0482*** (0.0135)	-0.0381** (0.0151)
Observations	73,407	71,706	60,624	60,624
Station Openings	637	637	620	620
Station x Dist FE	✓	✓	✓	✓
Cal. and Evt. Time FE	✓	✓	✓	✓
City x Date FE		✓		
Station x Date FE			✓	✓
Dist x Date FE				✓

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table contains coefficient estimates and standard errors from equation (4) for trips 300–400 meters from the station under varying fixed effects. These coefficients estimate the effect of a new train station opening on Uber trips 300–400 meters away in every month leading up to and after the station opening relative to Uber trips 1100–1200 meters away. Coefficients are relative to two months prior to the station opening. All specifications include station and calendar and event-time fixed effects. Column (2) adds city by calendar time fixed effects. Column (3) adds station by date fixed effects, and is our main specification. Column (4) includes distance by date fixed effects. All errors are clustered at the station level.

Table A4: Quadrant III - Effects of rail station opening on citywide Uber trips for 0–500 m POI sample

	(1)	(2)	(3)	(4)
	Citywide Uber trips	Citywide Uber trips	Citywide Uber trips	Citywide Uber trips
-6 months after	-0.0110 (0.00728)	-0.00301 (0.00575)	-0.00200 (0.00492)	0.00727 (0.00678)
-5 months after	-0.0107* (0.00546)	-0.00606 (0.00419)	-0.00401 (0.00357)	0.00391 (0.00456)
-4 months after	-0.00633 (0.00543)	-0.00393 (0.00413)	6.75e-05 (0.00357)	0.00639 (0.00427)
-3 months after	0.00118 (0.00462)	0.00137 (0.00373)	0.00325 (0.00368)	0.00715* (0.00413)
-1 months after	-0.00370 (0.00486)	-0.00754** (0.00371)	-0.00696** (0.00328)	-0.0103** (0.00418)
Month of expansion	-0.0145** (0.00627)	-0.0148*** (0.00381)	-0.0134*** (0.00361)	-0.0194*** (0.00438)
1 months after	-0.0282*** (0.00777)	-0.0253*** (0.00461)	-0.0244*** (0.00401)	-0.0318*** (0.00440)
2 months after	-0.0265*** (0.00906)	-0.0218*** (0.00569)	-0.0198*** (0.00491)	-0.0288*** (0.00498)
3 months after	-0.0317*** (0.00961)	-0.0242*** (0.00534)	-0.0228*** (0.00445)	-0.0336*** (0.00555)
4 months after	-0.0361*** (0.0112)	-0.0236*** (0.00564)	-0.0217*** (0.00415)	-0.0360*** (0.00648)
5 months after	-0.0349*** (0.0108)	-0.0226*** (0.00546)	-0.0215*** (0.00430)	-0.0385*** (0.00761)
6 months after	-0.0343*** (0.0126)	-0.0237*** (0.00673)	-0.0215*** (0.00521)	-0.0421*** (0.00804)
Observations	63,701	62,090	54,492	54,492
Station Openings	555	555	529	529
Station x POI FE	✓	✓	✓	✓
Cal. and Evt. Time FE	✓	✓	✓	✓
City x Date FE		✓		
Station x Date FE			✓	✓
POI x Date FE				✓

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table contains coefficient estimates and standard errors from equation (4) for individuals with a POI near the station, and with varying fixed effects. These coefficients estimate the effect of a new train station opening on *citywide* Uber trips in every month leading up to and after the station opening, for those with a POI within 500 m of the station opening relative to those with a POI within 800–1200 meters. Coefficients are relative to two months prior to the station opening. All specifications include station and calendar and event-time fixed effects. Column (2) adds city by calendar time fixed effects. Column (3) adds station by date fixed effects, and is our main specification. Column (4) includes distance by date fixed effects. All errors are clustered at the station level.

A.3 Model estimates and counterfactuals

Table A5: Model fit

Moment	Data	Model
<i>Ia: Trips near station</i>		
0–100 m	0.388	0.384
100–200 m	0.047	0.034
200–300 m	−0.028	−0.049
300–400 m	−0.034	−0.042
400–500 m	−0.008	−0.029
500–600 m	−0.019	−0.019
600–700 m	−0.013	−0.011
700–800 m	−0.003	−0.005
800–900 m	−0.004	−0.002
900–1000 m	−0.003	0.000
1000–1100 m	0.005	0.002
1100–1200 m	-	0.002
<i>Ib: Trip length near station</i>		
0–100 m	−0.168	−0.169
100–200 m	−0.063	−0.085
1100–1200 m	-	−0.000
<i>II: Local trips (POI)</i>		
100–200 m	−0.078	−0.080
200–300 m	−0.083	−0.054
300–400 m	−0.062	−0.038
400–500 m	−0.043	−0.024
500–600 m	−0.012	−0.013
600–700 m	−0.015	−0.003
700–800 m	−0.010	0.003
800–1200 m	-	0.012
<i>IIIa: Non-Local trips (POI)</i>		
0–500 m	0.053	0.027
500–800 m	0.030	0.037
800–1200 m	-	0.023
<i>IIIb: Citywide trips (POI)</i>		
0–500 m	−0.022	−0.048
500–800 m	−0.002	−0.001
800–1200 m	-	0.013

Notes: This table compares the moments used in estimation with those implied by the model. Each Panel corresponds to a different set of moments. Column (1) indicates the relevant distance group for that moment. Column (2) shows the estimate from the reduced-form results in Section 5. Column (3) shows the model-implied equivalent moment.

Table A6: Distribution of welfare changes in fiscally neutral policy mix


Distance (m)	Percentiles (\$)					Mean (\$)
	P10	P25	Median	P75	P90	
0	-0.68	-0.37	-0.05	+0.02	+0.03	-0.19
500	-0.70	-0.66	-0.41	+0.08	+0.22	-0.31
1000	-0.70	-0.69	-0.65	-0.32	+0.77	-0.34
1500	-0.71	-0.70	-0.67	-0.55	+0.09	-0.35
2000	-0.71	-0.70	-0.68	-0.60	-0.26	-0.40
2500	-0.71	-0.70	-0.69	-0.63	-0.37	-0.46

Notes: This table calculates welfare changes in WTP terms for the policy in the second to last row of Table 6. The policy mix solves for the fiscally neutral set of last-mile subsidies, ride-hailing taxes, and congestion taxes that would deliver a 5 pp increase in average transit use, while minimizing welfare losses. Each column shows a different percentile of the welfare change distribution. Each row shows a different distance from the station those distributions are calculated at.

B Supplemental figures

B.1 Descriptives

100 m rings around stations



174 m hexes around stations





Figure A1: Spatial sampling for Toronto-York Spadina subway extension in 2017

Notes: These maps depict two ways to sample Uber trips in our data. Blue dots indicate the locations of 6 subway openings along the 1-line in Toronto's subway system. The map on the left uses concentric rings of 100 meter width from the location of the station. The map on the right uses fixed width hexagons at varying distances from the new station. We use these to determine points-of-interest (POIs) for individual riders.

Figure A2: Transit expansions over time



Notes: This figure plots the time series of expansions and station openings in our sample. The left-side y-axis contains counts of rail expansions, which are plotted in blue. The right-side y-axis contains counts of stations associated with those expansions, which are plotted in red.





Figure A3: Distribution of days and months of station openings

Notes: These histograms plot the temporal distribution of station openings. The x-axis is the fraction of station openings in our sample that fall within a certain day/month. The top figure plots the distribution of numeric day within the month of expansion. The bottom figure plots the distribution of month within year of expansion.

B.2 Quadrant I Results




Figure A4: Results over time for each 100 m ring around station

Notes: These subfigures plot the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips for the labeled distance bands. The coefficients are plotted for each month between six months prior and six months after an expansion. The regression is a dynamic difference-in-differences model of station openings on Uber trips that either originated or terminated at a given distance band from the opening, and the regression equation is given by equation (4). The model contains station by distance band and station by time fixed effects; additionally, fixed effects for distance band by more than 6 months before or after an expansion are included. All errors are clustered at the station level.




Figure A5: Event study estimates of rail station opening on Uber trips 1100–1200 m from station using PPML estimator

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips at our farthest distance band: 1100–1200 meters. It is a complement to the event study analysis depicted in Figure 6. The coefficients are plotted for each month between six months prior and six months after an expansion. The PPML estimates follow an event study model of station openings on Uber trips that either originated or terminated at a given distance band from the opening. The model contains station and city by time fixed effects; additionally, fixed effects for more than 6 months before or after an expansion are included. All errors are clustered at the station level.




Figure A6: Event study estimates of rail station opening on Uber trips 1100–1200 m from station using two-way fixed effects linear regression and inverse hyperbolic sine transformed Uber trips (in hundreds)

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips at our farthest distance band: 1100–1200 meters. It is a complement to the event study analysis depicted in Figure 6, and corresponds to the simpler two-way fixed effects linear regression model. The coefficients are plotted for each month between six months prior and six months after an expansion. The regression is an event study model of station openings on Uber trips that either originated or terminated at a given distance band from the opening. The model contains station and city by time fixed effects; additionally, fixed effects for more than 6 months before or after an expansion are included. All errors are clustered at the station level.





Figure A7: Effects of rail station opening on Uber pickups and drop-offs by distance band (average effect months 1-6)

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating equation (5) separately for pickups and drop-offs. These coefficients estimate the effect of a new train station opening on Uber pickups and drop-offs at varying distances from the train station during months 1 through 6, relative to the farthest distance band, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station. All errors are clustered at the station level.

Effects of rail station opening on Uber total fares by distance band (average effect months 1-6)



Effects of rail station opening on total kilometers traveled on Uber trips by distance band (average effect months 1-6)




Figure A8: Estimates for total fares and total distance on Uber trips

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating equation (5) for the total fare paid (top figure) and kilometers traveled (bottom figure) along Uber trips. These coefficients estimate the effect of a new train station opening on total Uber fares and trip lengths at varying distances from the train station during months 1 through 6, relative to the farthest distance band, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station. All errors are clustered at the station level.





Figure A9: Effects of rail station opening on Uber trips by distance band (average effect months 1-6) for terminal and non-terminal stations

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating equation (5) separately for stations that were a terminal stop on their train line. The coefficients estimate the effect on Uber trips at varying distances from the train station opening, separately for each type of train station. All errors are clustered at the station level.

Panel A: Interaction of rail opening with normalized travel time from CBD by distance from station



Panel B: Interaction of rail opening with normalized travel distance from CBD by distance from station




Figure A10: Differential effects for peripheral stations

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating a version of equation (5) that interacts each effect with a normalized measure of distance from the CBD. The measure uses travel time in Panel A and physical distance in Panel B. The coefficients estimate the different effect on Uber trips at varying distances from the train station opening. All errors are clustered at the station level.

B.3 Quadrant II and III results




Figure A11: Results over time for nearby trips in POI sample for each 100 m ring around station

Notes: This figure contains event study estimates of the effect of rail station openings on Uber trips for riders whose POI is near a station. Each panel shows a separate difference-in-differences estimate comparing trips in a given 100 meter distance band to a control group of trips farther from the station. The first four panels correspond to riders whose POI is within 500 meters of the station (POI 0–500 m group), the next three to riders whose POI is between 500 and 800 meters (POI 500–800 m group), and the final three to riders whose POI is between 800 and 1200 meters (POI 800–1200 m group). For the POI 0–500 m and POI 500–800 m groups, the control group is trips in the 900–1200 meter distance band within the POI 800–1200 m group. For the POI 800–1200 m group, the control group is the 1100–1200 meter distance band within the same group. All regressions include station-by-distance and station-by-date fixed effects, as well as common event-time fixed effects. Standard errors are clustered at the station level. The omitted period is two months prior to opening.




Figure A12: Results over time for citywide, nearby and non-local trips in POI sample

Notes: This figure presents event study estimates of the effect of rail station openings on Uber trips for riders in the POI sample. Each panel shows coefficients from equation (4), estimated separately for each month relative to station opening. The first row corresponds to riders whose pre-opening point of interest (POI) is within 500 meters of the station (POI 0–500 m group) and the second row to riders whose POI is between 500 and 800 meters (POI 500–800 m group). Columns show results for citywide, local, and non-local trips respectively. Citywide trips are all Uber trips taken by riders in the POI sample regardless of where they occur. Local trips are near-station trips within the rider’s own POI distance band; non-local trips are trips occurring farther from the station than the outer edge of the rider’s POI distance band. The control group in all specifications is the POI 800–1200 m group. All specifications include station-by-POI, station-by-date, and event-time fixed effects. Standard errors are clustered at the station level. The omitted period is two months prior to opening. Bars denote 95% confidence intervals.

B.4 Robustness




Figure A13: Placebo tests for all 100 m distance bands, up through 1000–1100 m

Notes: This figure plots the results of a placebo test where we randomly assign placebo opening dates to each transit station. To implement this test, we randomly assign each station a new opening date with uniform probabilities from 6 months after the station’s first observation to 6 months prior to the station opening. This ensures that we have sufficient observations to estimate the event-time specific coefficients above. Also, it ensures that the new station opening is prior to any actual effects on Uber trips. We then re-estimate our baseline model using the “fake” opening dates. We perform this 100 times and use a box-whisker plot to depict the distribution of these estimates. The blue circle denotes estimates from the true station opening dates. We perform this placebo analysis for each distance band.




Figure A14: Effects of rail station opening on Uber fare per trip by distance band (average effect months 1-6)

Notes: This figure plots the coefficients (circles and squares) and 95% confidence intervals (bars) from estimating equation (5) for the Uber fare in USD per trip. We look at a change in the fixed cost (the main effect, in blue) and the marginal cost per kilometer (the interacted effect with kilometers per trip, in red). These coefficients estimate the change in fare structure from a new train station opening at varying distances from the train station during months 1 through 6, relative to the farthest distance band, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station. All errors are clustered at the station level.




Figure A15: Event study estimates of rail station opening on Uber trips estimated via PPML

Notes: These figures plot the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips for each 100 meter distance band. It relies on the simpler two-way fixed effects PPML regression model. The coefficients are plotted for each month between six months prior and six months after an expansion. The regression is an event study model of station openings on Uber trips that either originated or terminated at a given distance band from the opening. The model contains station and city by time fixed effects; additionally, fixed effects for more than 6 months before or after an expansion are included. All errors are clustered at the station level.





Figure A16: Event study estimates of rail station opening on inverse hyperbolic sine Uber trips (in hundreds) estimated via linear regression

Notes: These figures plot the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips for each 100 meter distance band. It relies on the simpler two-way fixed effects linear regression model and the inverse hyperbolic sine transformation of the outcome Uber trips. The coefficients are plotted for each month between six months prior and six months after an expansion. The regression is an event study model of station openings on Uber trips that either originated or terminated at a given distance band from the opening. The model contains station and city by time fixed effects; additionally, fixed effects for more than 6 months before or after an expansion are included. All errors are clustered at the station level.

Effects of rail station opening on Uber trips relative to adjacent 100-meter band as control instead of the 1100–1200 m band (Marginal effects)



Effects of rail station opening on Uber trips (Cumulative effects)




Figure A17: Difference-in-differences estimates using the adjacent bin as a control group

Notes: These figures plot the coefficients (circles) and 95% confidence intervals (bars) from estimating a version of equation (5) that relies on adjacent distance bins as the control group. In the top figure, these coefficients estimate the effect of a new train station opening on Uber trips at varying distances from the train station during months 1 through 6, relative to the *next* farthest distance band, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station relative to 400 and 500 m away. In the bottom figure, those coefficients are summed over the entire range to account for the effect the new station opening has on the “control” group at close enough distances. This procedure is described in Section 5 as well as Appendix D.2. All errors are clustered at the station level.




Figure A18: Effects of rail station opening on Uber trips by distance band using hexagons instead of rings

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) from estimating a version of equation (5) estimated on data aggregated to the hexagon level. Hexagons are assigned distances from the station according to their centroid. These coefficients estimate the effect of a new train station opening on Uber trips at varying distances from the train station during months 1 through 6, relative to the farthest distance band, and relative to months -6 through -3. Each distance is measured over a 100 m band ending at the given distance; for example, the coefficient at 400 m reports the change in trips between 300 and 400 m away from a train station. All errors are clustered at the station level.




Figure A19: Effects of rail station opening on Uber trips using an 18-month window

Notes: This figure plots the coefficients (circles) and 95% confidence intervals (bars) for the percentage effect of a new train station opening on Uber trips 0–100 m from the station. The coefficients are plotted for each month between 18 months prior and 18 months after an expansion. The regression is a dynamic difference-in-differences model of station openings on Uber trips that either originated or terminated at a given distance band from the opening, and the regression equation is given by equation (4). The model contains station by distance band and station by time fixed effects; additionally, fixed effects for distance band by more than 18 months before or after an expansion are included. All errors are clustered at the station level.






Figure A20: Placebo rail openings for citywide, local, and non-local trips

Notes: This figure presents placebo tests for the difference-in-differences estimates of the effect of rail station openings on Uber trips in the POI sample. For each of 100 simulations, we randomly reassign opening dates to stations within a window of 6 months before the true opening date to 6 months after the first observation for that station and re-estimate the average treatment effect over months 1–6 following the placebo opening. Each panel shows the distribution of placebo average treatment effects as a histogram, with the actual estimated average treatment effect marked by a red vertical line. The first row corresponds to riders whose pre-opening point of interest (POI) is within 500 meters of the station (POI 0–500 m group) and the second row to riders whose POI is between 500 and 800 meters (POI 500–800 m group). Columns show results for citywide, local, and non-local trips respectively. The control group in all specifications is the POI 800–1200 m group. All specifications include station-by-POI, station-by-date, and event-time fixed effects. Standard errors are clustered at the station level.


Figure A21: Heterogeneous effects of rail station opening by city for POI sample




(a) Citywide, POI 0–500 m




(b) Citywide, POI 500–800 m




(c) Local, POI 0–500 m



(d) Local, POI 500–800 m



(e) Non-local, POI 0–500 m




(f) Non-local, POI 500–800 m

Notes: Each panel is a funnel chart showing city-level average treatment effect estimates for the effect of rail station openings on Uber trips in the POI sample. Each circle represents one city, with the horizontal axis showing the square root of the estimated average treatment effect (months 1–6 after opening) and the vertical axis showing the square root of the standard error. Circle size is proportional to baseline trip volume in the pre-opening period. The red triangle marks the aggregate average treatment effect across all cities. Shaded regions indicate the 10%, 5%, and 1% significance thresholds. The left column corresponds to riders whose pre-opening point of interest (POI) is within 500 meters of the station (POI 0–500 m group) and the right column to riders whose POI is between 500 and 800 meters (POI 500–800 m group). Rows show results for citywide, local, and non-local trips respectively. The control group in all specifications is the POI 800–1200 m group. All specifications include station-by-POI, station-by-date, and event-time fixed effects. Standard errors are clustered at the station level.

B.5 Counterfactuals

Figure A22: Model-predicted change in transit use under transportation taxes



Notes: Each panel plots the model-predicted percent change in transit use relative to baseline (with the new transit station open) as the indicated instrument is varied from a \$2.50 subsidy to a \$5.00 tax, in 10-cent increments. Panel (a) taxes (subsidizes, if negative) all ride-hailing trips, including last-mile access and egress trips. Panel (b) taxes (subsidizes, if negative) commute-leg car and ride-hailing trips only, leaving last-mile ride-hailing pricing unchanged.

C Model appendix

This appendix completes the model in three stages. We start by deriving the model's closed-form choice probabilities, which we then use to prove our propositions referenced in the text. Next, we derive results that help map the model to the data. Finally, we describe the model simulations.

C.1 Finding choice probabilities

In this section, we first derive the value function for the outbound journey using backward induction, and then compute conditional choice probabilities for all three trips, which have the familiar logit formulation. These calculations are conditional on the persistent preference vector η , which is known to the traveler prior to deciding which mode to take for the outbound trip. We then calculate unconditional probabilities by integrating over the distribution $F(\eta)$.

Throughout, we use the result that when idiosyncratic shocks are i.i.d. Type I Extreme Value,

$$E[\max_k (V_k + \epsilon_k)] = \ln \left[\sum_k \exp(V_k) \right] + \gamma, \quad (6)$$

where $\gamma \approx 0.5772$ is Euler's constant. Going forward, we omit Euler's constant from this equation, as it does not affect decisions.

C.1.1 Value function for outbound journey (Period 1)

When deciding which mode to take for the outbound journey, the traveler considers both the utility they receive in this period, as well as how their choice impacts their expected utility in later periods.

When making the return journey, those who drove for the outbound journey must drive home ($m_3 = \text{car}$ if and only if $m_1 = \text{car}$). Those who didn't drive for the outbound journey choose between transit and Uber; if they choose transit, they must also choose an egress mode (walk or Uber). The inclusive value of choosing transit, by equation (1) and (6), is:

$$\begin{aligned} IV^{\text{transit}}(d, \eta) &= \ln [\exp(\alpha_{\text{transit}} - d \cdot c_{\text{walk}} + \eta_{\text{transit}}) + \exp(\alpha_{\text{transit}} - f_{\text{Uber}} - d \cdot c_{\text{Uber}} + \eta_{\text{transit}})] \\ &= \alpha_{\text{transit}} + \eta_{\text{transit}} + \ln [\exp(-d \cdot c_{\text{walk}}) + \exp(-f_{\text{Uber}} - d \cdot c_{\text{Uber}})]. \end{aligned} \quad (7)$$

Again using equations (1) and (6), this gives us that the expected utility from the return journey, conditional on the outbound mode, is:

$$V_3(m_1, d, \eta) = \begin{cases} \alpha_{\text{car}} & \text{if } m_1 = \text{car}, \\ \ln[\exp(IV^{\text{transit}}(d, \eta)) + \exp(\alpha_{\text{Uber}} + \eta_{\text{Uber}})] & \text{if } m_1 \in \{\text{transit}, \text{Uber}\}. \end{cases} \quad (8)$$

With probability p_s , the traveler needs to make an intermediate trip. If she didn't take her personal automobile already, she cannot do so now. Using equations (2) and (6) gives the following

expected utility for the second period:

$$V_2(m_1) = p_s \cdot \begin{cases} \ln[\exp(\beta_{\text{car}}) + \exp(\beta_{\text{transit}}) + \exp(\beta_{\text{Uber}})] & \text{if } m_1 = \text{car} \\ \ln[\exp(\beta_{\text{transit}}) + \exp(\beta_{\text{Uber}})] & \text{if } m_1 \in \{\text{transit}, \text{Uber}\} \end{cases} \quad (9)$$

With the expected utilities from periods 2 and 3, we can now find the expected utility in period 1. In period 1, travelers choose between all three primary modes, and if using transit, choose between the two access modes. The inclusive value of choosing transit is the same as in the last period, so we can use (7).

The expected utility for each outbound mode, considering the utility in period one and in later periods, is:

$$V_1(m_1, d, \eta) = \begin{cases} \alpha_{\text{car}} + V_2(\text{car}) + V_3(\text{car}, d, \eta) & \text{if } m_1 = \text{car} \\ IV^{\text{transit}}(d, \eta) + V_2(\text{transit}) + V_3(\text{transit}, d, \eta) & \text{if } m_1 = \text{transit} \\ \alpha_{\text{Uber}} + \eta_{\text{Uber}} + V_2(\text{Uber}) + V_3(\text{Uber}, d, \eta) & \text{if } m_1 = \text{Uber} \end{cases} \quad (10)$$

Note that, per (9) and (8), $V_2(\text{transit}) = V_2(\text{Uber})$ and $V_3(\text{transit}, d, \eta) = V_3(\text{Uber}, d, \eta)$.

C.1.2 Conditional choice probabilities

We can now use the familiar logit formulation to derive choice probabilities conditional on the preference shocks η_m .

Outbound Journey (Period 1)

The probability of choosing mode m_1 for the outbound journey is:

$$P_1(m_1 | d, \eta) = \frac{\exp(V_1(m_1, d, \eta))}{\sum_{m' \in \{\text{car}, \text{transit}, \text{Uber}\}} \exp(V_1(m', d, \eta))} \quad (11)$$

For travelers who choose transit, the probability of using access mode a_1 is:

$$P_1(a_1 | m_1 = \text{transit}, d, \eta) = \frac{\exp(\alpha_{\text{transit}} - f_{a_1} - d \cdot c_{a_1} + \eta_{\text{transit}})}{\sum_{a' \in \{\text{walk}, \text{Uber}\}} \exp(\alpha_{\text{transit}} - f_{a'} - d \cdot c_{a'} + \eta_{\text{transit}})} \quad (12)$$

Intermediate Trip (Period 2)

The mode choice probabilities for the intermediate trip depend on the mode chosen for the outbound journey.

For travelers who drove outbound:

$$P_2(m_2 | m_1 = \text{car}) = \frac{\exp(\beta_{m_2})}{\exp(\beta_{\text{car}}) + \exp(\beta_{\text{transit}}) + \exp(\beta_{\text{Uber}})} \quad \text{for } m_2 \in \{\text{car}, \text{transit}, \text{Uber}\} \quad (13)$$

For travelers who used transit or Uber outbound:

$$P_2(m_2 \mid m_1 \in \{\text{transit}, \text{Uber}\}) = \begin{cases} \frac{\exp(\beta_{m_2})}{\exp(\beta_{\text{transit}}) + \exp(\beta_{\text{Uber}})} & \text{if } m_2 \in \{\text{transit}, \text{Uber}\} \\ 0 & \text{if } m_2 = \text{car} \end{cases} \quad (14)$$

Return Journey (Period 3)

The mode choice probabilities for the return journey also depend on the mode chosen for the outbound journey. For travelers who drove outbound, the return mode is deterministic:

$$P_3(m_3 \mid m_1 = \text{car}) = \begin{cases} 1 & \text{if } m_3 = \text{car} \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

The conditional choice probabilities for the return journey for those who did not drive outbound are

$$P_3(m_3 \mid m_1 \in \{\text{transit}, \text{Uber}\}, d, \eta) = \begin{cases} \frac{\exp(IV^{\text{transit}}(d, \eta))}{\exp(IV^{\text{transit}}(d, \eta)) + \exp(\alpha_{\text{Uber}} + \eta_{\text{Uber}})} & \text{if } m_3 = \text{transit} \\ \frac{\exp(\alpha_{\text{Uber}} + \eta_{\text{Uber}})}{\exp(IV^{\text{transit}}(d, \eta)) + \exp(\alpha_{\text{Uber}} + \eta_{\text{Uber}})} & \text{if } m_3 = \text{Uber} \\ 0 & \text{if } m_3 = \text{car} \end{cases} \quad (16)$$

For travelers who choose transit for the return journey, the egress mode probabilities are

$$P_3(a_3 \mid m_3 = \text{transit}, d, \eta) = \frac{\exp(\alpha_{\text{transit}} - f_{a_3} - d \cdot c_{a_3} + \eta_{\text{transit}})}{\sum_{a' \in \{\text{walk}, \text{Uber}\}} \exp(\alpha_{\text{transit}} - f_{a'} - d \cdot c_{a'} + \eta_{\text{transit}})} \quad (17)$$

C.1.3 Unconditional choice probabilities

The unconditional choice probabilities can be found by integrating over possible realizations of η_m . Let $F(\eta)$ denote the joint distribution over all preference shocks η_m . Then the unconditional choice probabilities for the outbound trip and access modes are

$$P_1(m_1 \mid d) = \int P_1(m_1 \mid d, \eta) dF(\eta), \text{ and} \\ P_1(a_1 \mid m_1 = \text{transit}, d) = \int P_1(a_1 \mid m_1 = \text{transit}, d, \eta) dF(\eta).$$

Similarly, the unconditional choice probabilities for the return trip and egress modes are

$$P_3(m_3 \mid m_1, d) = \int P_3(m_3 \mid m_1, d, \eta) dF(\eta), \text{ and} \\ P_3(a_3 \mid m_3 = \text{transit}, d) = \int P_3(a_3 \mid m_3 = \text{transit}, d, \eta) dF(\eta).$$

C.2 Theoretical results

This subsection restates and proves the proposition and corollary from the main text.

Proposition 1 (Three-channel decomposition (restated)) *For each period $t \in \{1, 2, 3\}$, let $P_t(\text{Uber}_{\text{total}} | d, \eta)$ denote the probability that the trip uses Uber either as the primary mode or as a transit access or egress mode, conditional on (d, η) . Its derivative with respect to α_{transit} can be written as*

$$\frac{\partial P_t(\text{Uber}_{\text{total}} | d, \eta)}{\partial \alpha_{\text{transit}}} = \underbrace{S_t(d, \eta)}_{\text{Direct substitution} \leq 0} + \underbrace{D_t(d, \eta)}_{\text{Dynamic choice-set effect} \geq 0} + \underbrace{L_t(d, \eta)}_{\text{Last-mile complementarity} \geq 0}.$$

For $t \in \{1, 3\}$, there are parameter values for which this derivative is strictly negative and others for which it is strictly positive. For period 2, the probability is conditional on the intermediate trip occurring and $L_2 \equiv 0$ because that trip has no access mode. Because $S_2 = 0$, the period-2 derivative reduces to $D_2 \geq 0$.

Proof of Proposition 1. Throughout, we condition on (d, η) and suppress that conditioning to simplify the notation. Define

$$\Phi(d, \eta) \equiv P_3(\text{transit} | m_1 \neq \text{car}) \geq 0.$$

Outbound trip (period 1). The pairwise logit derivatives established in the main-text proof sketch give

$$\begin{aligned} \frac{\partial P_1(\text{Uber})}{\partial \alpha_{\text{transit}}} &= -P_1(\text{Uber})P_1(\text{transit}) + P_1(\text{Uber})P_1(\text{car})\Phi, \\ \frac{\partial P_1(\text{transit})}{\partial \alpha_{\text{transit}}} &= P_1(\text{transit})P_1(\text{Uber}) + P_1(\text{transit})P_1(\text{car})(1 + \Phi). \end{aligned}$$

Because the conditional access-mode probability is unaffected by α_{transit} , combining these derivatives gives the full outbound decomposition:

$$\begin{aligned} \frac{\partial P_1(\text{Uber}_{\text{total}})}{\partial \alpha_{\text{transit}}} &= \underbrace{-P_1(\text{Uber})P_1(\text{transit})}_{S_1 \leq 0} + \underbrace{P_1(\text{Uber})P_1(\text{car})\Phi}_{D_1 \geq 0} \\ &\quad + \underbrace{P_1(a_1 = \text{Uber} | m_1 = \text{transit})P_1(\text{transit}) \times [P_1(\text{Uber}) + P_1(\text{car})(1 + \Phi)]}_{L_1 \geq 0}. \quad (18) \end{aligned}$$

The first term, S_1 , is the loss of direct Uber trips when travelers substitute toward transit. The second term, D_1 , is the gain in direct Uber trips relative to car because not driving now gives travelers access to improved non-car choices later. The final term, L_1 , is the increased use of Uber for access trips. We highlight again that this final category occurs because *transit ridership* increases due to substitution and dynamic-choice set effects, which then leads to more first- and last-mile trips.

Before we continue with the proof for periods 2 and 3, it is useful to note that

$$\frac{\partial P_1(\text{car})}{\partial \alpha_{\text{transit}}} = -P_1(\text{car}) [P_1(\text{transit}) + (1 - P_1(\text{car}))\Phi] \leq 0. \quad (19)$$

Return trip (period 3). For the return trip, the same three channels operate, but the dynamic choice-set effect is now backward-looking. Only travelers who did not drive outbound can choose Uber, so total return-trip Uber use is

$$P_3(\text{Uber}_{\text{total}}) = (1 - P_1(\text{car})) \left[P_3(\text{Uber} \mid m_1 \neq \text{car}) + P_3(a_3 = \text{Uber} \mid m_3 = \text{transit}) P_3(\text{transit} \mid m_1 \neq \text{car}) \right]. \quad (20)$$

Among travelers already in the non-car state, a transit improvement shifts primary-mode probability from Uber to transit:

$$\begin{aligned} \frac{\partial P_3(\text{Uber} \mid m_1 \neq \text{car})}{\partial \alpha_{\text{transit}}} &= -P_3(\text{Uber} \mid m_1 \neq \text{car}) P_3(\text{transit} \mid m_1 \neq \text{car}), \\ \frac{\partial P_3(\text{transit} \mid m_1 \neq \text{car})}{\partial \alpha_{\text{transit}}} &= -\frac{\partial P_3(\text{Uber} \mid m_1 \neq \text{car})}{\partial \alpha_{\text{transit}}}. \end{aligned}$$

The egress probability is fixed for the same reason as the outbound access probability. However, increasing transit quality increases the mass of travelers who can choose Uber or transit on the return trip. Differentiating (20) yields

$$\frac{\partial P_3(\text{Uber}_{\text{total}})}{\partial \alpha_{\text{transit}}} = S_3 + D_3 + L_3, \quad (21)$$

where

$$\begin{aligned} S_3 &\equiv (1 - P_1(\text{car})) \frac{\partial P_3(\text{Uber} \mid m_1 \neq \text{car})}{\partial \alpha_{\text{transit}}} \leq 0, \\ D_3 &\equiv \frac{\partial (1 - P_1(\text{car}))}{\partial \alpha_{\text{transit}}} P_3(\text{Uber} \mid m_1 \neq \text{car}) \geq 0, \\ L_3 &\equiv P_3(a_3 = \text{Uber} \mid m_3 = \text{transit}) \left[\frac{\partial (1 - P_1(\text{car}))}{\partial \alpha_{\text{transit}}} P_3(\text{transit} \mid m_1 \neq \text{car}) \right. \\ &\quad \left. + (1 - P_1(\text{car})) \frac{\partial P_3(\text{transit} \mid m_1 \neq \text{car})}{\partial \alpha_{\text{transit}}} \right] \geq 0. \end{aligned}$$

The first term is substitution from direct Uber to transit. The second is the direct Uber use of travelers who can now choose a non-car return mode because they did not drive outbound. The final term captures Uber egress rides generated both when existing non-car travelers switch from Uber to transit and when travelers newly induced not to drive outbound choose transit.

Intermediate trip (period 2). For the intermediate trip, there are only two channels because there is no access mode. Let $B \in \{m_1 = \text{car}, m_1 \neq \text{car}\}$ index the two branches determined by the outbound choice. Conditional on the trip occurring, the law of total probability gives

$$P_2(\text{Uber}_{\text{total}}) = \sum_B \Pr(B) P_2(\text{Uber} \mid B). \quad (22)$$

Within either branch, a transit improvement does not shift the probability from Uber to transit because there is no change in the within period 2 relative utilities.

$$\frac{\partial P_2(\text{Uber} \mid B)}{\partial \alpha_{\text{transit}}} = 0.$$

Differentiating (22) yields

$$\frac{\partial P_2(\text{Uber}_{\text{total}})}{\partial \alpha_{\text{transit}}} = S_2 + D_2 + L_2, \quad (23)$$

where

$$\begin{aligned} S_2 &\equiv 0, \\ D_2 &\equiv \sum_B \frac{\partial \Pr(B)}{\partial \alpha_{\text{transit}}} P_2(\text{Uber} \mid B) \\ &= \frac{\partial P_1(\text{car})}{\partial \alpha_{\text{transit}}} P_2(\text{Uber} \mid m_1 = \text{car}) + \frac{\partial [1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_2(\text{Uber} \mid m_1 \neq \text{car}) \\ &= \underbrace{\frac{\partial [1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}}}_{\geq 0} \underbrace{[P_2(\text{Uber} \mid m_1 \neq \text{car}) - P_2(\text{Uber} \mid m_1 = \text{car})]}_{\geq 0} \\ &\geq 0, \\ L_2 &\equiv 0. \end{aligned}$$

The first term, S_2 , is zero because there is no substitution from Uber to transit within the two period-2 choice sets. The second, D_2 , is the dynamic choice-set effect, now backward-looking. Increasing transit quality reduces the probability of using car for the outbound journey, changing the choice set for the intermediate trip. Because removing car from the intermediate trip raises Uber's conditional choice probability, this shift increases Uber use for the intermediate trip. The final term, L_2 , is zero because the intermediate trip has no access mode.

Expected Uber rides before the traveler learns whether an intermediate trip occurs equal $p_s P_2(\text{Uber}_{\text{total}})$; differentiating this object multiplies every term in (23) by p_s .

Equality across periods 1 and 3. Next we show that $P_1(\text{Uber}_{\text{total}}) = P_3(\text{Uber}_{\text{total}})$. The model assigns the outbound and return trips the same primary-mode utilities, access or egress costs, and distributions of transitory shocks.

Because car, Uber, and transit exhaust the period-1 primary modes,

$$1 - P_1(\text{car}) = P_1(\text{Uber}) + P_1(\text{transit}).$$

The two ratios below are therefore the period-1 probabilities of Uber and transit conditional on not choosing car. Their relative odds are

$$\begin{aligned} \frac{P_1(\text{Uber})}{P_1(\text{transit})} &= \exp[V_1(\text{Uber}) - V_1(\text{transit})] \\ &= \exp\left[\alpha_{\text{Uber}} + \eta_{\text{Uber}} - IV^{\text{transit}}(d, \eta) \right. \\ &\quad \left. + V_2(\text{Uber}) - V_2(\text{transit}) + V_3(\text{Uber}, d, \eta) - V_3(\text{transit}, d, \eta)\right] \\ &= \exp\left[\alpha_{\text{Uber}} + \eta_{\text{Uber}} - IV^{\text{transit}}(d, \eta)\right]. \end{aligned}$$

The final equality uses $V_2(\text{Uber}) = V_2(\text{transit})$ and $V_3(\text{Uber}, d, \eta) = V_3(\text{transit}, d, \eta)$: choosing either Uber or transit outbound produces the same later choice sets. Conditional on not driving outbound, the period-3 Uber-to-transit odds are the same,

$$\frac{P_3(\text{Uber} \mid m_1 \neq \text{car})}{P_3(\text{transit} \mid m_1 \neq \text{car})} = \exp\left[\alpha_{\text{Uber}} + \eta_{\text{Uber}} - IV^{\text{transit}}(d, \eta)\right].$$

Both pairs of conditional probabilities sum to one. Equal odds therefore imply

$$\begin{aligned} \frac{P_1(\text{Uber})}{1 - P_1(\text{car})} &= P_3(\text{Uber} \mid m_1 \neq \text{car}), \\ \frac{P_1(\text{transit})}{1 - P_1(\text{car})} &= P_3(\text{transit} \mid m_1 \neq \text{car}). \end{aligned} \tag{24}$$

The outbound access probability also equals the return egress probability. Combining this equality with (24) gives

$$\begin{aligned} P_1(\text{Uber}_{\text{total}}) &= P_1(\text{Uber}) + P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})P_1(\text{transit}) \\ &= (1 - P_1(\text{car})) \left[P_3(\text{Uber} \mid m_1 \neq \text{car}) \right. \\ &\quad \left. + P_3(a_3 = \text{Uber} \mid m_3 = \text{transit})P_3(\text{transit} \mid m_1 \neq \text{car}) \right] \\ &= P_3(\text{Uber}_{\text{total}}). \end{aligned} \tag{25}$$

Thus the total-Uber probabilities, and hence their derivatives, are identical in periods 1 and 3. Their decompositions nevertheless describe the dynamic channel from different temporal perspectives: it is forward-looking in period 1 and backward-looking in period 3. It therefore suffices to construct both signs for period 1.

Improving transit can increase or decrease Uber usage. Finally, we show that $\partial P_t / \partial \alpha_{\text{transit}}$ can be positive or negative for $t \in \{1, 3\}$. Recall that $\Phi \geq 0$ and fix (d, η) .

As $\alpha_{\text{car}} \rightarrow -\infty$, $P_1(\text{car}) \rightarrow 0$, while the other period-1 probabilities remain strictly positive. Therefore, (18) implies

$$\begin{aligned} \lim_{\alpha_{\text{car}} \rightarrow -\infty} \frac{\partial P_1(\text{Uber}_{\text{total}})}{\partial \alpha_{\text{transit}}} &= -P_1(\text{Uber})P_1(\text{transit}) + \underbrace{P_1(\text{Uber})P_1(\text{car})\Phi}_{\rightarrow 0} \\ &\quad + P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})P_1(\text{transit}) \left[P_1(\text{Uber}) + \underbrace{P_1(\text{car})(1 + \Phi)}_{\rightarrow 0} \right] \\ &= -P_1(\text{Uber})P_1(\text{transit}) + P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})P_1(\text{transit})P_1(\text{Uber}) \\ &= -[1 - P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})] \times P_1(\text{Uber})P_1(\text{transit}) < 0. \end{aligned}$$

As $\alpha_{\text{Uber}} \rightarrow -\infty$, $P_1(\text{Uber}) \rightarrow 0$, while $P_1(\text{car})$ and $P_1(\text{transit})$ remain strictly positive. Therefore, (18) implies

$$\begin{aligned} \lim_{\alpha_{\text{Uber}} \rightarrow -\infty} \frac{\partial P_1(\text{Uber}_{\text{total}})}{\partial \alpha_{\text{transit}}} &= \underbrace{-P_1(\text{Uber})P_1(\text{transit})}_{\rightarrow 0} + \underbrace{P_1(\text{Uber})P_1(\text{car})\Phi}_{\rightarrow 0} \\ &\quad + P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})P_1(\text{transit}) \left[\underbrace{P_1(\text{Uber})}_{\rightarrow 0} + P_1(\text{car})(1 + \Phi) \right] \\ &= P_1(a_1 = \text{Uber} \mid m_1 = \text{transit})P_1(\text{transit})P_1(\text{car})(1 + \Phi) > 0. \end{aligned}$$

Finally, the preceding results condition on (d, η) . Let $\bar{P}_t(\text{Uber}_{\text{total}} \mid d)$ denote the corresponding probability after integrating over the persistent shocks. Because $F(\eta)$ is independent of α_{transit} , the logit probabilities and their derivatives are bounded, and differentiation may be passed under the integral. Hence the unconditional derivative has the same decomposition, with each component integrated over $F(\eta)$, and the component-wise signs are preserved. For period 2, ex ante expected Uber use and its derivative are multiplied by p_s . ■

Corollary 1 (Mapping mode switches to mechanisms (restated)) *Conditional on (d, η) , the marginal response of mode probabilities to an increase in α_{transit} maps to the three channels in the proposition above as follows:*

1. *From Uber to transit-by-walk: direct substitution only.*
2. *From car to transit-by-Uber: last-mile complementarity only.*
3. *From Uber to transit-by-Uber: direct substitution and last-mile complementarity.*
4. *From car to Uber: dynamic choice-set effect only.*

In the intermediate trip (period 2), there is no Uber access leg, so clauses 2 and 3 do not arise. Because the station opening does not directly change intermediate-trip flow utilities, clause 1 does not arise either; only the car-to-Uber margin in clause 4 remains.

Proof of Corollary 1. Each channel operates through a distinct part of the model: direct substitution through transit’s period utility, the dynamic choice-set effect through the constraint the outbound choice places on later choice sets, and last-mile complementarity through the Uber access and egress rides attached to transit trips.³¹ Each clause of the corollary asserts not only that certain channels operate on a margin, but also that the others do not. The proof of the proposition shows which channels are present; here we show that, on each margin, the remaining pathways are closed. Three observations, each established in the proof of the proposition, do this work:

- (i) Choosing Uber or transit leaves the traveler with the same choice sets in periods 2 and 3, so the two modes’ continuation values respond identically to α_{transit} and cancel from the Uber–transit margin.
- (ii) Neither car’s nor Uber’s period utility depends directly on α_{transit} .
- (iii) The access-mode share is unaffected by α_{transit} , so Uber access rides change only in proportion to transit ridership.

The Uber-to-transit-by-walk margin (clause 1) carries only direct substitution. By (i), the continuation values cancel, so the margin moves only through period-1 utility, where $\partial(V_1(\text{Uber}) - V_1(\text{transit}))/\partial\alpha_{\text{transit}} = -1$. Because the traveler walks to transit, her switch generates no Uber access rides, so last-mile complementarity is absent as well.

The car-to-Uber margin (clause 4) carries only the dynamic choice-set effect. By (ii), the margin moves only through continuation values, where $\partial(V_1(\text{Uber}) - V_1(\text{car}))/\partial\alpha_{\text{transit}} = \Phi$. Because neither endpoint is a transit trip, the switch generates no Uber access rides.

The two margins terminating at transit-by-Uber (clauses 2 and 3) are the only ones that carry last-mile complementarity. By (iii), every flow into transit generates Uber access rides in proportion to the fixed access share, and no other flow generates any. A traveler who switches from Uber (clause 3) also gives up a direct Uber trip; this loss is the same direct substitution as in clause 1, so the margin carries both channels. A traveler who switches from car (clause 2) gives up no direct Uber trip. Her switch to transit reflects both transit’s gain in period-1 utility and its gain in continuation value (the $1 + \Phi$ in the proof of the proposition), but the channels partition the change in *Uber* use, and the only Uber rides her switch creates are access rides. We therefore count this margin as last-mile complementarity only.

The remaining margin, from car to transit-by-walk, involves no Uber use at either endpoint and generates no access rides, so the list is exhaustive.

³¹Recall that *transit* usage increases due to substitution and dynamic choice-set effects, which then increases last-mile trips. Our decomposition is focused on the impact of a transit improvement on Uber, and so we call all these trips “last-mile complementarity.”

The return trip is symmetric, with the dynamic choice-set effect operating through the realized outbound choice rather than continuation values.

For the intermediate trip there is no access mode, so clauses 2 and 3 do not arise. Because there is no within-choice-set substitution toward transit, clause 1 also does not arise. Clause 4 arises from the change in choice-set composition: improved transit causes fewer travelers to drive outbound, and these travelers are more likely to choose Uber now that using their car is not in their choice set. ■

Proposition 2 Assume $\alpha_m = \tilde{\alpha}_m - \rho(d, \eta) \cdot \pi_m$ and $\beta_m = \tilde{\beta}_m - \rho(d, \eta) \cdot \pi_m$, for average mode prices π_m and for a traveler at distance d and with preferences η . Let $\varepsilon^D(d, \eta)$ be the elasticity of demand for Uber trips with respect to Uber price π_{Uber} . Let $Q(\text{Uber}|d, \eta)$ be total expected direct Uber trips over the outbound, intermediate, and return trips. Then $\rho(d, \eta)$ can be expressed as

$$\rho(d, \eta) = -\frac{\varepsilon^D(d, \eta)}{\pi_{\text{Uber}}} \cdot \frac{Q(\text{Uber}|d, \eta)}{Z(d, \eta)}$$

where $Z(d, \eta)$ is a function of model probabilities given in the proof.

Proof of Proposition 2. Expected direct Uber trips, summed over the outbound, intermediate, and return legs, are

$$\begin{aligned} Q(\text{Uber}|d, \eta) &= P_1(\text{car}|d, \eta) p_s P_2(\text{Uber}|m_1 = \text{car}, d, \eta) \\ &\quad + P_1(\text{transit}|d, \eta) \left[p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) \right] \\ &\quad + P_1(\text{Uber}|d, \eta) \left[1 + p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) \right]. \end{aligned}$$

A traveler who drives outbound still faces an Uber option on the intermediate trip, so the continuation value of the car branch responds to π_{Uber} :

$$\frac{\partial V_1(\text{car}|d, \eta)}{\partial \pi_{\text{Uber}}} = -\rho(d, \eta) p_s P_2(\text{Uber}|m_1 = \text{car}, d, \eta).$$

By the same logic, the continuation values of the transit and Uber outbound branches are

$$\begin{aligned} \frac{\partial V_1(\text{transit}|d, \eta)}{\partial \pi_{\text{Uber}}} &= -\rho(d, \eta) \left[p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) \right], \\ \frac{\partial V_1(\text{Uber}|d, \eta)}{\partial \pi_{\text{Uber}}} &= -\rho(d, \eta) \left[1 + p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) \right], \end{aligned}$$

the leading 1 in the last line reflecting that choosing Uber outbound means paying π_{Uber} directly. Taking derivatives over the choice probabilities yields

$$\frac{\partial P_1(m_1|d, \eta)}{\partial \pi_{\text{Uber}}} = P_1(m_1|d, \eta) \left[\frac{\partial V_1(m_1|d, \eta)}{\partial \pi_{\text{Uber}}} - \sum_{m'_1} P_1(m'_1|d, \eta) \frac{\partial V_1(m'_1|d, \eta)}{\partial \pi_{\text{Uber}}} \right].$$

The bracketed sum is $\sum_{m'_1} P_1(m'_1|d, \eta) \partial V_1(m'_1|d, \eta) / \partial \pi_{\text{Uber}} = -\rho(d, \eta) Q(\text{Uber}|d, \eta)$, since it is exactly the branch-probability-weighted sum of the three continuation-value derivatives above, matching $Q(\text{Uber}|d, \eta)$ term by term. Substituting,

$$\frac{\partial P_1(\text{car}|d, \eta)}{\partial \pi_{\text{Uber}}} = \rho(d, \eta) P_1(\text{car}|d, \eta) \left[Q(\text{Uber}|d, \eta) - p_s P_2(\text{Uber}|m_1 = \text{car}, d, \eta) \right],$$

with the analogous expressions for the transit and Uber branches.

Within each outbound branch, the intermediate and return mode choices are themselves logit choices with direct own-price sensitivity:

$$\frac{\partial P_2(\text{Uber}|m_1, d, \eta)}{\partial \pi_{\text{Uber}}} = -\rho(d, \eta) P_2(\text{Uber}|m_1, d, \eta) [1 - P_2(\text{Uber}|m_1, d, \eta)],$$

$$\frac{\partial P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta)}{\partial \pi_{\text{Uber}}} = -\rho(d, \eta) P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) [1 - P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta)].$$

Applying the product rule to each of the three terms in $Q(\text{Uber}|d, \eta)$ and collecting terms,

$$\frac{\partial Q(\text{Uber}|d, \eta)}{\partial \pi_{\text{Uber}}} = -\rho(d, \eta) Z(d, \eta),$$

where

$$\begin{aligned} Z(d, \eta) = & P_1(\text{car}|d, \eta) \left[p_s P_2(\text{Uber}|m_1 = \text{car}, d, \eta) - Q(\text{Uber}|d, \eta) \right]^2 \\ & + P_1(\text{transit}|d, \eta) \left[p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) - Q(\text{Uber}|d, \eta) \right]^2 \\ & + P_1(\text{Uber}|d, \eta) \left[1 + p_s P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) + P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) - Q(\text{Uber}|d, \eta) \right]^2 \\ & + p_s P_1(\text{car}|d, \eta) P_2(\text{Uber}|m_1 = \text{car}, d, \eta) [1 - P_2(\text{Uber}|m_1 = \text{car}, d, \eta)] \\ & + p_s [1 - P_1(\text{car}|d, \eta)] P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta) [1 - P_2(\text{Uber}|m_1 \neq \text{car}, d, \eta)] \\ & + [1 - P_1(\text{car}|d, \eta)] P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta) [1 - P_3(\text{Uber}|m_1 \neq \text{car}, d, \eta)]. \end{aligned}$$

The first three terms capture substitution among outbound alternatives; the last three capture substitution within the intermediate and return choice sets. Every term is nonnegative, so $Z(d, \eta) \geq 0$, with strict inequality except in degenerate cases. Consequently

$$\rho(d, \eta) = -\frac{\partial Q(\text{Uber}|d, \eta)}{\partial \pi_{\text{Uber}}} \cdot \frac{1}{Z(d, \eta)} = -\frac{\epsilon^D(d, \eta)}{\pi_{\text{Uber}}} \cdot \frac{Q(\text{Uber}|d, \eta)}{Z(d, \eta)}.$$

■

Corollary 2 Assume the elasticity of demand for Uber trips is common across travelers, $\epsilon^D(d, \eta) = \epsilon^D$ for every (d, η) . Let (d, η) be distributed according to $G(d)F(\eta)$ in the population, and define the average

price coefficient

$$\bar{\rho} \equiv \iint \rho(d, \eta) dG(d) dF(\eta).$$

Then

$$\bar{\rho} = -\frac{\varepsilon^D}{\pi_{\text{Uber}}} \iint \frac{Q(\text{Uber}|d, \eta)}{Z(d, \eta)} dG(d) dF(\eta).$$

Proof of Corollary 2. Under the assumption that $\varepsilon^D(d, \eta) = \varepsilon^D$ for every (d, η) , Proposition 2 gives, for every (d, η) ,

$$\rho(d, \eta) = -\frac{\varepsilon^D}{\pi_{\text{Uber}}} \cdot \frac{Q(\text{Uber}|d, \eta)}{Z(d, \eta)}.$$

Integrating both sides over the population distribution of (d, η) and applying linearity of the integral,

$$\bar{\rho} = \iint \rho(d, \eta) dG(d) dF(\eta) = -\frac{\varepsilon^D}{\pi_{\text{Uber}}} \iint \frac{Q(\text{Uber}|d, \eta)}{Z(d, \eta)} dG(d) dF(\eta),$$

as claimed. ■

Remark 1 Corollary 2 identifies $\bar{\rho}$, the simple average of the price coefficient across travelers, from a single external estimate of the common elasticity ε^D , model-implied choice probabilities, and the price π_{Uber} . We use $\hat{\rho}$, the sample analogue of $\bar{\rho}$, as the price coefficient throughout the willingness-to-pay calculations and counterfactuals that follow.

C.3 Mode-by-mechanism decomposition behind Figure 10

This subsection derives the transit and car responses that accompany the three Uber channels in Proposition 1, and collects them into the mode-by-mechanism accounting used in Figure 10. The proof of Proposition 1 in Appendix C.2 already derives the Uber terms for each period and the outbound car response; the remaining steps are to derive the transit and car responses for the return and intermediate trips, add the three periods, and verify that each mechanism reallocates primary-mode probability without creating or destroying it. Appendix E then explains how the resulting derivatives are integrated over the estimated model to produce the percentages reported in the figure.

Uber derivatives already defined in the proposition. Continue to condition on (d, η) and suppress that conditioning, as in the proof. The proof defines

$$\Phi = P_3(\text{transit} \mid m_1 \neq \text{car}),$$

as well as the period-specific direct-substitution, option-value, and last-mile terms S_t , D_t , and L_t . No additional notation is needed for the three whole-day Uber contributions:

$$\begin{aligned} S_1 + S_3 = & -P_1(\text{Uber})P_1(\text{transit}) \\ & - \{1 - P_1(\text{car})\}P_3(\text{Uber} \mid m_1 \neq \text{car})P_3(\text{transit} \mid m_1 \neq \text{car}), \end{aligned} \quad (26)$$

$$\begin{aligned} D_1 + p_s D_2 + D_3 = & P_1(\text{car})P_1(\text{Uber})\Phi \\ & + \frac{\partial[1 - P_1(\text{car})]}{\partial\alpha_{\text{transit}}} \left[P_3(\text{Uber} \mid m_1 \neq \text{car}) \right. \\ & \left. + p_s \{P_2(\text{Uber} \mid m_1 \neq \text{car}) - P_2(\text{Uber} \mid m_1 = \text{car})\} \right], \end{aligned} \quad (27)$$

$$\begin{aligned} L_1 + L_3 = & P_1(a_1 = \text{Uber} \mid m_1 = \text{transit}) \\ & \times \frac{\partial P_1(\text{transit})}{\partial\alpha_{\text{transit}}} \\ & + P_3(a_3 = \text{Uber} \mid m_3 = \text{transit}) \\ & \times \frac{\partial[\{1 - P_1(\text{car})\}P_3(\text{transit} \mid m_1 \neq \text{car})]}{\partial\alpha_{\text{transit}}}. \end{aligned} \quad (28)$$

There is no S_2 because transit quality does not change relative utilities within either intermediate-trip choice set, and there is no L_2 because the intermediate trip has no access or egress leg. The factor p_s accounts for the probability that the intermediate trip occurs.

Adding the three expressions gives the derivative of total Uber ridership:

$$\frac{\partial \text{total Uber ridership}}{\partial\alpha_{\text{transit}}} = (S_1 + S_3) + (D_1 + p_s D_2 + D_3) + (L_1 + L_3). \quad (29)$$

Primary-mode derivatives, period by period. We next derive the primary-mode responses for transit and car separately in each period, displaying the direct-Uber derivative alongside them in period 1 to make the pairwise flows transparent. As in the proof, all probabilities remain conditional on (d, η) . Within-period flows caused directly by transit's higher utility are classified as direct substitution. Later-period flows caused by the outbound choice changing the traveler's available modes are classified as option value. Last-mile complementarity affects the number of Uber access and egress rides, not which primary mode is used.

Outbound trip (period 1). The proof sketch in Section 2 writes the Uber and transit derivatives as sums of pairwise logit margins. Applying the same identity to car puts all three primary-mode

derivatives in a common form:

$$\begin{aligned}
\frac{\partial P_1(\text{Uber})}{\partial \alpha_{\text{transit}}} &= P_1(\text{Uber})P_1(\text{transit}) \frac{\partial [V_1(\text{Uber}) - V_1(\text{transit})]}{\partial \alpha_{\text{transit}}} \\
&\quad + P_1(\text{Uber})P_1(\text{car}) \frac{\partial [V_1(\text{Uber}) - V_1(\text{car})]}{\partial \alpha_{\text{transit}}} \\
&= \underbrace{-P_1(\text{Uber})P_1(\text{transit})}_{\text{direct substitution to transit}} + \underbrace{P_1(\text{Uber})P_1(\text{car})\Phi}_{\text{option value from car}},
\end{aligned} \tag{30}$$

$$\begin{aligned}
\frac{\partial P_1(\text{transit})}{\partial \alpha_{\text{transit}}} &= P_1(\text{transit})P_1(\text{Uber}) \frac{\partial [V_1(\text{transit}) - V_1(\text{Uber})]}{\partial \alpha_{\text{transit}}} \\
&\quad + P_1(\text{transit})P_1(\text{car}) \frac{\partial [V_1(\text{transit}) - V_1(\text{car})]}{\partial \alpha_{\text{transit}}} \\
&= \underbrace{P_1(\text{transit})P_1(\text{Uber})}_{\text{direct substitution from Uber}} + \underbrace{P_1(\text{transit})P_1(\text{car})}_{\text{direct substitution from car}} \\
&\quad + \underbrace{P_1(\text{transit})P_1(\text{car})\Phi}_{\text{option value from car}},
\end{aligned} \tag{31}$$

$$\begin{aligned}
\frac{\partial P_1(\text{car})}{\partial \alpha_{\text{transit}}} &= P_1(\text{car})P_1(\text{Uber}) \frac{\partial [V_1(\text{car}) - V_1(\text{Uber})]}{\partial \alpha_{\text{transit}}} \\
&\quad + P_1(\text{car})P_1(\text{transit}) \frac{\partial [V_1(\text{car}) - V_1(\text{transit})]}{\partial \alpha_{\text{transit}}} \\
&= \underbrace{-P_1(\text{car})P_1(\text{transit})}_{\text{direct substitution to transit}} + \underbrace{-P_1(\text{car})P_1(\text{Uber})\Phi}_{\text{option value to Uber}} \\
&\quad + \underbrace{-P_1(\text{car})P_1(\text{transit})\Phi}_{\text{option value to transit}}.
\end{aligned} \tag{32}$$

The second equality in each case substitutes the relative-value derivatives from the proof sketch in Section 2:

$$\begin{aligned}
\frac{\partial [V_1(\text{Uber}) - V_1(\text{transit})]}{\partial \alpha_{\text{transit}}} &= -1, \\
\frac{\partial [V_1(\text{Uber}) - V_1(\text{car})]}{\partial \alpha_{\text{transit}}} &= \Phi, \\
\frac{\partial [V_1(\text{transit}) - V_1(\text{car})]}{\partial \alpha_{\text{transit}}} &= 1 + \Phi.
\end{aligned}$$

The reverse-order derivatives carry the opposite signs. The labeled terms then display the same pairwise flows once from the gaining mode's perspective and once from the losing mode's perspective.

Period-1 accounting check. Because Uber, transit, and car exhaust the period-1 primary-mode choice

set,

$$P_1(\text{Uber}) + P_1(\text{transit}) + P_1(\text{car}) = 1.$$

Differentiating this identity shows formally that equations (30)–(32) must sum to zero:

$$\frac{\partial P_1(\text{Uber})}{\partial \alpha_{\text{transit}}} + \frac{\partial P_1(\text{transit})}{\partial \alpha_{\text{transit}}} + \frac{\partial P_1(\text{car})}{\partial \alpha_{\text{transit}}} = 0. \quad (33)$$

The simplified expressions verify the identity term by term: the direct Uber–transit and car–transit flows each enter once positively and once negatively, as do the two option-value flows out of car.

Return trip (period 3). Travelers who drove outbound must return by car. Travelers who did not drive choose between Uber and transit. Thus

$$\begin{aligned} P_3(\text{transit}) &= \{1 - P_1(\text{car})\} P_3(\text{transit} \mid m_1 \neq \text{car}), \\ P_3(\text{car}) &= P_1(\text{car}). \end{aligned}$$

Differentiating gives

$$\begin{aligned} \frac{\partial P_3(\text{transit})}{\partial \alpha_{\text{transit}}} &= \underbrace{\{1 - P_1(\text{car})\} P_3(\text{Uber} \mid m_1 \neq \text{car}) P_3(\text{transit} \mid m_1 \neq \text{car})}_{\text{direct substitution} = -S_3} \\ &\quad + \underbrace{\frac{\partial [1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_3(\text{transit} \mid m_1 \neq \text{car})}_{\text{option value}}, \end{aligned} \quad (34)$$

$$\begin{aligned} \frac{\partial P_3(\text{car})}{\partial \alpha_{\text{transit}}} &= \underbrace{0}_{\text{direct substitution}} - \underbrace{\frac{\partial [1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}}}_{\text{option value}}. \end{aligned} \quad (35)$$

The direct-substitution term moves existing non-car return travelers from Uber to transit. The option-value terms arise because the transit improvement moves more outbound travelers into the non-car state, making Uber and transit—and not car—available for their return trip. This is why the entire change in return-trip car use is classified as option value even though the change in outbound non-car mass itself reflects both terms in equation (32).

Intermediate trip (period 2). The traveler has access to car on the intermediate trip only after driving outbound. The unconditional primary-mode probabilities, conditional on the intermediate trip occurring, are therefore

$$\begin{aligned} P_2(\text{transit}) &= P_1(\text{car}) P_2(\text{transit} \mid m_1 = \text{car}) \\ &\quad + \{1 - P_1(\text{car})\} P_2(\text{transit} \mid m_1 \neq \text{car}), \\ P_2(\text{car}) &= P_1(\text{car}) P_2(\text{car} \mid m_1 = \text{car}). \end{aligned}$$

The conditional period-2 shares do not depend on α_{transit} . Hence

$$\frac{\partial P_2(\text{transit})}{\partial \alpha_{\text{transit}}} = \underbrace{0}_{\text{direct substitution}} + \underbrace{\frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} [P_2(\text{transit} \mid m_1 \neq \text{car}) - P_2(\text{transit} \mid m_1 = \text{car})]}_{\text{option value}}, \quad (36)$$

$$\frac{\partial P_2(\text{car})}{\partial \alpha_{\text{transit}}} = \underbrace{0}_{\text{direct substitution}} - \underbrace{\frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_2(\text{car} \mid m_1 = \text{car})}_{\text{option value}}. \quad (37)$$

Both changes arise entirely from the shift in outbound travelers from the car branch to the non-car branch. When forming expected daily usage, multiply both period-2 derivatives by p_s .

Adding the three periods. Adding the period-specific terms yields the direct-substitution contribution to transit usage,

$$\frac{\partial \text{transit usage} \mid \text{direct substitution}}{\partial \alpha_{\text{transit}}} = -(S_1 + S_3) + P_1(\text{car})P_1(\text{transit}), \quad (38)$$

and its option-value contribution,

$$\begin{aligned} \frac{\partial \text{transit usage} \mid \text{option value}}{\partial \alpha_{\text{transit}}} &= P_1(\text{car})P_1(\text{transit})\Phi \\ &+ \frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_3(\text{transit} \mid m_1 \neq \text{car}) \\ &+ p_s \frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_2(\text{transit} \mid m_1 \neq \text{car}) \\ &- p_s \frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} P_2(\text{transit} \mid m_1 = \text{car}). \end{aligned} \quad (39)$$

These are simply the direct-substitution and option-value terms, respectively, from equations (31), (34), and (36).

For car usage, the corresponding sums are

$$\frac{\partial \text{car usage} \mid \text{direct substitution}}{\partial \alpha_{\text{transit}}} = -P_1(\text{car})P_1(\text{transit}), \quad (40)$$

$$\begin{aligned} \frac{\partial \text{car usage} \mid \text{option value}}{\partial \alpha_{\text{transit}}} &= -P_1(\text{car})\{1 - P_1(\text{car})\}\Phi \\ &\quad - \frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} [1 + p_s P_2(\text{car} \mid m_1 = \text{car})] \\ &= -(D_1 + p_s D_2 + D_3) \\ &\quad - \frac{\partial \text{transit usage} \mid \text{option value}}{\partial \alpha_{\text{transit}}}. \end{aligned} \quad (41)$$

The first equality displays the period-1, period-3, and expected period-2 car losses directly. The second writes the same quantity as the negative of the corresponding Uber and transit gains.

Last-mile complementarity contributes $L_1 + L_3$ to Uber ridership and zero to primary transit and car usage.

Complete mode-by-mechanism accounting. Every entry in the following table is a derivative with respect to α_{transit} . The Uber column counts all Uber rides; the transit and car columns count primary-mode usage.

Mechanism	Uber ridership	Transit usage	Car usage
Direct substitution	$S_1 + S_3$	equation (38)	equation (40)
Option value	$D_1 + p_s D_2 + D_3$	equation (39)	equation (41)
Last-mile complementarity	$L_1 + L_3$	0	0
Total	sum of the three entries above	sum of the first two entries above	sum of the first two entries above

The accounting can also be checked period by period. In period 1, the direct-substitution flows satisfy

$$-P_1(\text{Uber})P_1(\text{transit}) + P_1(\text{transit})\{P_1(\text{Uber}) + P_1(\text{car})\} - P_1(\text{car})P_1(\text{transit}) = 0,$$

and the option-value flows satisfy

$$P_1(\text{car})P_1(\text{Uber})\Phi + P_1(\text{car})P_1(\text{transit})\Phi - P_1(\text{car})\{1 - P_1(\text{car})\}\Phi = 0.$$

In period 3, direct substitution moves probability only from Uber to transit, while the option-value

terms sum to

$$\frac{\partial[1 - P_1(\text{car})]}{\partial \alpha_{\text{transit}}} [P_3(\text{Uber} \mid m_1 \neq \text{car}) + P_3(\text{transit} \mid m_1 \neq \text{car}) - 1] = 0.$$

In period 2, the corresponding bracket is also zero because Uber and transit exhaust the non-car branch, whereas Uber, transit, and car exhaust the car branch. Thus each period and each mechanism separately sum to zero across the three primary modes. The same is consequently true after adding periods 1 and 3 and the p_s -weighted period 2. Last-mile rides are excluded from this check: an access or egress ride is an additional Uber ride attached to a primary transit trip rather than a reassignment of that primary trip.

C.4 Mapping the model to the data

It will be simpler to discuss the choice of transit plus an access mode as a single choice, and so note that

$$\begin{aligned} P_1(\text{transit-by-Uber} \mid d) &= P_1(a_1 = \text{Uber} \mid m_1 = \text{transit}, d) \\ &\quad \times P_1(m_1 = \text{transit} \mid d), \\ P_3(\text{transit-by-Uber} \mid m_1 \neq \text{car}, d) &= P_3(a_3 = \text{Uber} \mid m_3 = \text{transit}, d) \\ &\quad \times P_3(m_3 = \text{transit} \mid m_1 \neq \text{car}, d). \end{aligned}$$

In our data, we observe the number of Uber trips near a transit stop and the number of Uber trips in the rest of the city. If N is the total number of travelers within $\bar{d}$ kilometers of the transit stop, then the number of Uber trips occurring between d_1 and d_2 kilometers of the transit stop is

$$\begin{aligned} \int_{d=d_1}^{d_2} &\left(P_1(\text{Uber} \mid d) + P_1(\text{transit-by-Uber} \mid d) \right. \\ &\quad \left. + [1 - P_1(\text{car} \mid d)] [P_3(\text{Uber} \mid m_1 \neq \text{car}, d) \right. \\ &\quad \left. + P_3(\text{transit-by-Uber} \mid m_1 \neq \text{car}, d)] \right) f(d) dd, \quad (42) \end{aligned}$$

where $f(d)$ is the population density function. Assuming uniform population density, $f(d)$ simplifies to $2Nd/\bar{d}^2$.

The number of Uber trips happening in the rest of the city by people living between d_1 and d_2 kilometers of the transit stop is

$$\begin{aligned} \int_{d=d_1}^{d_2} &\left(P_2(\text{Uber} \mid m_1 = \text{car}) \times P_1(\text{car} \mid d) \right. \\ &\quad \left. + P_2(\text{Uber} \mid m_1 \neq \text{car}) \times P_1(m_1 \neq \text{car} \mid d) \right) f(d) dd. \end{aligned}$$

C.5 Model simulations

Table A7: Model parameters for simulation in Figure 2

Parameter	Positive impact on Uber trips	Negative impact on Uber trips	Mixed impact on Uber trips
α_{car}	0	0	0
α_{transit}	-1	0.7	1.5
α_{Uber}	-1.5	-1.4	-2
β_{car}	0	0	0
β_{transit}	0.5	0.2	0.5
β_{Uber}	-1.5	-1.4	-2
f_{Uber}	3	10	3
c_{walk}	3	7	4
c_{Uber}	0.1	4	0.3
p_s	0.75	0.50	0.50
p_r	1	1	1
$\sigma_{\alpha_{\text{transit}}}$	0	0.5	0.5
$\sigma_{\alpha_{\text{Uber}}}$	1	2	2
ϕ	0	0	0
n_{sim}	1000	1000	1000

Notes: Each model simulates 1000 individuals with random preferences over modes, and uses distance bands from 0 to 5000 meters for probability calculations. However, all plots are restricted to 0–1200 meters from the station. Parameter definitions are discussed in Section 2.

D Empirics appendix

D.1 Event study for far distances

We adapt our empirical approach slightly to assess the validity of our control group. Because we can no longer use a farther distance band as a control, we instead use untreated observations across all stations and rely on city-by-month fixed effects to absorb secular trends. We implement the estimator of [de Chaisemartin and D’Haultfœuille \(2026\)](#), using a linear model with the inverse hyperbolic sine of trips (in hundreds) as the outcome.³² The event-study specification for a given distance d^* is:

$$y_{dit} = \gamma_{c(i)t} + \delta_{di} \sum_{j \in \{-6, -5, \dots, 6\} \setminus \{-2\}} \alpha_{d^*, j} \times \mathbb{1}\{\tau_{it} = j\} \times \mathbb{1}\{d = d^*\} \\ + (\beta_{1, d^*} \times \mathbb{1}\{\tau_{it} < -6\} + \beta_{2, d^*} \times \mathbb{1}\{\tau_{it} > 6\}) \mathbb{1}\{d = d^*\} + \epsilon_{dit}, \quad \text{where } d = d^*$$

D.2 Adjacent distance bins methodology

Total effects by distance calculation

Our main difference-in-differences estimates come from (5), given by

$$y_{dit} = \gamma_{it} + \delta_{di} + \alpha_{\tilde{d}} \times \mathbb{1}_{\tau_{it} \in \{1, \dots, 6\}} \times \mathbb{1}_{d=\tilde{d}} \\ + \left(\beta_{1, \tilde{d}} \times \mathbb{1}_{\tau_{it} < -6} + \beta_{2, \tilde{d}} \times \mathbb{1}_{\tau_{it} > 6} + \beta_{3, \tilde{d}} \times \mathbb{1}_{\tau_{it} \in \{-2, -1, 0\}} \right) \mathbb{1}_{d=\tilde{d}} + \epsilon_{dit}, \\ \forall d \in \{\tilde{d}, \tilde{d} + 1\}.$$

Each regression is estimated separately by $\tilde{d}$, producing an array of expansion effects at varying distances relative to the next distance band $\tilde{d} + 1$. The α_d can be estimated jointly by stacking all such distance pairs p , where $p \in \{100, \dots, 1100\}$ refers to the “treated” distance in each pair. In the new data set, distance bands 0–100 m and 1100–1200 m will appear once, while all others will appear twice (once in the control group and once in the treatment group). The specification can then be written as follows.

$$y_{dipt} = \gamma_{ipt} + \delta_{dip} + \sum_{\tilde{d}} \alpha_{\tilde{d}} \times \mathbb{1}_{\tau_{it} \in \{1, \dots, 6\}} \times \mathbb{1}_{d=\tilde{d}} \times \mathbb{1}_{p=\tilde{d}} \\ + \sum_{\tilde{d}} \left(\beta_{1, \tilde{d}p} \times \mathbb{1}_{\tau_{it} < -6} + \beta_{2, \tilde{d}p} \times \mathbb{1}_{\tau_{it} > 6} + \beta_{3, \tilde{d}p} \times \mathbb{1}_{\tau_{it} \in \{-2, -1, 0\}} \right) \mathbb{1}_{d=\tilde{d}} \times \mathbb{1}_{p=\tilde{d}} + \epsilon_{dipt}$$

As before, the α estimate the effect at one distance relative to another. In order to estimate the total effect of an expansion on Uber trips at a given distance, we sum all estimates from farther

³²Log or IHS transformations are not scale-invariant ([Chen and Roth, 2024](#)); we normalize trips by 100 so that estimated effects at 0–100 m align with our PPML estimates.

distances. To implement this, we simplify the above regression as follows.

$$\begin{aligned}
y_{dipt} = & \gamma_{ipt} + \delta_{dip} + \sum_{\tilde{d}} \alpha_{\tilde{d}}^* \times \mathbb{1}_{\tau_{it} \in \{1, \dots, 6\}} \times \mathbb{1}_{d=\tilde{d}} \\
& + \sum_{\tilde{d}} \left(\beta_{1,\tilde{d}} \times \mathbb{1}_{\tau_{it} < -6} + \beta_{2,\tilde{d}} \times \mathbb{1}_{\tau_{it} > 6} + \beta_{3,\tilde{d}} \times \mathbb{1}_{\tau_{it} \in \{-2, -1, 0\}} \right) \mathbb{1}_{d=\tilde{d}} + \epsilon_{dipt}
\end{aligned}$$

We no longer estimate the α 's by comparisons within distance-pairs. By not fully-saturating on pairs p , we allow the treatment effect from one distance to pass through into the next, closer distance. Numerically, $\alpha_{\tilde{d}}^* = \sum_{d \geq \tilde{d}} \alpha_d$. Additionally, the regression implementation of summing the coefficients correctly calculates standard errors. The average effect for 1-6 months after the transit opening are reported in Figure A17.

E Model estimation and counterfactuals appendix

This appendix provides technical details for the SMM estimation in Section 6.

Simulation procedure. The model includes persistent unobserved preference heterogeneity for transit and Uber (η_{transit} and η_{Uber}), drawn from a joint normal distribution with mean zero and variance-covariance matrix:

$$\Sigma_{\eta} = \begin{pmatrix} \sigma_{\text{Uber}}^2 & \phi \sigma_{\text{Uber}} \sigma_{\text{transit}} \\ \phi \sigma_{\text{Uber}} \sigma_{\text{transit}} & \sigma_{\text{transit}}^2 \end{pmatrix} \quad (43)$$

These preference shocks affect utility for outbound (Period 1) and return (Period 3) trips but not intermediate (Period 2) trips.

To simulate the model, we draw $N_{\text{sim}} = 500$ realizations of base shocks $\tilde{\eta} = (\tilde{\eta}_{\text{Uber}}, \tilde{\eta}_{\text{transit}})' \sim \mathcal{N}(0, I)$ once and hold them fixed across all parameter evaluations. For a given parameter vector θ , we recover correlated preference shocks via the Cholesky decomposition of Σ_{η} . Specifically, let $L(\theta)$ denote the lower-triangular Cholesky factor such that $L(\theta)L(\theta)' = \Sigma_{\eta}(\theta)$. The simulated preference shocks are then:

$$\begin{pmatrix} \eta_{\text{Uber}} \\ \eta_{\text{transit}} \end{pmatrix} = L(\theta) \begin{pmatrix} \tilde{\eta}_{\text{Uber}} \\ \tilde{\eta}_{\text{transit}} \end{pmatrix} \quad (44)$$

This common random numbers approach ensures that the same underlying draws are used across all parameter evaluations, with only the scale and correlation parameters changing as θ varies (Train, 2009). As a result, the simulated objective function is smooth in θ , which improves the numerical stability of the optimization. For each simulated individual, we compute optimal mode choices via backward induction as described in Appendix C.

Spatial assignment of last-mile trips. Quadrant I moments capture Uber trips observed near the station, which include both direct trips by nearby residents and last-mile trips from residents throughout the catchment area. Last-mile trips originate at the traveler's home but terminate near the station. We model the spatial distribution of these trip endpoints using an exponential density with rate parameter λ .

Formally, let $T^{\text{direct}}(d)$ denote Uber trips by residents at distance d , and $T^{\text{last-mile}}(\delta)$ denote last-mile trips by residents at distance δ . The observed Uber usage at distance d is:

$$T^{\text{observed}}(d) = T^{\text{direct}}(d) + \int_0^{D_{\text{max}}} T^{\text{last-mile}}(\delta) \cdot \rho(d) d\delta \quad (45)$$

where $\rho(d) = \frac{\lambda e^{-\lambda d}}{1 - e^{-\lambda D_{\text{max}}}}$ is the exponential density giving the probability a last-mile trip terminates at distance d from the station, normalized over $[0, D_{\text{max}}]$ with $D_{\text{max}} = 5$ km. We discretize this integral using 100 m distance bins, weighting last-mile trips from each origin distance bin by annulus area to account for radial population density.

The probability of intermediate trips. The model features p_s , the probability that an intermediate trip occurs. We set $p_s = 0.5$ in estimation. We arrive at this value by checking that it implies a similar non-local trip share to the one implied by our citywide, local, and non-local Uber trip estimates. For example, for the 500–800 m POI group, the trimmed-window estimates used in Section 5.3 yield a citywide effect of -0.2% , a local effect of -1.2% , and a non-local effect of $+3.0\%$. Treating the citywide effect as a weighted average of the local and non-local effects implies a weight of 24% on non-local trips, close to the weight implied by $p_s = 0.5$ (i.e., $0.5/2.5 = 20\%$). The parameter estimates are not sensitive to this choice.

Population density and price heterogeneity. We model the residential population around each station as exponentially declining with distance: $\exp(-\kappa \cdot d)$, where κ is the population density gradient (per km) estimated within a 5 km radius. We calibrate κ using the ratio of Uber trips per square kilometer at 1000–1100 m to trips at 1100–1200 m prior to opening, which equals 1.07. This separates κ from c_{Uber} , which would otherwise have similar effects on the pattern of first- and last-mile trips.

To convert utility estimates to willingness-to-pay, we calibrate the marginal utility of income ρ using a demand elasticity of $\varepsilon^D = -0.6$ (Cohen et al., 2016, Castillo, 2025). The expression for ρ implied by the model is given in Proposition 2. Prices vary widely across our sample: the 10th percentile city by average fare is Hyderabad at \$2.35 per trip, while the 90th percentile is Stockholm at \$21.83. The average fare of \$7.84 is similar to prices in Panama City. While using a single price masks this heterogeneity, it allows us to assign a dollar value for the average context in our setting.

Estimation and standard errors. We minimize the GMM objective using L-BFGS-B with box constraints on parameters:

$$\hat{\theta} = \arg \min_{\theta} [m^{\text{data}} - m(\theta)]' W [m^{\text{data}} - m(\theta)] + \mu [n^{\text{calibrated}} - n(\theta)]' [n^{\text{calibrated}} - n(\theta)]$$

where W is the identity matrix and $\mu = 2$. Initial values are set based on the simulations in Section 2. We estimate $\log \lambda$ to ensure positivity and apply the delta method to report standard errors in natural units. We simulate data from 500 draws of η .

Standard errors are computed using the standard one-step GMM formula:

$$\text{Var}(\hat{\theta}) = \frac{1}{N} (G'WG)^{-1} G'W\Sigma WG (G'WG)^{-1} \quad (46)$$

where $N = 650$ station openings, G is the Jacobian of moments with respect to parameters computed via numerical differentiation, and Σ is the variance-covariance matrix of empirical moments from the first-stage PPML regressions.

Computing the mechanism decomposition. Figure 10 reports the mode-by-mechanism flows derived in Appendix C.3, integrated over the estimated model. Let $g(\alpha_{\text{transit}}, d, \eta)$ denote any of the conditional flow derivatives in that appendix. We evaluate its expectation over travelers and integrate along the path of the transit improvement,

$$\mathcal{I}[g] \equiv \int_0^{\hat{\Delta}_{\text{transit}}^\alpha} \mathbb{E}_{d,\eta} [g(\hat{\alpha}_{\text{transit}} + s, d, \eta)] ds, \quad (47)$$

using the same 500 draws of η as in estimation and weighting residence distances up to 2.5 km by annulus area times the estimated population gradient $\exp(-\hat{\kappa}d)$. The remaining parameters are the estimates in Table 5, with $\alpha_{\text{car}} = \beta_{\text{car}} = 0$ and $p_s = 0.5$. Because the path integral is applied to every term, the channel contributions sum exactly to the model-implied change in total Uber usage, whereas multiplying the derivatives at $\hat{\alpha}_{\text{transit}}$ by $\hat{\Delta}_{\text{transit}}^\alpha$ would give only a local-linear approximation.

Each entry in the figure is the corresponding path integral divided by the baseline expected daily usage of the mode in question, evaluated at $s = 0$. For transit and car, the baselines count primary trips over the outbound and return trips plus the p_s -weighted intermediate trip:

$$B_T \equiv \mathbb{E}_{d,\eta} [P_1(\text{transit}) + P_3(\text{transit}) + p_s P_2(\text{transit})], \quad (48)$$

$$B_C \equiv \mathbb{E}_{d,\eta} [P_1(\text{car}) + P_3(\text{car}) + p_s P_2(\text{car})], \quad (49)$$

where $P_3(\cdot)$ and $P_2(\cdot)$ are the unconditional period-3 and period-2 probabilities defined in Appendix C.3. The Uber baseline additionally counts transit access and egress rides,

$$B_U \equiv \mathbb{E}_{d,\eta} [P_1(\text{Uber}_{\text{total}}) + P_3(\text{Uber}_{\text{total}}) + p_s P_2(\text{Uber})], \quad (50)$$

with $P_t(\text{Uber}_{\text{total}})$ as defined in Proposition 1 and $P_2(\text{Uber})$ the unconditional intermediate-trip probability, defined analogously to $P_2(\text{transit})$. Thus the direct-substitution, option-value, and last-mile contributions to Uber usage in the bottom row of the figure are $100 \mathcal{I}[S_1 + S_3]/B_U$, $100 \mathcal{I}[D_1 + p_s D_2 + D_3]/B_U$, and $100 \mathcal{I}[L_1 + L_3]/B_U$, with the first negative because $S_1 + S_3 \leq 0$, and the transit and car boxes apply the same operation to the corresponding rows of the accounting table in Appendix C.3.