Optimal Shared Micromobility Taxes with Distributional Concerns*

Preliminary draft; please do not cite without permission.

Benjamin B. Lockwood, Peter Lugthart, and Arthur A. van Benthem†

November 27, 2024

Abstract

Dockless, shared e-bikes and scooters are a rapidly-growing segment of the local transportation market despite being taxed at higher rates per mile of travel than many conventional transportation modes, including motor vehicles. This paper asks what the optimal tax rate for shared bikes and scooters would be when accounting for environmental externalities and redistribution. We derive a sufficient-statistics model that describes how optimal tax rates depend on both extensive-margin (number of trips) and intensive-margin (duration of trips) demand responses to price changes, as well as the income of bike and scooter users and net externalities relative to other travel modes. We estimate the parameters of this model using trip-level data from a large bike and scooter company and find that in several U.S. cities, use of their vehicles is concentrated among low-income individuals. The model accordingly suggests that policymakers in these cities might want to subsidize shared bike and scooter use even if it generates only modest positive externalities by displacing trips by motor vehicles.

*We thank Michael Partridge for excellent research assistance.

†Lockwood: Wharton School, University of Pennsylvania and NBER. Email: ben.lockwood@wharton.upenn.edu. Lugthart: Wharton School, University of Pennsylvania. Email: lugthart@wharton.upenn.edu. van Benthem: Wharton School, University of Pennsylvania and NBER. Email: arthurv@wharton.upenn.edu.

1 Introduction

Individuals impose costs on society in the form of air pollution, greenhouse gas emissions, traffic congestion, and other externalities when traveling locally. The magnitude of these externalities depends on whether they choose to travel by foot, bike, bus, car, or other modes of transportation. In addition to correcting these externalities with a set of mode-specific Pigouvian taxes, a social planner might wish to consider the distributional consequences of these taxes if mode choice varies by travelers' characteristics. The interaction between these efficiency and equity considerations and the difficulty of measuring externalities makes setting optimal transportation taxes a challenging theoretical and empirical question.

This paper studies optimal transportation taxes in the context of dockless, shared bicycles and electric scooters, which are a growing segment of the transportation market in the United States. Users of these vehicles rent them on a trip-by-trip basis through a mobile app, typically for short journeys of around 10 minutes in duration. Both globally and in the United States, users of these "shared micromobility" (SMM) vehicles typically face taxes and fees that are more than an order of magnitude higher per mile than the taxes levied on motor vehicle travel (Thigpen, Fang, and MacArthur, 2023). The goal of this paper is to estimate the optimal tax or subsidy for SMM trips when accounting for their social cost relative to alternative modes of travel and policymakers' redistributive motives. Our results shed light on whether the large disparity in tax rates between SMM travel and motor vehicle travel is efficient and equitable.

This paper makes both a theoretical and an empirical contribution to the literature on optimal taxation of transportation. In terms of theory, we characterize optimal SMM pricing, using a sufficient statistics approach to identify the parameters that govern the optimal price given distributional considerations and externalities. These parameters include own-price demand responses for SMM use, the net externalities generated by SMM travel relative to travel on other modes of transportation, and the relationship between SMM usage and riders' incomes. Intuitively, if shared scooter and bike trips frequently displace trips with higher social costs, the optimal transportation tax regime would involve a lower effective tax rate on SMM vehicles than if they do so rarely. If lower-income travelers tend to prefer SMM travel, the social planner may want to tax its consumption less heavily than it would if it were only concerned with efficiency. Finally, consumers' demand responses to changes in the tax rate affect both corrective and redistributive taxation by influencing how efficiently the policymaker can change behavior and raise revenue for redistribution.

The paper's empirical contribution is to estimate the parameters of our theoretical model using trip-level microdata from the dockless SMM company, Lime. We use experimental variation in prices to identify own-price elasticities of both number of trips and trip length. To our knowledge, our study would be the first to calculate demand elasticities for SMM vehicles using data on realized consumer decisions. By matching trip data with income data in the areas where trips occur, we can estimate the association between traveler income and mode choice. Combining these estimates with evidence about the net externality of SMM trips relative to counterfactual transit modes provides the full set of statistics relevant for determining optimal policy.

The optimal fully-flexible tax on local transportation would be individual-, mode-, and trip-specific, but such fine-grained policies are logistically impractical to implement and are politically infeasible. We therefore consider simple, constrained-optimal SMM taxes that take tax rates on other modes of travel as given. These policies that can internalize a portion of the externality costs of transportation while meeting constraints such as budget balance and limits on mode-specific tax or subsidy rates. In the analysis below, we consider per-ride and per-minute taxes on SMM use, although these principles naturally extend. In future analysis we may consider other tax structures that are currently implemented by municipal governments.

2 Model

We present a model of consumer optimization and the social planner's problem in the context of per-minute and per-ride taxes on SMM use. These are not the only forms of taxation levied on SMM by local governments. Some municipalities also condition SMM operators' permits to do business on payment of one-time flat fees, flat annual fees, and per-vehicle flat fees. Annual and one-time flat fees do not affect the marginal cost and marginal revenue from each SMM ride purchased by a consumer, so in theory they should not directly affect the pricing decisions of a competitive or monopolistic SMM firm, although they may affect whether a firm decides to enter or exit each market. Per-vehicle flat fees may affect firms' marginal cost and marginal revenue because deploying more vehicles may allow a firm to sell more rides, and these fees raise the cost of doing so. The demand model presented here does not account for any of these supply-side effects, but we plan to extend it to do so in ongoing work.

2.1 Utility and planner's problem

Consumers. Consumers have heterogeneous types, indexed by $\theta \in \Theta \subset \mathbb{R}_+^N$ with population normalized to one and distributed according to $\mu(\theta)$. This type governs their income, ability, and welfare weight (described below). They derive utility from a numeraire good c and from consuming a trip via an SMM service, denoted s, which we interpret as measured in minutes of use:

$$U(c, s; \theta) = c + v(s; \theta) + \chi \cdot \mathbb{1}\{s > 0\}. \tag{1}$$

The function $v(s;\theta)$ represents the subutility of SMM consumption. It is smooth, concave in s, and satisfies $v(0;\theta) = 0$. The parameter χ represents the fixed cost of taking a micromobility trip of any length, including the hassle of finding and unlocking a bike or scooter. This cost is heterogeneous within each type θ, and its marginal distribution is denoted $F_\theta(\chi)$, i.e., $\Pr(\chi < x|\theta) = F_\theta(x)$. This utility function makes it explicit that each individual faces an extensive-margin decision about whether or not to take a single ride on a shared bike or scooter. Conditional on choosing to take a ride, each individual chooses the optimal length $s(\theta)$ of the ride in minutes; this optimal choice of $s(\theta)$ is a function of θ because it depends on $v(s;\theta)$. These two decisions are, respectively, analogous to the labor force participation and work intensity decisions modeled in Saez (2002a) and Saez (2000). We therefore follow Saez (2000) in defining the function $h(\theta)$ to denote the density of individuals of type θ who choose to take a ride.

The price of a micromobility trip consists of two parts: a fixed "unlock" price p_u

and a price p_m for each minute s of micromobility consumption. We consider two (for now) government-determined tax instruments, one for each of the components of the price: a tax t_u on each trip, and a tax t_m on each unit (minute) of usage. For now we assume these taxes are fully passed through to consumers. We assume labor supply is exogenous—meaning in this context that labor supply decisions are not sensitive to taxes on micromobility services—and income (net of any income taxes) is denoted z_θ. The consumer's budget constraint is therefore

$$c = z_\theta - (p_u + t_u) \cdot \mathbb{1}\{s > 0\} - (p_m + t_m)s. \tag{2}$$

This setup produces a model of demand with intensive- and extensive-margin elasticities. Conditional on taking a ride, the consumer's chosen ride length satisfies a first-order condition:

$$v_s'(s_+(\theta); \theta) = p_m + t_m, \tag{3}$$

where $s_+(\theta)$ denotes the optimal number of units consumed by a θ-type conditional on taking a trip; it is the implicit solution to equation (3).

The extensive (participation) margin for each type is governed by the marginal distribution of the fixed costs $F_\theta(\chi)$. Type θ takes a ride if and only if

$$v(s_+(\theta); \theta) + \chi \geq p_u + t_u + (p_m + t_m)s_+(\theta), \tag{4}$$

which implies that the share of θ-types that take rides is

$$h(\theta) := 1 - F_\theta \left(p_u + t_u + (p_m + t_m)s_+(\theta) - v(s_+(\theta); \theta)\right). \tag{5}$$

Combining these definitions, the average number of units consumed by each θ-type is

$$s(\theta) := s_+(\theta)h(\theta). \tag{6}$$

Averaging across θ-types, since the population is normalized to one, the total number of trips taken is

$$\bar{h} := \int_\Theta h(\theta) d\mu(\theta). \tag{7}$$

The total number of units (minutes) of SMM usage consumed by agents of type θ is the product of the length of ride that type θ individuals optimally take and the share $h(\theta)$ of θ-types who take a ride:

$$\bar{s} := \int_{\Theta} s(\theta) d\mu(\theta). \tag{8}$$

Finally, we denote the average trip length among those who take trips as

$$\bar{s}_+ := \frac{\int_{\Theta} s_+(\theta) h(\theta) d\mu(\theta)}{\int_{\Theta} h(\theta) d\mu(\theta)} = \frac{\bar{s}}{\bar{h}}. \tag{9}$$

Social planner. Consumption of each unit (minute) of micromobility travel generates an externality e. To conform with convention in the Pigouvian literature, we let $e > 0$ represent a *negative* externality, so that $e > 0$ corresponds to a positive corrective tax. This represents the net externality from micromobility travel relative to the counterfactual case, and thus it can be positive or negative for any given trip, depending for example on whether the traveler would otherwise have walked or driven. For modeling simplicity, we assume that this externality enters the government's budget constraint as a fiscal cost, and in our baseline we take this externality to be constant, equal to the average net externality across trips.[1]

The policymaker selects taxes t_u and t_m to maximize aggregate utility, weighted by type-specific Pareto weights $\alpha(\theta)$:

$$\max_{t_u, t_m} \left[\int_{\Theta} \alpha(\theta) U(c, s(\theta); \theta) d\mu(\theta) \right], \tag{10}$$

subject to a government budget constraint

$$t_u \bar{h} + (t_m - e)\bar{s} \geq R, \tag{11}$$

to which we assign a Lagrange multiplier λ representing the marginal value of public funds, and subject to consumer optimization:

$$\{c(\theta), s(\theta)\} = \arg\max_{\{s,c\}} U(c, s; \theta) \text{ s.t. } c = z_\theta - (p_u + t_u) \cdot \mathbb{1}\{s > 0\} - (p_m + t_m)s. \tag{12}$$

[1]In ongoing work, we plan to describe and measure empirically how e depends on externalities from counterfactual travel modes and mode-shift elasticities that capture substitution between modes.

2.2 Necessary conditions for optimal taxes

We derive necessary conditions for optimal micromobility taxes using a first-order approach: under the optimal policy, small changes in dt_u and dt_m in the available tax instruments produce no first-order change in total welfare. We provide a heuristic derivation of the optimal tax condition by decomposing the effects of incremental changes in dt_u and dt_m into several effects.

First, a tax increase mechanically transfers income from consumers to the government, holding fixed consumer behavior. The government's valuation of this revenue increase depends on its preferences for redistribution and the extent to which the tax burden falls on individuals to whom the government would like to redistribute. Redistributive preferences are encoded using social marginal welfare weights,

$$g(\theta) = \frac{\alpha(\theta) U_c'(\theta)}{\lambda}, \tag{13}$$

where $U_c'(\theta)$ denotes the marginal utility of consumption for type θ (at a given allocation) and $\alpha(\theta)$ denotes type-specific Pareto weights. Given the quasilinear specification for utility, $U_c'(\theta) = 1$ as in Saez (2001); redistributed preferences are therefore entirely controlled by the Pareto weights $\alpha(\theta)$, which can be interpreted as a reduced-form representation of the policymaker's inequality aversion. Welfare weights are normalized by the marginal value of public funds λ, and because labor supply is exogenous, this implies that welfare weights are normalized to have an average of one, i.e., $\bar{g} = \int_\Theta g(\theta) d\mu(\theta) = 1$. Following Saez (2002b) and Allcott, Lockwood, and Taubinsky (2019), we assume that welfare weights are constant among individuals with a given earnings at the optimum, so that with slight abuse of notation we can write $g(z)$.

Second, a tax increase makes SMM travel more expensive and leads consumers to reduce their consumption of it. This has no first-order impact on consumers' own utility due to the envelope theorem, but it may change welfare in two other ways. The first relates to environmental and congestion externalities: a reduction in SMM use may cause individuals to choose modes of travel that cause a different level of air pollution and/or traffic congestion. The second is a fiscal externality: since consumption of SMM decreases, the government makes less money from taxing it.

The size of these externalities depends on 1) how responsive individuals' consumption of SMM is to changes in taxes, 2) which travel modes individuals switch to, and 3) the difference in air pollution and congestion caused by travel on SMM vehicles and alternative modes.

Necessary conditions for the optimal taxes

We begin with the per-ride ("unlock") tax t_u. A small increase dt_u in the per-ride tax raises tax revenue at the expense of decreasing individuals' consumption, valued at $g(\theta)$, which produces a mechanical effect on welfare—i.e., holding fixed individuals' behaviors—of

$$dt_u \cdot \int_\Theta h(\theta)(1-g(\theta))d\mu(\theta) = dt_u \cdot \left(\bar{h} - E\left[h(\theta)g(\theta)\right]\right) \tag{14}$$

$$= dt_u \cdot \left(\bar{h} - \bar{h}\bar{g} - \text{Cov}\left[h(\theta), g(\theta)\right]\right) \tag{15}$$

$$= dt_u \cdot \left(-\text{Cov}\left[h(\theta), g(\theta)\right]\right), \tag{16}$$

where the final line uses the fact that $\bar{g}$ is normalized to one.

A tax change dt_u also affects consumer behavior through the change in the number of trips taken, which produces both a change in externality costs and fiscal externalities from tax revenues,

$$dt_u \cdot \int_\Theta \left(\frac{dh(\theta)}{dt_u}t_u + \frac{ds(\theta)}{dt_u}(t_m - e)\right)d\mu(\theta) = dt_u \cdot \left(t_u\frac{d\bar{h}}{dt_u} + (t_m - e)\frac{d\bar{s}}{dt_u}\right). \tag{17}$$

At the optimum, the sum of these components must produce no first-order change in social welfare, implying

$$-\text{Cov}\left[h(\theta), g(\theta)\right] + t_u\frac{d\bar{h}}{dt_u} + (t_m - e)\frac{d\bar{s}}{dt_u} = 0. \tag{18}$$

Rearranging, we have

$$t_u^* = \frac{-\text{Cov}\left[h(\theta), g(\theta)\right]}{\left|\frac{d\bar{h}}{dt_u}\right|} + (e - t_m)\frac{\left|\frac{d\bar{s}}{dt_u}\right|}{\left|\frac{d\bar{h}}{dt_u}\right|}. \tag{19}$$

This is a first-order condition for the optimal per-ride tax. The first term can be interpreted as an extensive-margin version of Diamond's "many person Ramsey rule"

(Diamond, 1975) reflecting the intuition that if trips are concentrated among high-welfare-weight consumers, then the optimal per-ride tax should be lower. The second term can be interpreted as a Pigouvian corrective term which raises the tax to the extent that the negative externality e exceeds the per-minute tax t_m, representing the net (of fiscal effects) externality. In the special case where the per-minute fee is set equal to the marginal externality e, then the second term is zero and only the first redistributive term matters.

We can also derive a condition for the optimal per-unit (per-minute) tax t_m. Such a reform produces a mechanical change in welfare of

$$dt_r \cdot \int_\Theta s(\theta)(1 - g(\theta))d\mu(\theta) = dt_r \cdot (\bar{s} - E[s(\theta)g(\theta)]) \tag{20}$$

$$= dt_r \cdot (\bar{s} - \bar{s}\bar{g} - \text{Cov}[s(\theta), g(\theta)]) \tag{21}$$

$$= dt_r \cdot (-\text{Cov}[s(\theta), g(\theta)]) . \tag{22}$$

The reform also produces welfare effects through the changes in fiscal and environmental externalities:

$$dt_m \cdot \int_\Theta dh(\theta)t_u + ds(\theta)(t_m - e)d\mu(\theta) = dt_m \cdot \left(t_u \frac{d\bar{h}}{dt_m} + (t_m - e)\frac{d\bar{s}}{dt_m} \right). \tag{23}$$

Again combining these effects, the optimal per-unit tax t_m necessarily satisfies

$$-\text{Cov}[s(\theta), g(\theta)] + t_u \frac{d\bar{h}}{dt_m} + (t_m - e)\frac{d\bar{s}}{dt_m} = 0, \tag{24}$$

which can be rearranged as

$$t_m^* = \frac{-\text{Cov}[s(\theta), g(\theta)] - t_u \left| \frac{d\bar{h}}{dt_m} \right|}{\left| \frac{d\bar{s}}{dt_m} \right|} + e$$

$$= \frac{-\text{Cov}[s(\theta), g(\theta)]}{\left| \frac{d\bar{s}}{dt_m} \right|} - t_u \frac{\left| \frac{d\bar{h}}{dt_m} \right|}{\left| \frac{d\bar{s}}{dt_m} \right|} + e. \tag{25}$$

Here too the first term represents a redistributive component: when s consumption is concentrated among high-welfare-weight consumers, the optimal per-minute tax is reduced. The second term represents a fiscal externality effect: if raising the per-minute tax t_m decreases the number of trips taken, then those foregone trips reduce fiscal revenues through t_u. And finally, the tax should adjust one-for-one with the externality e.

The difference in the way e enters equations (19) and (25) reflects the fact that the externality is understood to scale with minutes of useage, and thus it is optimally corrected by the per-minute tax t_m. In fact, substituting equation (25) into equation (19) yields

$$t_u^* = \frac{-\mathrm{Cov}\left[h(\theta), g(\theta)\right]}{\left|\frac{d\bar{h}}{dt_u}\right|} + \left(\frac{\mathrm{Cov}[s(\theta), g(\theta)]}{\left|\frac{d\bar{s}}{dt_m}\right|} + t_u^* \frac{\left|\frac{d\bar{h}}{dt_m}\right|}{\left|\frac{d\bar{s}}{dt_m}\right|}\right) \frac{\left|\frac{d\bar{s}}{dt_u}\right|}{\left|\frac{d\bar{h}}{dt_u}\right|}, \tag{26}$$

demonstrating that the optimal per-ride tax does not depend directly on the externality e when the per-minute tax t_m is set optimally.

2.3 Recovering demand responses from per-minute price changes

The empirical covariances and demand responses in equations (19) and (25) can be estimated from data on usage patterns and responses to price changes. Because we have microdata on usage patterns and consumers' characteristics, we can directly estimate the distributions of usage patterns—ride usage $h(\theta)$ and intensity $s_+(\theta)$—as functions of θ provided that relevant variation in types θ can be mapped to observable characteristics, an assumption we invoke for our empirical analysis.

The optimality conditions in equations (19) and (25) also depend on demand responses to both the per-ride tax t_u and the per-minute tax t_m. Our empirical data contains experimental pricing variation in p_m. Here we explain how that variation can be used to recover the demand responses relevant for optimal taxes.

First, we note that p_m and t_m enter identically and additively in intensive- and extensive margin demand in equations (5) and (4), and thus can simply use

$$\frac{ds(\theta)}{dt_m} = \frac{ds(\theta)}{dp_m} \tag{27}$$

and

$$\frac{dh(\theta)}{dt_m} = \frac{dh(\theta)}{dp_m}, \tag{28}$$

which allows us to recover the terms $\frac{d\bar{s}}{dt_m}$ and $\frac{d\bar{h}}{dt_m}$ in equations (19) and (25).

The recovery of the demand responses $\frac{d\bar{s}}{dt_u}$ and $\frac{d\bar{h}}{dt_u}$ is more nuanced because we do not directly observe variation in the unlock price (via either p_u or t_u) in the data. Instead, we invoke an implication of our utility specification, which implies a relationship between the behavioral responses to the unlock price $p_u + t_u$ and the per-minute price $p_m + t_m$. From the definition of $h(\theta)$ in equation (5), differentiating with respect to t_m yields

$$\frac{dh(\theta)}{dt_m} = -f_\theta(\tilde{\chi}) \cdot \left(s_+(\theta) + (p_m + t_m - v'_{s_+}(s_+(\theta); \theta)) \frac{s_+(\theta)}{dt_m} \right). \tag{29}$$

where $\tilde{\chi} = p_u + t_u + (p_m + t_m)s_+(\theta) - v(s_+(\theta); \theta)$. The term multiplying $\frac{s_+(\theta)}{dt_m}$ is equal to zero, implying that adjustments in the optimal ride length on the intensive margin do not create a first-order change in the appeal of taking a ride—a consequence of the envelope theorem—and thus they do not affect the extensive margin response. Thus

$$\frac{dh(\theta)}{dt_m} = -f_\theta(\tilde{\chi}) \cdot s_+(\theta). \tag{30}$$

Now differentiating equation (5) with respect to t_u yields

$$\frac{dh(\theta)}{dt_u} = -f_\theta(\tilde{\chi}). \tag{31}$$

Combining equations (30) and (31) delivers the relationship that allows us to recover the demand response to a change in the unlock price from variation in the per-minute price:

$$\frac{dh(\theta)}{dt_u} = \frac{\left(\frac{dh(\theta)}{dt_m}\right)}{s_+(\theta)}. \tag{32}$$

Integrating both sides of this equation and using equation (28) to replace $\frac{dh(\theta)}{dt_m}$ with $\frac{dh(\theta)}{dp_m}$ yields

$$\frac{d\bar{h}}{dt_u} = \int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta). \tag{33}$$

Because we can measure $\frac{dh(\theta)}{dp_m}$ and $s_+(\theta)$ in the data, this equation allows us to recover $\frac{d\bar{h}}{dt_u}$. Finally, we can also recover $\frac{d\bar{s}}{dt_u}$:

$$\begin{aligned}
\frac{d\bar{s}}{dt_u} &= \frac{d}{dt_u}\int_\Theta s_+(\theta)h(\theta)d\mu(\theta) &(34)\\
&= \int_\Theta s_+(\theta)\frac{dh(\theta)}{dt_u}d\mu(\theta) &(35)\\
&= \int_\Theta \frac{dh(\theta)}{dt_m}d\mu(\theta) &(36)\\
&= \frac{d\bar{h}}{dt_m}, &(37)
\end{aligned}$$

where line (35) uses the fact from equation (4) that $s_+(\theta)$ is not sensitive to t_u.

With these relationships in hand, we are in a position to estimate the key inputs to the optimal tax formulas from empirical microdata.

3 Data

3.1 Lime bike and scooter trips

The key to our empirical work is trip-level data from the shared micromobility operator Lime, which provides e-scooter and bike share services in more than 50 U.S. cities. In order to estimate own-price demand responses, we use data from three cities where Lime has recently conducted pricing experiments: Columbus, OH; Detroit, MI; and Indianapolis, IN. These data are described in greater detail in Section 4.1.1.

We also examine the association between income and SMM use in the following 14 additional U.S. cities, for the time period January 1, 2024 to August 13, 2024: Atlanta, GA; Austin, TX; Boise, ID; Chicago, IL; Washington, D.C.; Denver, CO; Los Angeles, CA; Milwaukee, WI; Minneapolis, MN; Portland, OR; Salt Lake City, UT; Seattle, WA; San Francisco, CA; and Spokane, WA.

Descriptive statistics for the Lime dataset are shown in Table 1. These data describe a rich set of trip characteristics, including the latitude and longitude of the start and end locations, the date and time that the trip started and ended, the distance traveled, whether the device was an e-scooter or a bike, and a unique user identifier.

Lime's trip data also includes complete information about each trip's price. Lime offers four pricing plans. Most trips are taken by riders who use a pay-as-you-go (PAYG) plan, in which they pay a fixed unlock fee (usually $1) to rent a bike or scooter plus a variable per-minute price that typically falls between $0.25 and $0.75. Lime offers two forms of passes that make it less expensive to take many trips or longer trips. "LimePrime" (LP) is a monthly subscription targeted to commuters that waives the fixed cost of unlocking bikes or scooters, allows users to reserve bikes and scooters for 30 minutes prior to a ride rather than the standard 10 minutes, and charges users the regular per-minute price. "Minute-capped ride passes" (MCRP) allow users to buy a block of minutes at a discounted rate and provide free unlocks while the block of minutes lasts. We observe the pricing plan for each trip, and the resulting price.

We also observe rides taken by users through the Lime Access program, which is a means-tested program that provides discounts for low-income riders. The pricing structure for Lime Access users varies by city. Finally, the data describe how prices are affected by coupons that users occasionally apply to their rides. These coupons typically take the form of refer-a-friend discounts and are relatively rare in the experimental cities, comprising 2.3% of all rides.

3.2 Income

We use data from the U.S. Census' 2022 American Communities Survey (ACS) to measure income. We do not observe income and information for individual users of Lime vehicles, so we match trips taken by SMM vehicles with income from the ACS based on the census block group where trips begin and end. To create user-level measures of income, we take the median of the CBG median income in the CBGs where each user started and ended their trips. If a user takes trips with multiple origin and destination CBGs, this income measure is the median of the median income across the set of all CBGs in which their trips started and ended, weighted by the number of times a trip started and ended in a particular CBG.

4 Estimating key parameters for optimal SMM taxes

In this section we describe causal identification of the key demand responses described in section 2.3, $\frac{d\bar{h}}{dp_m}$ and $\frac{d\bar{s}}{dp_m}$, using experimental variation in SMM trip prices. We also describe the empirical covariance between income and SMM use, which determines the extent to which SMM taxes are redistributive.

4.1 Demand responses to changes in per-minute prices

Our estimates of demand responses for scooter trips come from randomized experiments conducted by Lime's pricing team to inform its pricing strategy.[2] In these experiments, Lime randomly assigns treated PAYG users to receive a discounted per-minute price for rides for an extended time period, and compares the count and duration of treated users' scooter trips over time to trips by control users who face typical, undiscounted per-minute prices.

4.1.1 Experimental data description

We use trip-level data from Lime's experiments in Columbus, OH; Detroit, MI; and Indianapolis, IN. Our data for these cities begins on May 26, 2024, and the experiments began on July 26, 2024 and continue through the end of our data on Oct 10, 2024. Treatment affected prices for PAYG trips only, but was offered with equal probability to users who typically purchased pass and PAYG trips (shown in Table 3; see table description below). Lime assigned treatment randomly at the user level at the moment that each user opened the app for the first time after the experiment began. Within each city, 90% of users were assigned the lower price; they constitute the treatment group. The remaining control group users were charged the standard price that all users faced prior to the pricing experiment. The levels of the treatment and control prices are shown in Table 2.

[2] Experimental variation in bike prices is not available at this time. Scooter trips constitute about 89% of PAYG trips across the 14 cities for which we have both scooter and bike trip data, as shown in Table 1.

To estimate the extensive-margin response $\frac{d\bar{h}}{dp_m}$, we first separately aggregated PAYG rides and pass rides (both MCRP and LP) at the user by date level to obtain counts of each ride type. We observe the date of treatment assignment, so the first observation for a user is from the day on which they were assigned to treatment or control by opening the app, even if they did not take a ride on that day. The resulting intermediate dataset is an unbalanced panel, beginning at treatment assignment for each user, in which the unit of observation is a user by date pair. To construct the dataset used for analysis, we balanced this panel by imputing zeros for the ride counts on days after treatment assignment when a user did not take a ride.

Since pass purchases of all kinds account for a minority of observed trips and are not included in Lime's pricing experiments, the sample for our main analysis is PAYG trips. We show in section 4.1.2 that our results are not driven by substitution between pricing schemes. We also exclude Lime Access users from our estimation, since they make up a small minority of Lime users and were not included in Lime's pricing experiments. While the pricing experiment may have caused PAYG users to substitute towards or away from Lime Access, this substitution is most likely limited because Lime Access is a means-tested program.

Intensive-margin demand responses $\frac{d\bar{s}}{dp_m}$ are estimated directly from a trip-level dataset that excludes rides taken with passes and rides taken by Lime Access users.

Table 3 shows pre-treatment outcomes for users assigned to treatment vs. control in order to validate the randomization. This table describes users who took at least one trip on a Lime scooter between May 26, 2024 and the date of their treatment assignment. These users are a small subset of the experimental sample, since most users who were assigned to a treatment group either never took a trip or only took trips after treatment assignment. Nevertheless, we check for balance on observable characteristics of trips, user-by-date observations that match the unit of analysis for the data that we used to estimate treatment effects, and user-level observations that match the level at which treatment is assigned.

Panels A and B of this table document modest differences in trip distance and duration between the treatment and control groups. Among user-level outcomes, which we consider most important for evaluating whether the user-level randomization was carried out correctly, only 2 of 24 variables across 3 cities show statistically significant differences between the treatment and control groups for all three cities. This suggests that the randomization was implemented successfully.

4.1.2 Extensive margin estimation and results

We estimate $\frac{d\bar{h}}{dp_m}$ using the following baseline regression specification, where i indexes individual Lime users and c indexes cities:

$$y_i = \alpha_c + \beta_1 p_i + \epsilon_i. \tag{38}$$

The unit of observation is a user-date pair, but we suppress a t subscript because there is no variation in p within a user, since individuals are either assigned to treatment or control and the estimation sample includes data from the treatment period only. The city fixed effect α_c ensures that variation in p_i used for identification comes only from treatment assignment, rather than from differences in baseline prices across cities, which is crucial when estimating this regression in a pooled, multi-city specification. We also estimate this regression separately for each experimental city, in which case we omit α_c.

In the baseline results, y_i is a count of daily PAYG trips for each user. Since prices are randomized, β_1 is the causal treatment effect of a \$1 increase in price per minute on the daily count of PAYG trips taken by the average user. The coefficient β_1 corresponds directly to $\frac{d\bar{h}}{dp_m}$. To compare this result to estimates from the literature, we transform β_1 into a price elasticity by multiplying by the ratio of the per-minute price to the mean PAYG trip count for the control group.

Baseline estimates of the extensive-margin demand response are shown in the odd-numbered columns of Table 4. The coefficient on price per minute indicates that a \$1 increase in price per minute results in -0.054 trips per user per day in Columbus, -0.015 trips per user per day in Detroit, and -0.029 trips per user per day in Indianapolis. The mean response across all cities is -0.030, which corresponds to a price elasticity of -0.40, indicating that a 1% increase in the scooter per-minute price results in a 0.4% decrease in scooter trips demanded. The estimates are statistically significant at the 5% level in all three cities.

We believe that this is the first study to calculate SMM demand elasticities using observed SMM consumption, and these estimates suggest that SMM users' demand is more inelastic than suggested by prior studies on SMM use. Prior studies have used survey-based or simulated measures of demand, which range from -1.94 to -5.0 for bicycles (Kaviti and Venigalla, 2019; Manout, Diallo, and Gloriot, 2023). Notably, our estimates are similar to estimates of elasticities for public transit fares that use trip-level microdata and an observational research design. For example, Kholodov et al. (2021) find an elasticity of -0.46 for riders on the Stockholm metro, and Wang et al. (2018) find an elasticity of -0.32 for riders on the Beijing metro.

The even-numbered columns of Table 4 show how the extensive-margin demand response varies with income. Across all cities, higher-income users appear more inelastic. The pooled specification in Column (8) indicates that riders with below-median incomes have an elasticity of -0.52, compared to an elasticity of -0.05 for users with above-median incomes. This difference is statistically significant at the 0.1% level; results from individual cities are similar in magnitude and are all statistically significant at the 5% level.

The income matching procedure described in section 3.2 results in a smaller sample than in the baseline specifications shown in odd-numbered columns, since many users who are assigned to either the treatment or the control group do not take a ride during the sample period. The fact that users may select into the estimation sample for the even-numbered specifications may introduce bias in the estimates, since users who are more elastic are more likely to take a ride and appear in the sample when assigned to treatment. To alleviate concerns that bias from selection is driving the results, we re-estimate the specifications in Table 4 using income measured only in the pre-treatment period. These results are shown in Table A1 in the Appendix. This restriction causes a substantial reduction in sample size since most users did not take a pre-treatment ride, and the difference in demand responses between high- and low-income riders is no longer statistically significant. However, the magnitude of the point estimates is very similar to the magnitudes from the full sample, suggesting that bias driven by selection into taking a ride is not driving the results.

The interpretation of our estimates of $\frac{d\bar{h}}{dp_m}$ from the experimental results in Table 4 as the causal effect of a change in price on overall consumption of Lime scooter trips is not valid if users who are assigned to the control group substitute to pass trips at a higher rate than treated users (or vice versa). Higher rates of substitution to pass trips by control group users, the expected direction, would mean that users are more elastic than our estimates suggest. To understand whether this substitution occurs, we estimate the causal effect of treatment on consumption of pass trips using equation (38) as in Table 4, where the outcome y_{ic} is now the daily count of trips by each user for which payment was completed using either a minute-capped ride pass or LimePrime.

The results in Table 5 indicate that substitution to pass trips is not caused by treatment in Detroit and Indianapolis, where we find relatively precise zero effects. In Columbus, it appears that treated users (who face lower prices) are somewhat *more* likely to buy passes than control users, although this effect is not statistically significant. Overall, we interpret these results to indicate that substitution to other forms of payment does not drive our results; if anything, lower prices may encourage users to purchase passes, perhaps by changing their travel patterns.

4.1.3 Intensive margin estimation and results

We use trip-level data from our experimental cities to estimate intensive margin, own-price demand response $\frac{d\bar{s}}{dp_m}$ for scooters. The basic idea of this analysis is to ask whether, conditional on taking a ride, an individual i in city c takes a longer ride when assigned to the treatment group rather than the control group. The baseline regression specification is:

$$s_i = \alpha_c + \beta_2 \ln p_i + \epsilon_i, \tag{39}$$

where s_i is the duration of the scooter trip in minutes, α_c is a city fixed effect (omitted for city-specific regressions), and p_i is the per-minute price. The coefficient β_2 is the intensive-margin demand response.

Table 6 shows the estimates of intensive-margin demand responses for scooters. When the point estimates for these demand responses are translated to elasticities, they are much smaller in magnitude than the extensive margin elasticities across all cities, and none are statistically significant. The elasticity point estimate from the pooled sample of all cities in column (4) is -0.040, which corresponds to a 0.914

minute reduction in scooter trip duration in response to a $1 increase in the per-minute price.

These estimates are consistent with the idea that individuals wish to travel a fixed distance and choose a travel mode based on total cost of taking a trip on a given mode.

4.2 Income associations

In theory, the relationship between micromobility use and income may vary separately on the extensive and intensive margins, as represented by $h(\theta)$ and $s(\theta)$ in the model. We find that this is the case empirically, and that the association between income and SMM consumption also varies by city.

We plot these relationships for six large U.S. cities in Figures 1 and 2. These figures were created by matching data on monthly trip counts and durations, aggregated to the census block group (CBG) level according to the start location of each trip and divided by the CBG population in thousands of people, with CBG-level income data from the American Communities Survey.

Among the fourteen cities in our data, these six cities show a range of variation in the relationship between income and SMM usage. To illustrate this variation, consider Chicago and Los Angeles. Figure 1 shows that in Chicago, CBGs with median income of $50,000 have have trip counts of around 80 trips per 1,000 residents per month, compared to about 120 trips per month in CBGs with a median income of $150,000. In Los Angeles, the highest and lowest-income neighborhoods have similar income levels to Chicago but show the reverse pattern of SMM use: the lowest-income CBGs have about 140 trips per month on average, while the highest-income neighborhoods have fewer than 30 trips per month. Other cities show completely different patterns, such as the non-monotonic relationships in Atlanta, Chicago, and Washington, DC.

On the intensive margin, Figure 2 shows that in most cities, the relationship between income and total trip duration per 1,000 CBG residents is similar to the relationship between trip counts and income. But this is not necessarily the case; in Chicago, the relationship between total trip minutes and income is less monotonic.

In Los Angeles and Seattle, it is noisier than the relationship between trip counts and income.

The association between income and SMM use is, of course, endogenous to the tax system, details of local regulation of micromobility companies, infrastructure such as public transit and bike lanes, and the density and location within each city of high- and low-income CBGs. The associations shown in Figures 1 and 2 do not, therefore, represent causal relationships, but they do describe how the burden of taxes on SMM use are currently distributed across neighborhoods of different income levels.

5 Optimal tax computations

We compute the optimal per-minute and per-ride taxes for the six U.S. cities discussed in 4.2, under two scenarios: 1) a fully-optimal scenario in which the policymaker jointly optimizes both taxes, and 2) a constrained-optimal scenario in which the policymaker sets either t_u^* or t_m^* conditional on the other tax being zero.

5.1 Empirical per-ride tax formulas

To calculate t_u^* when the policymaker optimally chooses t_m^* and t_u^* jointly, we rewrite the joint optimal per-ride tax equation (26) in terms of the empirical demand responses described by equations (27), (28), (33), and (36), to obtain

$$t_u^* \mid_{t_m^*} = \frac{-\mathrm{Cov}[h(\theta), g(\theta)]}{\left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|} + \frac{\mathrm{Cov}[s(\theta), g(\theta)] \left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|} + \frac{t_u^* \left|\frac{d\bar{h}}{dp_m}\right| \left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|}. \tag{40}$$

Rearranging terms to isolate t_u^* produces

$$t_u^* \mid_{t_m^*} = \frac{\frac{-\mathrm{Cov}[h(\theta), g(\theta)]}{\left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|} + \frac{\mathrm{Cov}[s(\theta), g(\theta)] \left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|}}{\left(1 - \frac{\left|\frac{d\bar{h}}{dp_m}\right| \left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|}\right)} \tag{41}$$

The term $1/\left(\left|\frac{d\bar{s}}{dp_m}\right|\left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|\right)$ can be factored from the numerator and denominator and canceled to produce a simplified expression that we estimate in the data, which is

$$t_u^* \mid_{t_m^*} = \frac{\mathrm{Cov}[s(\theta), g(\theta)] \left|\frac{d\bar{h}}{dp_m}\right| - \mathrm{Cov}[h(\theta), g(\theta)] \left|\frac{d\bar{s}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right| - \left|\frac{d\bar{h}}{dp_m}\right|^2} \tag{42}$$

When $t_m = 0$, the optimal per-ride tax equation (19) and equations (27), (28), (33), and (36) imply that

$$t_u^* \mid_{t_m=0} = \frac{-\mathrm{Cov}\left[h(\theta), g(\theta)\right]}{\left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|} + e\left(\frac{\left|\frac{d\bar{h}}{dp_m}\right|}{\left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right|}\right). \tag{43}$$

5.2 Empirical per-minute tax formulas

To obtain an expression for the optimal per-minute tax when t_u^* is set optimally, we use the empirical demand responses described by equations (27), (28), (33), and (36) to rewrite the optimal per-minute tax equation (25) in terms of estimated parameters, which yields

$$t_m^* \mid_{t_u^*} = \frac{-\mathrm{Cov}[s(\theta), g(\theta)]}{\left|\frac{d\bar{s}}{dp_m}\right|} - t_u^* \frac{\left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right|} + e \tag{44}$$

Substituting equation (42) for t_u^* provides the expression for $t_m^* \mid_{t_u^*}$ that we estimate in the data, which is:

$$t_m^* \mid_{t_u^*} = \frac{-\mathrm{Cov}[s(\theta), g(\theta)]}{\left|\frac{d\bar{s}}{dp_m}\right|} - \left[\frac{\mathrm{Cov}[s(\theta), g(\theta)] \left|\frac{d\bar{h}}{dp_m}\right| - \mathrm{Cov}[h(\theta), g(\theta)] \left|\frac{d\bar{s}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right| \left|\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)\right| - \left|\frac{d\bar{h}}{dp_m}\right|^2}\right] \frac{\left|\frac{d\bar{h}}{dp_m}\right|}{\left|\frac{d\bar{s}}{dp_m}\right|} + e \tag{45}$$

When $t_u = 0$, equation (44) implies that

$$t_m^* \mid_{t_u=0} = \frac{-\text{Cov}[s(\theta), g(\theta)]}{\left|\frac{d\bar{s}}{dp_m}\right|} + e \tag{46}$$

5.3 Parameter calibration and estimation

Computing covariances. To calculate $\text{Cov}[h(\theta), g(\theta)]$ and $\text{Cov}[s(\theta), g(\theta)]$, we first create a balanced panel at the user-by-week level, imputing zeros in gaps representing weeks during which an individual did not take a trip and taking weekly averages of trip counts and minutes of scooter travel by user. We then collapse this dataset to the user level to compute $s(\theta)$ and $h(\theta)$. In our static model from Section 2, the objective and budget constraint equations can be interpreted as steady-state versions of these equations, evaluated over a given time period such as a day, week, or month. Because the optimal policies are invariant to rescaling all s and h values by an amount, the choice of time period is not consequential, provided that behavior is consistent across time periods. As a result, h can be interpreted equally as the probability that type θ takes a trip in some short time period, e.g., a day, or as the number of trips that type θ takes over a longer time period, e.g., a week, because converting from the former to the latter effectively amounts to rescaling by 7. Similarly, s can be interpreted as the number of minutes taken in a single trip during some short period, or the aggregate number of trip minutes over a longer period. We use week-level aggregates for our calculations. For each user θ, the term $s(\theta)$ is the average number of SMM minutes consumed per week during the sample period. The term $h(\theta)$ is the average number of trips taken per user per week.

Estimating demand responses. We use the average demand responses across Columbus, Indianapolis, and Detroit estimated from the experimental data. On the intensive margin, this is $\frac{d\bar{s}}{dp_m} = -0.914$, which corresponds to the estimate from column (4) of Table 6. On the extensive margin, we estimate the change in trips per week resulting from a \$1 change in price per minute by aggregating the experimental data to the week level and re-running the specification shown in column (8) of Table 4. The resulting estimate is $\frac{d\bar{h}}{dp_m} = -0.183$.

Calculating income and welfare weights. As in section 4.1.2, each users' net-of-income-tax income z_θ is proxied by the median CBG income[3] across all the CBGs in which the user started or ended a trip during the sample period, on the assumption that the most common trip start/end point represents the user's residence. Welfare weights $g(\theta)$ are proportional to $\frac{1}{z_\theta}$ and normalized to have a mean of one.

Estimating $\int_\Theta \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)$. To predict $\frac{dh(\theta)}{dp_m}$ for each user based on their income, we use the fitted values from a version of the regression specification shown in column (8) of Table 4 with a continuous income measure. This regression is given by

$$y_i = \alpha_c + \alpha_c \text{Income}_i + \gamma_1 p_i + \gamma_2 \text{Income}_i + \gamma_3 p_i \times \text{Income}_i + \epsilon_i. \tag{47}$$

where y_i is the weekly count of trips per user, α_c is a city fixed effect, p_i is price per minute, and Income_i is user income, de-meaned by city in order to make an out-of-sample prediction. From this regression, $\frac{dh(\theta)}{dp_m} = \gamma_1 + \gamma_3 \times \text{Income}_i$. We predict trip duration for weeks during which a user did not take a ride using the fitted values from a regression of trip duration on income and week of the year, and compute $s_+(\theta)$ by taking the mean trip duration for each user.

Computing externalities. We use estimates from the literature of mode-switching and externalities from vehicle travel in order to calculate the per-minute externality from scooter travel. Kuhmann et al. (2024) use a difference-in-differences design to estimate the effect of SMM travel during a pilot period in Chicago, during which SMM use was restricted to a cluster of neighborhoods, on vehicle travel by ride-hailing services. These results are marginally statistically significant but to our knowledge represent the best estimates of mode-switching elasticities from SMM use using observed switching between modes rather than surveys which ask about counterfactual mode choice. From the results in Table 4 of this paper,[4] we calculate that each minute of SMM travel displaces 0.018 miles of vehicle travel. This is most

[3] We take median income within each CBG as our measure of CBG income.

[4] We use the midpoint of each distance range given in Table 3, and apply each estimate to the base number of trips falling in each distance bin, given in Table 3. To calculate the displacement of rideshare miles traveled on a per-SMM-mile-traveled basis, we use the base of 820,000 SMM trips taken during the study period and a mean ride duration of 12.32 minutes calculated for Chicago in the data provided to us by Lime.

likely a conservative estimate, since it only accounts for rideshare trips displaced by SMM use and does not account for SMM users who would counterfactually have taken trips by private car.

To find the externality cost of this displacement, we use estimates of externalities from vehicle travel in terms of greenhouse gas emissions and local air pollution, traffic congestion, and traffic-related fatalities from Langer, Maheshri, and Winston (2017). Adjusted for inflation, the 2024 EPA estimate of the social cost of carbon, and the 2024 fleet-wide fuel economy, this estimate suggests that each mile of vehicle travel costs society \$0.38. The implied externality per minute of SMM travel is 0.018 miles*(-\$0.38) = -\$0.0068.

5.4 (Preliminary) Results

We first note that when the social planner does not have preferences for redistribution, (i.e., $g(\theta) = 1$ for all $\theta \in \Theta$), from equation (42) it follows directly that $t_u^* \mid_{t_m^*} = 0$ and from equation (45) it follows that $t_m^* \mid_{t_u^*} = e = -\$0.0068$. That is, SMM travel receives a subsidy of about 0.7 cents per minute when the social planner sets both taxes optimally in the absence of a redistributive motive. Notably, this is a much lower tax rate than the rates currently in place in many U.S. cities.

The city-specific parameters and estimates of optimal taxes in Table 7 indicate that the externality-based subsidy is very small relative to the magnitude of the taxes and subsidies in Panel B, which account for the social planner's preferences for redistribution. This implies that policymakers may want to take seriously concerns about equity when choosing tax rates for SMM. The magnitudes of these preliminary estimates should be interpreted cautiously, but the direction of the results indicates several key dynamics captured by these taxes.

First, Panel A shows that covariances between SMM consumption and welfare weights largely map to the income associations shown in Figures 1 and 2. Nine of these twelve covariances are positive, which suggests that redistributive motives by policymakers should often point in the direction of lower taxes, or subsidies, on SMM use. This suggestion is confirmed by the constrained-optimal taxes shown in the first two rows of Panel B. If extensive-margin SMM consumption covaries positively with welfare weights, then $t_u^* \mid_{t_m==0} < 0$; that is, a per-ride subsidy is optimal, and vice versa if extensive-margin SMM consumption covaries negatively with welfare weights. Similarly, $t_m^* \mid_{t_u==0} < 0$ if minutes of SMM consumption are concentrated among low-income users. Both of these patterns are evident in comparing constrained-optimal taxes for Chicago and Atlanta. In Chicago, $\text{Cov}[h(\theta), g(\theta)] < 0$, so $t_u^* \mid_{t_m==0} > 0$, and the reverse is true in Atlanta. Meanwhile, the covariance between welfare weights and intensive-margin use is positive in both cities, indicating that both cities should offer per-minute subsidies in the constrained-optimal case.

Second, since demand responses are assumed to be the same across cities, the relative magnitude of the covariances is a major determinant of differences in the constrained-optimal tax rates across cities. For example, since $\text{Cov}[s(\theta), g(\theta)]$ is more positive in Atlanta than in Chicago, $t_m^* \mid_{t_u==0}$ indicates a larger per-minute subsidy in Atlanta.

Turning to jointly optimal taxes, the patterns caused by the policymaker's tradeoffs are more nuanced. In all 5 cities in which either mean trip duration or mean trip count positively covaries with income, the net tax at the mean trip length is negative. However, in many cities the social planner accomplishes this redistribution via a positive per-ride tax and a negative per-minute tax, or vice versa. We save a thorough discussion of the tradeoffs for policymakers in each city that cause the observed tax patterns for a later version of the paper.

References

Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax". *The Quarterly Journal of Economics* 134, no. 3 (2019): 1557–1626.

Diamond, Peter A. "A many-person Ramsey tax rule". *Journal of Public Economics* 4, no. 4 (1975): 335–342.

Kaviti, Shruthi, and Mohan M. Venigalla. "Assessing service and price sensitivities, and pivot elasticities of public bikeshare system users through monadic design and ordered logit regression". *Transportation Research Interdisciplinary Perspectives* 1 (2019): 100015.

Kholodov, Yaroslav, Erik Jenelius, Oded Cats, Niels van Oort, Niek Mouter, Matej Cabecauer, and Alex Vermeulen. "Public transport fare elasticities from smart-card data: Evidence from a natural experiment". *Transport Policy* 105 (2021): 35–43.

Kuhmann, Julius, Tim Haarhaus, Wolfgana Ketter, and Christian Schwens. "Scoot Over, Uber: Assessing E-Scooter Sharing's Impact on Ride-Hailing Usage". *ICIS 2024 Proceedings* 14 (2024).

Langer, Ashley, Vikram Maheshri, and Clifford Winston. "From gallons to miles: A disaggregate analysis of automobile travel and externality taxes". *Journal of Public Economics* 152 (2017): 33–46.

Manout, Ouassim, Azise Oumar Diallo, and Thibault Gloriot. "Implications of pricing and fleet size strategies on shared bikes and e-scooters: a case study from Lyon, France". *hal-04017908* (2023).

Saez, Emmanuel. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses". *Working Paper no. 7708, National Bureau of Economic Research* (2000).

— . "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses". *The Quarterly Journal of Economics* 117, no. 3 (2002a): 1039–1073.

— . "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes". *Journal of Public Economics* 83, no. 2 (2002b): 217–230.

— . "Using elasticities to derive optimal income tax rates". *The review of economic studies* 68, no. 1 (2001): 205–229.

Thigpen, Calvin, Kevin Fang, and John MacArthur. "Shared micromobility fees: current patterns, impacts, and best practices". *Transportation Research Board Annual Meeting* (2023).

Wang, Zi-jia, Feng Chen, Bo Wang, and Jian-ling Huang. "Passengers' response to transit fare change: an ex post appraisal using smart card data". *Transportation* 45 (2018): 1559–1578.

6 Tables and Figures

Table 1: Descriptive statistics for Lime bike and scooter trips

	Mean	Std. dev.	Min	Max	N
Panel A: Trip-level data for scooters					
Trip duration if PAYG (mins)	11.79	20.40	0.00	42,745.18	11,274,253
Trip distance if PAYG (km)	1.85	2.06	0.00	48.28	11,274,253
Trip duration if pass (mins)	10.82	12.20	0.07	4,941.25	4,074,015
Trip distance if pass (km)	1.98	2.13	0.00	48.28	4,074,015
Variable Price if PAYG ($/min)	0.44	0.12	0.00	1.00	11,274,039
Unlock fee if PAYG ($)	0.98	0.08	0.00	1.09	11,274,253
Panel B: Trip-level data for bikes					
Trip duration if PAYG (mins)	12.75	13.41	0.00	1,266.83	1,383,419
Trip distance if PAYG (km)	2.34	2.64	0.00	99.78	1,383,419
Trip duration if pass (mins)	11.80	11.76	0.05	114.58	576,609
Trip distance if pass (km)	2.46	2.55	0.00	99.78	576,609
Variable Price if PAYG ($/min)	0.45	0.11	0.00	0.90	1,383,048
Unlock fee if PAYG ($)	0.96	0.15	0.00	1.09	1,383,419
Panel C: User x date-level data for A/B experiment cities only					
Count of PAYG scooter trips	0.09	0.41	0.00	19.00	2,835,329
Count of pass scooter trips	0.04	0.39	0.00	27.00	2,835,329

Panel A presents descriptive statistics for scooter trips using trip-level data from the 14 U.S. cities included in the non-experimental sample. Pay-as-you-go (PAYG) trips are those for which users pay a fixed unlock fee and a variable price per minute. Pass trips include LimePrime and Minute-capped ride passes, which are described in Section 3.1. Panel B presents descriptive statistics for trip-level data for bikes, which are subject to the same pricing schedules with price levels that are typically slightly higher. The data in Panel C is from trips in the cities of Columbus, Detroit, and Indianapolis where Lime conducted pricing experiments. We aggregate these data from the trip level to the user $\times$ date level in order to estimate demand responses, and impute zeros for all days with no trips between the first and last date that we observe a trip for a user.

Table 2: Characteristics of Lime pricing experiments

	Columbus	Detroit	Indianapolis
Control (typical) variable price ($/minute)	0.49	0.54	0.44
Treatment (discounted) variable price ($/minute)	0.31	0.35	0.29
Number of users assigned to control	5,233	7,713	5,222
Number of users assigned to treatment	46,707	69,372	47,250

This table describes the experiments that Lime uses to estimate price elasticities in the three cities in our experimental sample. Treatment (lower, discounted prices) and control (typical) prices were assigned at the moment that users opened the app in a 9:1 ratio.

Table 3: Lime A/B testing experiments descriptive statistics

	Columbus			Detroit			Indianapolis		
	Control	Treat	p-val	Control	Treat	p-val	Control	Treat	p-val
Panel A: Trip-level data									
Variable price if PAYG ($/min)	0.45 (0.23)	0.45 (0.25)	0.727	0.40 (0.22)	0.40 (0.22)	0.748	0.50 (0.23)	0.50 (0.23)	0.311
Trip duration (minutes)	7.89 (7.98)	8.38 (9.29)	<0.001	11.40 (13.38)	12.06 (13.79)	0.002	10.46 (11.95)	10.25 (11.67)	0.384
Trip distance (km)	1.59 (1.75)	1.61 (1.76)	0.475	1.92 (2.34)	2.05 (2.44)	<0.001	1.87 (1.93)	1.81 (1.95)	0.140
N	5,945	49,091		4,701	40,165		2,597	22,121	
Panel B: User x date-level data									
Count of PAYG trips	0.04 (0.28)	0.04 (0.27)	0.002	0.04 (0.31)	0.04 (0.29)	0.010	0.03 (0.26)	0.03 (0.24)	0.003
Count of discounted PAYG trips	0.00 (0.04)	0.00 (0.05)	0.640	0.00 (0.05)	0.00 (0.05)	0.993	0.00 (0.04)	0.00 (0.04)	0.090
Count of trips with LP	0.05 (0.45)	0.04 (0.40)	<0.001	0.02 (0.31)	0.02 (0.30)	0.891	0.02 (0.29)	0.03 (0.32)	<0.001
Count of trips with MCRP	0.00 (0.03)	0.00 (0.03)	0.001	0.00 (0.04)	0.00 (0.03)	0.283	0.00 (0.03)	0.00 (0.03)	0.578
N	67,573	661,000		75,340	673,849		47,682	382,024	
Panel C: User-level data									
Count of PAYG trips	4.25 (8.83)	3.88 (7.89)	0.259	4.06 (8.57)	3.81 (8.51)	0.445	3.53 (7.99)	3.13 (6.14)	0.207
Count of discounted PAYG trips	0.19 (0.92)	0.20 (0.74)	0.796	0.23 (0.84)	0.23 (1.27)	0.966	0.19 (0.89)	0.15 (0.63)	0.268
Count of trips with LP	5.03 (21.64)	3.93 (19.55)	0.182	2.07 (14.13)	2.11 (13.31)	0.944	2.11 (11.04)	2.78 (13.97)	0.320
Count of trips with MCRP	0.12 (0.52)	0.08 (0.34)	0.008	0.13 (0.42)	0.12 (0.42)	0.422	0.09 (0.33)	0.10 (0.37)	0.686
Ever took a PAYG trip	0.85 (0.35)	0.88 (0.32)	0.022	0.90 (0.30)	0.91 (0.29)	0.446	0.84 (0.37)	0.85 (0.36)	0.420
Ever used a coupon for PAYG	0.11 (0.32)	0.13 (0.34)	0.140	0.13 (0.33)	0.14 (0.34)	0.357	0.12 (0.32)	0.11 (0.31)	0.455
Ever used a LP	0.22 (0.42)	0.22 (0.41)	0.738	0.17 (0.38)	0.16 (0.37)	0.451	0.24 (0.43)	0.23 (0.42)	0.448
Ever used a MCRP	0.09 (0.42)	0.07 (0.41)	0.014	0.11 (0.38)	0.10 (0.37)	0.355	0.08 (0.43)	0.08 (0.42)	0.943
N	636	6,251		755	6,689		456	3,704	

This table presents control means, standard deviations (in parentheses), and balance tests for the T - C difference in means for the pricing experiments conducted by Lime. Panel A presents statistics from the trip-level dataset, and Panels B and C show statistics from data created by aggregating to the user x date and user levels. All data is from trips taken during the pre-treatment period for users who were eventually assigned to treatment or control. Since randomization was conducted at the moment when users opened the app, the pre-treatment period varies by user; we show all trips that happened after May 26, 2024 and prior to treatment assignment. The first date of treatment assignment is July 26, 2024 and the last date of treatment assignment is October 10, 2024.

Table 4: Estimates of extensive-margin demand responses from pricing experiments

	Columbus		Detroit		Indianapolis		All cities	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variable price ($/min)	-0.054***	-0.115***	-0.015*	-0.049*	-0.029***	-0.056***	-0.030***	-0.070***
	(0.012)	(0.032)	(0.007)	(0.020)	(0.008)	(0.011)	(0.005)	(0.013)
$\mathbf{1}\{\text{Income}>\text{median}\}$		-0.099***		-0.044***		-0.018**		-0.024***
		(0.012)		(0.008)		(0.007)		(0.005)
$\mathbf{1}\{\text{Income}>\text{median}\} \times$ Variable price		0.103**		0.052*		0.043*		0.066***
		(0.037)		(0.022)		(0.021)		(0.015)
Cluster-level	User	User	User	User	User	User	User	User
City FE	No	No	No	No	No	No	Yes	Yes
City $\times$ $\mathbf{1}\{\text{Income}>\text{median}\}$ FE	No	No	No	No	No	No	No	Yes
N	2,265,679	1,444,123	3,389,989	2,239,717	2,390,190	1,776,699	8,045,858	5,460,539
Mean elasticity	-.54	—	-.25	—	-.4	—	-.4	—
Low-income elasticity	—	-0.540	—	-0.440	—	-0.560	—	-.52
High-income elasticity	—	-0.110	—	0.040	—	-0.130	—	-.05

Standard errors are shown in parentheses ($*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$) and are clustered at the user level. The outcome variable is pay-as-you-go (PAYG) trip counts. In odd columns, we transform point estimates of the effect of variable price on trip count into the extensive-margin demand elasticity by multiplying the coefficient β_1 on variable price by the ratio of variable price to mean PAYG trip count for the control group. In even columns, we include an indicator for the user earning above median income along with an interaction term between this indicator and variable price. For these columns, income is imputed for users using all pre- and post-treatment trips in the sample. A positive coefficient on this interaction term indicates that higher income individuals are more inelastic. We compute elasticities for low-income earners using the coefficient on variable price, and for high-income earners we add the coefficients on the variable price coefficient and the interaction term coefficient before multiplying by the ratio of variable price to mean PAYG trip count for the control group.

Table 5: Estimates of substitution to pass-based payment

	Pass trip count			
	(1) Columbus	(2) Detroit	(3) Indianapolis	(4) All Cities
Variable price ($/min)	-0.032	-0.003	-0.002	-0.011
	(0.017)	(0.009)	(0.012)	(0.007)
Cluster-level	User	User	User	User
City FE	No	No	No	Yes
N	2,265,679	3,389,989	2,390,190	8,045,858
Elasticity	-.65	-.14	-.06	-.36

Standard errors are shown in parentheses ($*$ p $<$ 0.05, $**$ p $<$ 0.01, $***$ p $<$ 0.001) and are clustered at the user level. The outcome variable is trips taken with a pass, which includes LimePrime and Minute-capped ride passes. We transform point estimates of the effect of variable price on trip counts into the extensive-margin demand elasticity by multiplying by the ratio of per-minute price to mean PAYG trip count for the control group.

Table 6: Estimates of intensive-margin demand responses

	Conditional trip duration ($\bar{s}_+$)				Total daily minutes ($\bar{s}$)
	(1) Columbus	(2) Detroit	(3) Indianapolis	(4) All Cities	(5) All Cities
Variable price ($/min)	-0.722	-2.103	1.399	-0.914	-0.448***
	(0.935)	(1.874)	(1.661)	(0.940)	(0.075)
Cluster-level	User	User	User	User	User
City FE	No	No	No	Yes	Yes
N	128,980	120,578	86,913	336,471	8,045,859
Elasticity	-.04	-.08	.05	-.04	-.37

Standard errors are shown in parentheses ($*$ p $<$ 0.05, $**$ p $<$ 0.01, $***$ p $<$ 0.001) and are clustered at the user level. The outcome variable is trip duration in minutes. We transform point estimates of demand responses in levels into price elasticities by multiplying by the ratio of variable price to mean PAYG trip count for the control group.

Table 7: (Preliminary) Tax parameters & and optimal tax levels

Panel A: Parameters

Statistic	Atlanta	Chicago	Los Angeles	San Francisco	Seattle	Washington DC	All Cities
$\int_{\Theta} \frac{\left(\frac{dh(\theta)}{dp_m}\right)}{s_+(\theta)} d\mu(\theta)$	-0.012	-0.013	-0.013	-0.013	-0.017	-0.014	-0.013
$\text{Cov}\left[h(\theta), g(\theta)\right]$	0.006	-0.006	0.025	-0.003	0.005	0.001	0.008
$\text{Cov}\left[s(\theta), g(\theta)\right]$	0.051	0.030	0.092	-0.045	0.004	0.004	0.043

Panel B: Taxes

Tax	Atlanta	Chicago	Los Angeles	San Francisco	Seattle	Washington DC	All Cities
$t_u^* \mid t_m = 0$	-0.641	0.364	-2.037	0.116	-0.375	-0.173	-0.757
$t_m^* \mid t_u = 0$	-0.031	-0.021	-0.050	0.014	-0.009	-0.008	-0.027
$t_u^* \mid t_m^*$	0.462	-3.524	6.210	0.404	-3.709	0.535	1.518
$t_m^* \mid t_u^*$	-0.071	0.284	-0.588	-0.021	0.312	-0.055	-0.158

Panel A lists city-varying parameters of the inputs to the optimal SMM taxes that are shown in Panel B. We also use several parameters across all cities. Externalities are $e = -\$0.0067$ per minute of SMM travel. Welfare weights $g(\theta)$ are proportional to $\frac{1}{z_\theta}$ and normalized to have a mean of one, where z_θ is each consumer's net-of-income-tax income, proxied by the median of the CBG median incomes from all of the CBGs where each user started their trips during the sample period. We use the average value of demand responses across Columbus, Indianapolis, and Detroit, which are given by $\frac{d\bar{h}}{dp_m} = -0.183$ and $\frac{d\bar{s}}{dp_m} = -0.914$.

Figure 1: Total scooter trip counts by trip start census block group income level

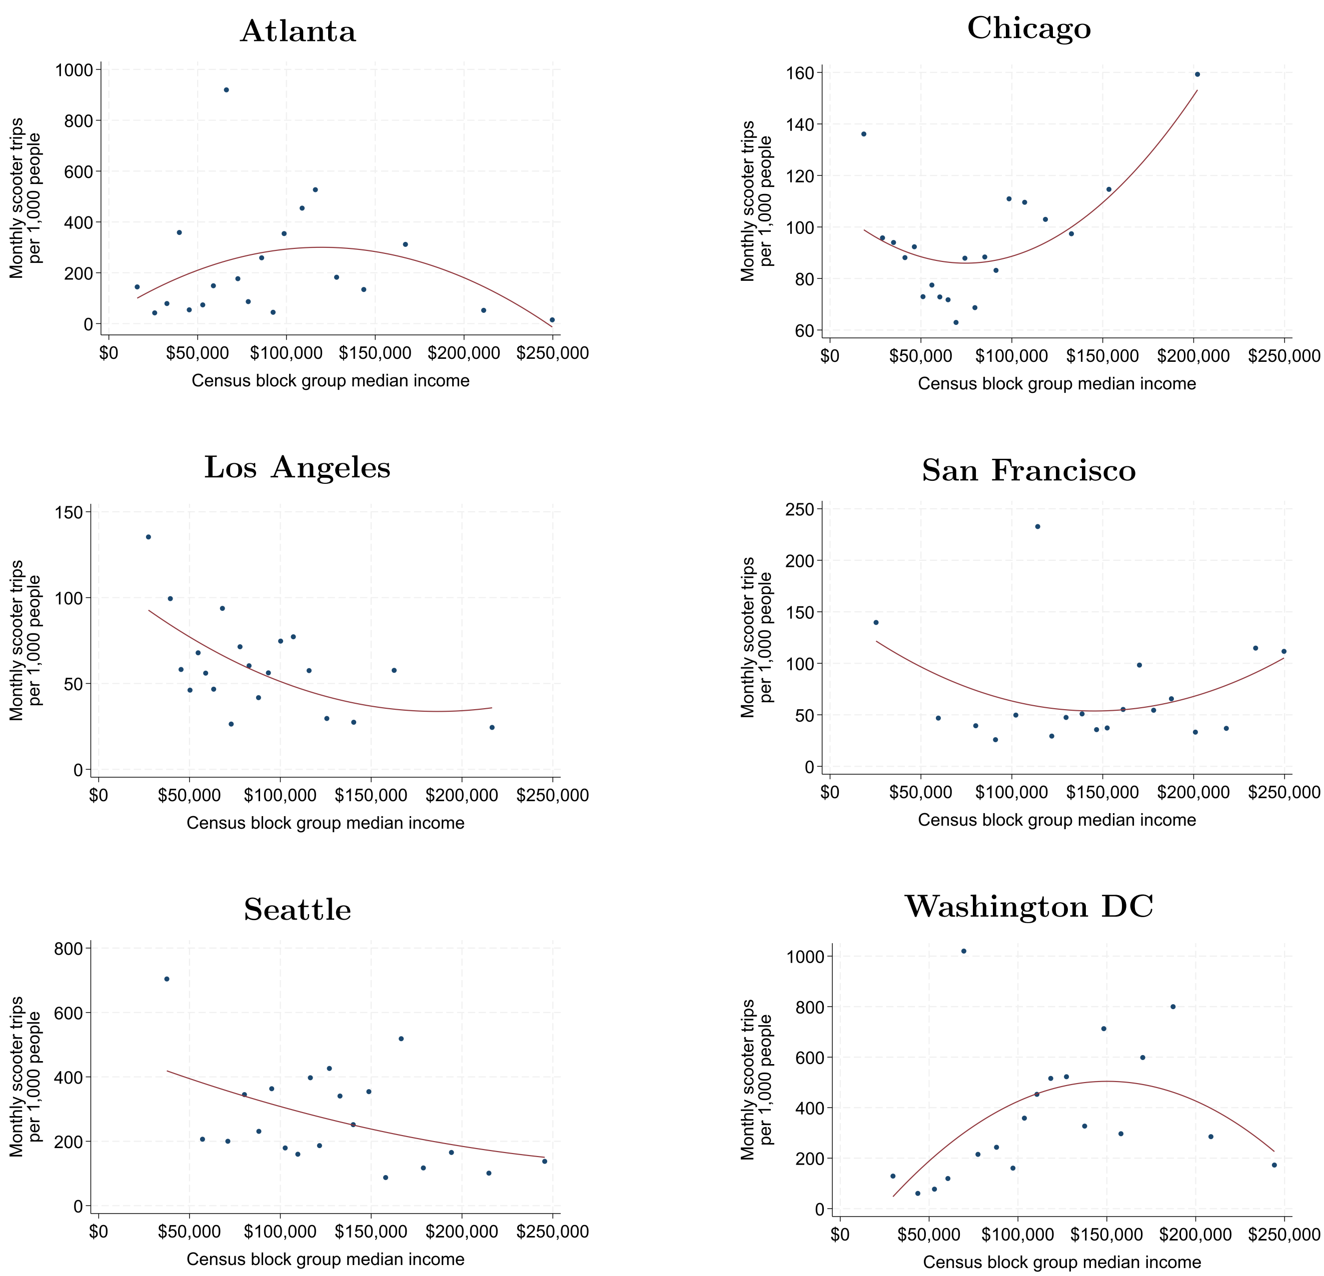

This figure presents scooter trip counts per 1,000 census block group (CBG) residents where the trip began by the median income within the CBG across several U.S. cities. Trips in this figure include pay-as-you-go and pass trips that were taken by non-Access users.

Figure 2: Total scooter trip duration by trip start census block group income level

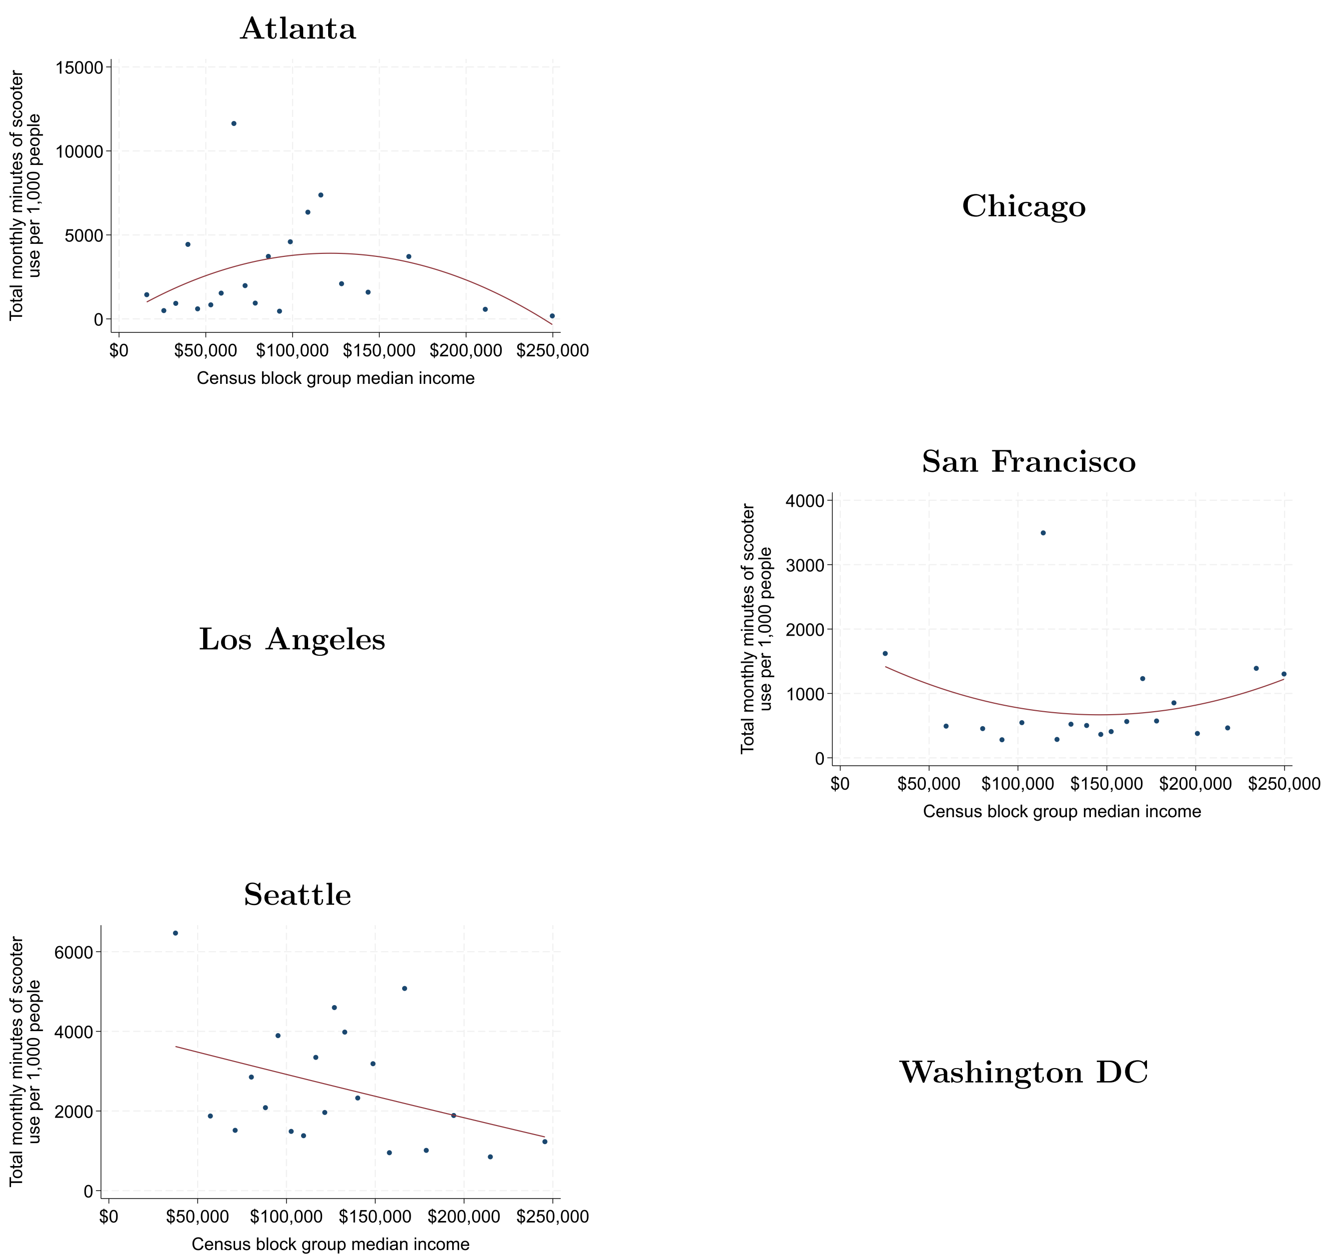

This figure depicts total scooter trip duration in minutes per 1,000 CBG residents where the trip began by the median income within the CBG across several U.S. cities. Trips in this figure include pay-as-you-go and pass trips that were taken by non-Access users.

7 Appendix

Table A1: Extensive-margin demand responses using pre-treatment income

	Columbus	Detroit	Indianapolis	All cities
	(1)	(2)	(3)	(4)
Variable price (\$/min)	-0.090	-0.112	-0.160*	-0.110
	(0.105)	(0.099)	(0.075)	(0.062)
1{Income>median}	-0.090*	-0.112*	-0.059	-0.029
	(0.043)	(0.047)	(0.043)	(0.025)
1{Income>median} × Variable price	0.115	0.111	0.234	0.134
	(0.131)	(0.126)	(0.137)	(0.079)
Cluster-level	User	User	User	User
City FE	No	No	No	Yes
City × **1**{Income>median} FE	No	No	No	Yes
N	310,753	278,528	177,599	766,880
Mean elasticity	—	—	—	—
Low-income elasticity	-.34	-.43	-1.03	-.91
High-income elasticity	.13	.00	.29	.25

Standard errors are shown in parentheses ($*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$) and are clustered at the user level. The outcome variable is pay-as-you-go (PAYG) trip count. In odd columns, we transform point estimates of the effect of variable price on trip count into the extensive-margin demand elasticity by multiplying the coefficient β_1 on variable price by the ratio of variable price to mean PAYG trip count for the control group. In even columns, we include an indicator for the user earning above median income along with an interaction term between this indicator and variable price. For these columns, income is imputed for users using all pre- and post-treatment trips in the sample. A positive coefficient on this interaction term indicates that higher income individuals are more inelastic. We compute elasticities for low-income earners using the coefficient on variable price, and for high-income earners we add the coefficients on variable price and the interaction term before multiplying by the ratio of variable price to mean PAYG trip count for the control group.