Travel Mode Choice and Distributional Impacts of Congestion Charge Policy in New York City

Yikuan Ji*

October 12, 2025

Click here for the latest version.

Abstract

This paper evaluates the distributional and welfare impacts of New York City's recently implemented congestion pricing program, the first in the United States. Using the 2022 Citywide Mobility Survey, I estimate a discrete choice model of travel mode with heterogeneous responses to cost and time, embedded in a congestion feedback framework. The analysis distinguishes between two channels: a price-side effect, capturing the direct burden of the toll on travelers, and a revenue-side effect, reflecting how toll revenues are recycled. The results show that the current toll ($9 peak / $2.25 off-peak) reduces private vehicle trips involving the Manhattan Core by about 19,000 per weekday, with most shifts to subways and walking. Accounting for congestion feedback raises driving speeds and substantially increases welfare relative to a fixed-traffic analysis. The toll alone lowers consumer surplus by $0.32 per person-day, with larger relative burdens on lower-income travelers, while a uniform rebate raises net welfare by $2.24 and makes the policy mildly progressive. Distributional analysis shows that equity outcomes hinge on these two effects: the price-side burden follows income patterns, while the revenue-side outcome depends on policy design. Transit fare subsidies generate smaller average gains while reinforcing mode shifts. The findings highlight that the equity and efficiency of congestion pricing depend jointly on traveler behavior and revenue design, with lessons for other U.S. cities planning similar policies.

Keywords: Congestion Pricing, Distributional Effects, Revenue Recycling

*Department of Agricultural and Resource Economics, University of Maryland, College Park (Email: yk68@umd.edu). I am deeply grateful to my advisors Joshua Linn, James Archsmith, and Cinzia Cirillo for their invaluable guidance and support. I thank Shanjun Li, Lawrence White, Cory Smith, Roger von Haefen, and Mark Anderson for helpful mentoring and comments. I would also like to thank Madalyn Romberger, Yumin Hong, Ana Milan Hinostroza, and seminar participants at the WEAI Graduate Student Workshop, AERE Summer Conference, North American UEA Meeting, APPAM Fall Conference, University of Maryland and Camp Resources for helpful comments. All errors are my own.

1 Introduction

A central challenge in public finance is that many consumption-based taxes and user fees are regressive, placing a heavier burden on lower-income households in relative terms. Gasoline taxes and pollution taxes are well-documented examples: although they improve efficiency by internalizing externalities, they often raise equity concerns when the direct cost of the tax, or price-side effect, falls more heavily on poorer households (Kayser, 2000; Sipes and Mendelsohn, 2001; West, 2004; Grainger and Kolstad, 2010). A large literature shows that the overall distributional impact depends critically on the use of revenues. Through revenue-side measures such as lump-sum transfers, targeted subsidies, or cuts in distortionary taxes, governments can recycle revenues in ways that offset or even reverse the regressive price-side impact, leading to a more equitable overall outcome (Johnstone and Alavalapati, 1998; Chiroleu-Assouline and Fodha, 2014; Fried, Novan, and Peterman, 2021).

A related body of work examines the distributional impacts of road pricing. Both theoretical models and empirical studies find that outcomes vary widely depending on traveler heterogeneity, access to alternatives, and revenue use (Arnott, De Palma, and Lindsey, 1994; Levinson, 2010). Some settings show that congestion tolls place relatively larger burdens on low-income drivers (Cain and Jones, 2008; Manville and Goldman, 2018), while others—such as Stockholm—demonstrate that pricing can be progressive when revenues are reinvested in public transit (Eliasson and Mattsson, 2006). Recent applications in United States (U.S.) cities confirm that results are highly context-specific (Ecola and Light, 2009; Baghestani et al., 2020; Maheshwari et al., 2024). Overall, this literature underscores that the equity of congestion pricing depends as much on revenue allocation and local travel patterns as on the charge itself.

Congestion charge[1] thus provides a natural test case for the broader lessons of regressive taxation and revenue allocation. Economists have long advocated it as a clean solution to traffic externalities (Pigou, 1920; Vickrey, 1963). The idea is straightforward: charge drivers a fee that reflects the cost they impose on others [2]. In addition to correcting congestion externalities, tolls generate substantial revenues that can be recycled. The distributional and efficiency impacts of the policy thus depend critically on how these revenues are allocated, making it a particularly relevant context in which to study distributional impacts.

In the U.S., this question is especially salient. Traffic congestion imposes billions of dollars in annual costs through lost time, wasted fuel, and increased pollution (Schrank et

[1] also referred to as congestion pricing or congestion tolls

[2] In theory, the *first-best* policy would be a finely tuned, time- and location-specific congestion toll equal to the marginal external cost imposed by each additional driver. Such a toll would vary continuously with traffic density and road capacity, ensuring efficiency at every moment. In practice, this design is infeasible due to limitations in monitoring, enforcement, and political acceptability. As a result, real-world policies are necessarily *second-best*. See Appendix A for a discussion of the theoretical first-best framework and why implementation constraints lead to simpler tolling systems.

al., 2024). Experiences in London, Stockholm, and Singapore demonstrate that congestion charges can deliver substantial improvements in traffic conditions, safety, and environmental outcomes (Leape, 2006; Green, Heywood, and Paniagua, 2020; Herzog, 2024; Eliasson, 2009; Simeonova et al., 2021).Yet adoption in the U.S. has been slow, in large part due to concerns about distributional fairness: tolls may disproportionately burden low-income households or commuters from outer neighborhoods, even when they succeed in reducing congestion.

In this context, New York City (NYC) offers the first opportunity to evaluate congestion pricing in the United States. In January 2025, the city launched a cordon-based[3] program requiring drivers to pay when entering Manhattan's Central Business District (CBD). The rollout was turbulent, including a temporary pause, but toll collection has since resumed. This provides a rare opportunity to examine both traveler responses to congestion pricing and how policy design—particularly the use of revenues—shapes distributional outcomes. Building on this context, the paper is structured around three research questions that also constitute its main contributions.

The first question concerns the distributional impacts of congestion pricing across income, demographic, and geographic groups. While much of the literature has emphasized efficiency, showing that congestion charges can reduce traffic and improve travel times (Small & Verhoef, 2007; Parry, Walls & Harrington, 2007), fewer studies have systematically examined distributional outcomes. Evidence from London, Stockholm, and Singapore points to mixed effects: some studies suggest that lower-income households bear a relatively larger burden, while others emphasize the role of transit availability in shaping equity (Eliasson & Mattsson, 2006; Santos & Rojey, 2004).

In the U.S., most analyses rely on simulations with stylized assumptions or aggregate data (Levinson, 2010; Manville & Goldman, 2012; Bagheshtani et al., 2022). Recent research has begun to study NYC directly. Brian Yueshuai He et al. (2021) develop a multi-agent simulation to evaluate congestion pricing policies across population segments by time-of-day, while Cook et al. (2025) document the short-run effects of the policy on traffic volumes and travel times. Yet both studies focus primarily on aggregate efficiency outcomes, leaving open the question of how costs and benefits are distributed across different types of travelers.

My contribution is to provide the first evidence based on individual travel diary data from the New York City Citywide Mobility Survey (NYCCMS). By estimating a discrete choice model of travel mode with heterogeneous responses to cost and time, and embedding it in a congestion feedback framework, I quantify how the toll reshapes both behavior and welfare across diverse traveler groups.

The second question examines how congestion feedback reshapes policy impacts. Most empirical studies treat traffic conditions as fixed, assuming that travel times remain un-

[3]A cordon-based system charges vehicles for crossing into or out of a designated zone, typically a central urban area. This contrasts with distance-based pricing, which charges per mile traveled, or time-of-day pricing, which varies by congestion level.

changed when congestion charges are introduced. Yet economic theory makes clear that congestion pricing alters both demand and supply: by reducing vehicle trips, it increases road speeds, which in turn affect mode utilities and subsequent travel behavior (Kenneth A Small, Verhoef, and Lindsey, 2004). Ignoring these feedbacks risks mismeasuring welfare impacts, since part of the efficiency benefit comes precisely from improved traffic conditions. Recent work highlights the importance of such equilibrium adjustments in measuring congestion externalities(Wang, 2023; Almagro et al., 2024; Barwick et al., 2024). My contribution is to link a discrete choice model of travel mode with a congestion feedback mechanism, closing the loop between travel demand and traffic conditions. This integration makes it possible to evaluate equilibrium outcomes in which mode choices and traffic speeds are jointly determined, offering a richer picture of how congestion pricing affects both efficiency and equity.

The third research question examines how revenue allocation shapes the distributional consequences of congestion pricing. A consistent finding in the broader taxation and environmental policy literature is that the regressivity of externality pricing depends critically on how revenues are recycled (Chiroleu-Assouline and Fodha, 2014; Fried, Novan, and Peterman, 2021). In the transport context, European evidence shows that reinvesting toll revenues in public transit can mitigate regressive effects (Börjesson and Kristoffersson, 2014), while Sen Li, Poolla, and Varaiya (2021) demonstrate that coupling congestion charges with wage policies can enhance equity for ride-hailing drivers. Tarduno (2024) further highlights how imperfect pricing design constrains the equity and efficiency of cordon systems, underscoring the broader importance of policy structure. Yet despite these insights, no study has systematically compared alternative revenue allocation schemes in the U.S. congestion pricing context. My contribution is to fill this gap by simulating multiple designs—including uniform rebates, transit fare subsidies, and targeted discounts—using the estimated mode choice model with congestion feedback. This analysis shows how recycling choices can fundamentally reshape both the efficiency and the equity of congestion pricing.

This paper addresses these questions using detailed trip diary data from the NYCCMS, combined with a discrete choice model of travel mode embedded in a congestion feedback framework. The analysis reveals two distinct channels through which congestion pricing affects welfare and equity. The price-side effect, driven by household travel behavior, income, and vehicle ownership, captures the direct burden of the toll. Lower-income households experience larger welfare losses in relative terms, reflecting their greater reliance on car travel for essential trips and limited flexibility to adjust travel mode or timing. In contrast, higher-income households can more easily substitute toward other modes or absorb the additional cost with smaller proportional impacts.

Quantitatively, the price-side effect of the current toll ($9 peak / $2.25 off-peak) reduces

private vehicle trips involving the Manhattan Core by roughly 19,000 per weekday, with most travelers shifting to subways or walking. Accounting for congestion feedback raises average road speeds and increases total welfare relative to an analysis that holds traffic conditions fixed. On average, the toll alone lowers consumer surplus by about $0.32 per person-day, with losses ranging from $0.17 for low-income households to $0.43 for high-income households. In absolute terms, higher-income travelers lose more, but when welfare changes are normalized by daily income, the pattern becomes regressive: the losses equal roughly 0.20 percent of daily income for low-income households and less than 0.05 percent for high-income households. This contrast underscores that the price-side burden falls disproportionately on lower-income travelers.

The revenue-side effect, in turn, depends on how policymakers recycle toll revenues—through uniform rebates, transit fare discounts, or targeted transfers. These choices determine whether the overall policy becomes progressive or remains regressive. When revenues are returned uniformly to residents, welfare increases substantially: a lump-sum rebate raises net welfare by about $2.24 per person-day, offsetting the direct burden across all income groups. Under this policy, lower-income households gain slightly more than higher-income ones, rendering the overall outcome mildly progressive. Transit fare discounts also improve welfare but generate smaller aggregate gains (about $0.301 per person-day) while reinforcing shifts toward public transit.

Together, these findings show that the distributional outcome of congestion pricing is governed by two distinct mechanisms—the price-side burden, which reflects existing income-based travel behavior, and the revenue-side adjustment, which depends on policy design and redistribution decisions. While the toll itself is regressive when evaluated on the price side, recycling revenues through rebates or transit subsidies can more than offset the burden on lower-income households, transforming congestion pricing from a regressive to a broadly progressive policy.

The remainder of the paper is organized as follows. Section 2 introduces the institutional background. Section 3 describes the data and construction of key variables. Section 4 outlines the theoretical framework. Section 5 presents the empirical estimation results. Section 6 reports the counterfactual simulations, welfare analysis, and distributional effects. Section 7 concludes with policy implications.

2 Background

NYC's CBD Tolling Program is the first cordon-based congestion pricing scheme in the United States. The idea of charging for access to Manhattan dates back to the 1970s, when Mayor John Lindsay first proposed tolls on the East River bridges. It resurfaced in 2007,

when Mayor Michael Bloomberg's *PlaNYC* included a congestion pricing plan, though the proposal stalled in the state legislature. Governor Andrew Cuomo revived the policy in 2017 as part of a broader effort to secure funding for the Metropolitan Transportation Authority (MTA). The turning point came in April 2019, when congestion pricing was formally included in the New York State budget, establishing the legal framework for implementation.

Despite this legislative success, implementation was delayed by environmental reviews, legal challenges, and political disputes between 2020 and 2023. The first rollout was scheduled for June 30, 2024, but Governor Kathy Hochul announced a pause just weeks before the start date, citing equity and economic concerns. The program was finally launched on January 5, 2025, only to face renewed uncertainty when the U.S. Department of Transportation rescinded federal approval in February 2025, requiring toll collection to stop by March. In response, the MTA resumed collection in spring 2025 despite the federal directive, and as of this writing, legal and political disputes remain ongoing.

Under the program, all vehicles entering Manhattan south of 60th Street are charged once per day, with exceptions for the FDR Drive and the West Side Highway (see Figure 1). Passenger cars and small commercial vehicles pay $9 during peak hours with E-ZPass, while the overnight rate is $2.25. Larger trucks and buses are charged between $14.40 and $21.60 depending on vehicle size. Motorcycles are charged $4.50. Taxis and rideshare vehicles are subject to per-trip fees of $0.75 and $1.50, respectively, for any trip entering or crossing the CBD. Vehicles without E-ZPass face higher mailed charges. Emergency vehicles are exempt, and credits are applied for certain tolled bridges and tunnels to avoid double charging.

To address equity concerns, the program offers a low-income discount, providing eligible households with a 50 percent reduction after ten monthly trips. Annual revenues are projected to exceed $1 billion, earmarked by law for MTA capital investments such as subway modernization, bus expansion, and accessibility improvements.

Recent evidence(MTA, 2025) underscores these policy developments. Six months after implementation, Governor Hochul and the MTA reported that congestion pricing reduced vehicle entries into the zone by 11 percent—over 10 million fewer vehicles compared to the previous year—while traffic delays fell by 25 percent in the Congestion Relief Zone and by nearly 10 percent regionwide. Transit ridership reached post-pandemic highs (subway +7%, bus +12%), crashes and injuries in the zone fell by about 15 percent, and pedestrian activity in business districts increased by more than 8 percent relative to 2024. Revenues are on track to fund $15 billion in MTA capital improvements, including subway modernization and bus expansion . To evaluate how these policy changes affect travel behavior and distributional outcomes, the next section describes the data sources used in this analysis

3 Data

This study integrates multiple data sources to capture both travel behavior and mode-specific attributes in New York City. The core dataset is the NYCCMS, which provides trip-level travel diaries and demographic information. To supplement these, I use National Household Travel Survey (NHTS) for household vehicle profiles, as well as external sources to construct travel times and costs across modes. Real-time traffic and routing data from Google Maps and INRIX are incorporated to capture congestion patterns. The following subsections describe each source and the construction of mode-specific attributes in detail.

3.1 Primary Data Sources: NYCCMS

The NYCCMS provides detailed trip diary data for residents across the five boroughs. Survey waves are available for 2017, 2018, 2019, and 2022, with data collection paused during the COVID-19 pandemic. I focus on the 2022 wave because the survey methodology changed substantially beginning in 2022, making earlier waves difficult to combine.

Before 2022, the NYCCMS used a traditional, recall-based questionnaire in which each respondent manually recorded a single day of travel. In contrast, the 2022 wave introduced a smartphone-based application that automatically tracked travel activity over a seven-day period. This new approach captures trips more continuously and accurately, but it also changes how trip duration, distance, and mode are measured, as well as the number of person-days observed per respondent.

Because of these methodological differences, merging the pre-2022 and post-2022 data would introduce inconsistencies and potential measurement bias. To ensure comparability and data quality, I use only the 2022 wave for estimation, while earlier waves are referenced only for descriptive or background comparisons.

Trips are reported at the survey zone level, with ten zones defined within NYC and an additional category for trips outside NYC (Figure 2). These zones group neighborhoods into larger geographic units such as Northern Bronx, Manhattan Core, and Outer Brooklyn, balancing geographic detail with sample size. The survey also records vehicle ownership, household demographics, and geographic locations, which are used to define feasible alternatives in the choice model.

While the NYCCMS captures reported trips and socio-demographics, it does not provide two key determinants of mode choice: generalized travel costs and travel times. To incorporate these attributes, I construct mode-specific costs and travel times using external sources. Section 3.2 details this construction and explains how these attributes are used both to define alternatives in the discrete choice estimation and to simulate counterfactual congestion pricing policies.

3.2 Construction of Travel OD Pairs and Times

The NYCCMS records trips at the level of ten survey zones rather than precise origins and destinations. To translate these zone-level trips into mode-specific attributes that can be compared consistently across alternatives, I construct a set of representative origin–destination (OD) pairs. These OD pairs serve as the basis for querying Google Maps and NYCCMS data to obtain travel times across modes.

The objective is to ensure that the constructed OD pairs capture the average travel distances and patterns observed in the survey while remaining tractable for repeated queries. For each zone pair, I select representative points whose inter-zone distances closely match the average trip distances reported in the NYCCMS. In making these selections, I prioritize connections that are especially relevant for travel behavior: links between major transit hubs (e.g., Penn Station, Grand Central Terminal), employment centers (e.g., Midtown and Downtown Manhattan), and high-demand corridors that carry substantial commuter traffic. This ensures that the constructed alternatives reflect the types of trips most commonly observed in the data, rather than atypical or fringe connections.

Once the representative OD pairs are defined, I collect travel time estimates using the Google Maps API. For each OD pair and mode, I query four representative times of day over a one-week period in February 2025: morning peak (8 AM), midday (1 PM), evening peak (6 PM), and late night (11 PM). These queries are supplemented with real-time traffic speed data from INRIX to better capture congestion patterns on the road network. For for-hire vehicles (FHVs), I adjust the resulting times using a regression of observed FHV durations on personal vehicle durations from the NYCCMS, in order to account for passenger wait and pickup times. Full details of the OD-pair construction procedure are provided in Appendix B.

3.3 Cost Construction

To compare alternatives consistently, I compute monetary costs for each travel mode. These constructed costs are combined with travel times to form generalized costs, which serve as key inputs into both the discrete choice estimation and the counterfactual simulations. Full details of the underlying formulas are provided in Appendix C.

Private vehicles. Costs include wear-and-tear, fuel, parking, and tolls. Fuel expenses are based on trip distance, imputed fuel economy from the NHTS, and average gasoline prices. Parking charges are assigned by zone group [4]. Wear-and-tear costs are applied on a per-mile

[4] In outer zones (Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn, Staten Island), I assume metered street parking dominates, with average hourly rates of \$1.25–\$1.50. In denser zones (Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, Inner Brooklyn), I assume private garage parking dominates, with average hourly rates of about \$11.50–\$25.63. see Appendix **??**

basis using estimates from the American Automobile Association (AAA). Tolls include the CBD congestion charge for trips entering Manhattan south of 60th Street and other bridge and tunnel charges where applicable.

For-hire vehicles (FHVs). Costs follow the New York City Taxi and Limousine Commission fare structure. Each trip includes a base fare, per-mile and per-minute charges, average tips, and New York State sales tax. For trips entering the CBD, I add the $1.50 congestion surcharge.

Public transit. Fares are set according to published schedules: $2.75 per trip for subways and local buses, and distance-based fares for commuter rail.

Active modes. Walking is assumed to have zero monetary cost. For biking, the cost is zero if the household reports owning a private bicycle. Otherwise, trip-specific charges are determined using the Citi Bike pricing structure. This distinction ensures that the model accounts for the difference between private and shared bicycle users.

Use in analysis. These constructed costs and times are merged with the NYCCMS to define the choice set for each observed trip. They enter the discrete choice model as key attributes that determine mode utilities, and they are also the basis for counterfactual simulations of congestion pricing policies. For example, when the CBD toll is introduced, the cost of private vehicle and FHVs are adjusted upward, which changes relative utilities and alters equilibrium mode shares through the congestion feedback mechanism.

3.4 Descriptive Statistics

A first look at the data reveals sharp contrasts between travel within the Congestion Tolling Zone (CTZ) and citywide travel patterns. Mode choice is strongly differentiated by geography (Figure 3). Within the CTZ, walking dominates at 64.2 percent of trips, followed by subway (19.2 percent) and private vehicles (5.6 percent). Citywide, private vehicles account for 27.7 percent of trips, with walking (41.2 percent) and subway (15.8 percent) playing smaller roles. These differences highlight how Manhattan is more transit- and pedestrian-oriented than the outer boroughs, where auto dependence is higher.

Trip purposes differ substantially between the CTZ and the rest of the city (Figure 4). CTZ trip origins are more work-oriented: 26.9 percent are for work or work-related purposes, and another 15.5 percent are transfers between modes. By contrast, 23.4 percent of citywide trip origins are home-based. Destination purposes show a similar pattern (Figure 5). CTZ-bound trips are concentrated in non-home-based activities, with 28 percent associated with

work or work-related purposes. In contrast, trips outside the CTZ are more likely to return home. These statistics reinforce Manhattan's role as a regional employment and activity center.

Focusing on vehicle trips into the CTZ, the distribution of drivers by household income reveals an equity concern (Figure 6). Nearly half (48.3 percent) of CTZ vehicle trips come from households earning $25,000–$49,999—well above this group's 13.8 percent share of trips citywide. High-income households ($200,000+) account for only 8.3 percent of CTZ vehicle trips, below their citywide share of 13.1 percent. Among CTZ drivers in the $25,000–$49,999 group, 77.6 percent of trips are work-related, 83.9 percent of drivers are aged 55–64, and 86.3 percent work in healthcare and social assistance. Nearly all (97.5 percent) are full-time workers. These patterns suggest that many potentially affected drivers are essential workers with limited travel flexibility.

Table 1 reports the monetary costs and travel durations across modes. FHV trips have the highest average cost ($25.79) among all modes, reflecting the base fare, per-mile and per-minute charges, and congestion surcharges. Average FHV trip duration is relatively short at 15 minutes, but with wide dispersion.

Rail trips are less costly on average ($5.37) but exhibit a strongly right-skewed distribution of durations. The mean is about 20 minutes with a large standard deviation, but most trips are much shorter: the median is 10 minutes, the 75th percentile is 20 minutes, and the 90th percentile is 36 minutes. A small number of very long commuter rail journeys (exceeding two hours) inflate the average and account for the extreme dispersion. Bus and subway trips both carry the flat fare of $2.75, with mean durations of about 20 minutes.

Private vehicle trips are the most expensive after FHVs, averaging $11.88 per trip when accounting for fuel, parking, tolls, and wear-and-tear. These trips also average longer durations (21.5 minutes).[5]

Biking costs average $2.68 per trip, though this reflects Citi Bike users; households with private bicycles face zero cost. Bike trips average about 20 minutes in duration, similar to bus and subway. Walking trips are assigned zero monetary cost, with mean durations of around 20 minutes.

An important detail is that the number of observations differs across modes because the choice set is not uniform: each trip includes a minimum of three modes and a maximum of seven modes. This variation arises from how the alternatives are defined given trip characteristics and household attributes, and it explains why the sample sizes differ across rows in Table 1.

Overall, these statistics highlight the wide variation in costs across modes—from free

[5]Reported cost and duration statistics for private vehicle trips are calculated only for households that own a vehicle. For households without a vehicle, the driving cost is artificially set to 1000 to exclude the driving option from the discrete choice model.

walking to costly FHVs and private vehicles—while durations are relatively clustered around 15–22 minutes but with substantial heterogeneity due to congestion and trip distance.

4 Theoretical Model

4.1 Model Trips

In this section, I present a discrete choice model for travelers' mode decisions. I define a tour as a series of linked trips where the origin of the first trip and the destination of the last trip are the same (Trips 1–6 in Figure 7). A tour that begins and ends at home is a home-based tour, while a tour that begins and ends at a non-home location is a non-home-based tour.

I assume that trips within a tour are correlated with one another, but that different tours are independent. This motivates an important assumption about driving consistency: if an individual drives her own vehicle on the first trip of a tour, the likelihood of continuing to drive for the rest of the tour is higher. A strong version of this assumption sets the probability equal to one, implying that once a tour starts with driving, the individual drives for the entire tour. I adopt the strong assumption in the theoretical model that once a tour begins with driving, the individual drives for the entire tour. This assumption captures the consistency of vehicle use within a travel sequence and provides a tractable way to link trips together conceptually [6].

In practice, however, the NYCCMS does not always observe complete tours, and many reported travel sequences span across multiple survey days. This makes empirical estimation at the tour level infeasible with the available data. For this reason, while the theoretical framework is presented at the tour level to highlight behavioral consistency, the empirical analysis is implemented at the trip level. The trip-level specification aligns more directly with the structure of the survey data and allows the construction of mode choice sets without imposing unobserved links across days.

4.2 Multinomial Logit Benchmark

I first present a standard multinomial logit (MNL) model as a baseline framework to analyze commuters' mode choices among a discrete set of alternatives. In this model, individuals are assumed to select the mode that maximizes indirect utility, which depends on observed attributes of the alternatives and an idiosyncratic error term assumed to be independently and identically distributed (i.i.d.) Type I Extreme Value.

[6] In some cases, a tour may include sub-tours. For instance, in Figure 7, Trips 3 and 4 form a sub-tour in which an individual leaves the workplace for a restaurant and then returns to work. The consistency assumption also applies to sub-tours. However, whether sub-tours should be modeled separately depends on their prevalence in the data. In the baseline specification, I exclude sub-tours.

The utility of individual i making trip t by mode j is:

$$\max_{j \in C_{it}} U_{itj} = \alpha_i \cdot Cost_{itj} + \beta_i \cdot Time_{itj} + \delta_{ij} + \mathbf{x}'_{itj}\kappa + \varepsilon_{itj}, \tag{1}$$

where C_{it} is the set of feasible travel modes for individual i, depending on trip attributes (origin, destination, departure time) and personal attributes (vehicle ownership, home location, travel preferences). $Cost_{itj}$ is the monetary cost of mode j, and $Time_{itj}$ is the travel time, which for driving personal vehicles, FHVs, and buses depends on traffic congestion. This captures the direct congestion effect of additional cars. For public transit more generally, travel time also includes schedule delay, access and egress time, waiting time, and in-vehicle ride time. δ_{ij} are alternative-specific constants, and $\mathbf{x}_{itj}$ includes other traveler and trip-level controls.

In addition to alternative-specific constants and trip-level controls, I include interactions between demographic characteristics and the cost and time variables. The motivation is that different groups of travelers may place different weights on money and time. For example, higher-income individuals may be less sensitive to monetary cost, and residents in different locations may have systematically different preferences. Including these interaction terms allows the model to capture heterogeneous tastes across groups, and it also makes it possible to compute the implied value of time (VOT) for each group in later analysis.

Formally, the specification can also be written more compactly as:

$$\max_{j \in C_{it}} U_{itj} = \delta_j + (\alpha_0 + \mathbf{z}'_i\alpha) \cdot Cost_{itj} + (\beta_0 + \mathbf{z}'_i\beta) \cdot Time_{itj} + \mathbf{z}'_i\theta_j + \mathbf{x}'_{itj}\kappa + \varepsilon_{itj}, \tag{2}$$

where $\mathbf{z}_i$ is a vector of demographic characteristics of individual i (such as income, age, and geographic location); α and β are vectors of interaction coefficients with cost and time; θ_j are mode-specific demographic effects; and $\mathbf{x}_{itj}$ are other mode–trip controls.

While the benchmark MNL provides a tractable framework, it restricts substitution patterns across alternatives. In Appendix D, I present extended nested logit (NL) specifications, including a baseline NL, alternative nesting structures, and an advanced joint mode–destination model. These models are useful for robustness checks and for illustrating richer substitution patterns across travel modes.

4.3 Congestion Feedback: Traffic Speed–Density Model

The total driving demand $D_c(v_c)$ in each area c depends on trip distances and the probability of choosing to drive:

$$D_c(v_c) = \sum_{i \in I_c} d_i \cdot P_i^{\text{drive}}(v_c), \tag{3}$$

where d_i is the trip distance for individual i, and $P_i^{\text{drive}}(v_c)$ is the probability of driving, which depends on travel costs, travel time (itself a function of average speed v_c), and other individual characteristics.

To link traffic density with travel speeds, I adopt a constant-elasticity speed–density relationship, following the standard traffic flow literature (Greenshields et al., 1935):

$$D_c = \exp(\alpha_0 + \pi \cdot \log(v_c)), \tag{4}$$

or equivalently,

$$v_c = \exp\left(\frac{\log(D_c) - \alpha_0}{\pi}\right) = e^{-\alpha_0/\pi} D_c^{1/\pi}, \tag{5}$$

where D_c is the driving demand (traffic density) in zone c, v_c is the corresponding average traffic speed, and π is the elasticity parameter calibrated from the empirical literature. This formulation directly links the demand for driving to road congestion: as more individuals choose to drive, traffic density rises and speeds fall.

Standard mode choice models often treat travel time as fixed, but in reality, travel times are endogenous: mode choices affect congestion, and congestion in turn affects travel times. By embedding the speed–density relationship from traffic literature, this model captures the feedback loop between demand and supply. Congestion pricing reduces driving, which raises average traffic speeds; faster speeds then feed back into travelers' utilities and influence future mode choices.

The elasticity parameter π is calibrated using evidence from Yang, Purevjav, and Shanjun Li (2020), who study a natural experiment in Beijing. They estimate $\pi = -1.1$, implying that a 1 percent increase in vehicle volume reduces average speed by roughly 1.1 percent. While this estimate is drawn from Beijing, the calibration is plausible for NYC given their similarities as dense global megacities: both have populations exceeding eight million within city limits, large numbers of daily vehicle entries into the central business district, and average commute times that are among the longest worldwide. These parallels suggest that the elasticity parameter from Beijing provides a reasonable proxy for NYC traffic conditions, while still acknowledging that the exact magnitude may vary with local infrastructure and driver behavior. This empirical grounding allows the model to capture system-wide effects of congestion pricing policies.

5 Empirical Estimation Strategy

5.1 Estimation of the Multinomial Logit Model

I estimate the MNL model implied by the utility specification in Equation 1. The parameters of interest capture how mode choice probabilities respond to monetary cost, travel time, and mode-specific constants, and are estimated using maximum likelihood with person weights provided by the NYCCMS.

An important quantity derived from the estimated coefficients is the VOT. VOT measures the marginal rate of substitution between travel time and money, reflecting how much travelers are willing to pay to reduce travel time by one unit. Given cost coefficient α_i and time coefficient β_i, the VOT for group i is defined as:

$$\text{VOT}_i = -\frac{\beta_i}{\alpha_i}. \tag{6}$$

If cost is measured in dollars and time in minutes, VOT_i is expressed in dollars per minute; multiplying by 60 yields dollars per hour. Because the MNL is estimated with interactions between demographics and cost/time, I compute group-specific VOTs by plugging in the relevant coefficients for each subgroup (e.g., income or age group).

5.2 Congestion Feedback and Counterfactual Speeds

The travel mode choice model is embedded within a congestion feedback structure, where mode shares determine traffic density and, in turn, average travel speed. In equilibrium, mode shares and traffic conditions must be jointly consistent.

The starting point is the observed baseline traffic speed v_1, which implies a baseline traffic density of

$$D_1 = \exp(a_0 + \pi \cdot \log(v_1)), \tag{7}$$

where π is the elasticity parameter from the speed–density model.

Suppose a policy such as congestion pricing changes the share of drivers. The resulting adjustment in driving demand produces a new counterfactual density D_2. The relationship between traffic density and speed implies

$$\log\left(\frac{D_1}{D_2}\right) = \pi \cdot \log\left(\frac{v_1}{v_2}\right), \tag{8}$$

which can be solved for the counterfactual speed v_2.

The updated speed feeds back into travel times for driving, FHV and bus modes, which enter the mode choice model and update the utilities of all modes. Revised mode shares then produce a new density, which implies a new speed, and the cycle repeats.

This process is iterated until the changes in mode shares, densities, and speeds become negligible from one step to the next. At that point, the model produces a counterfactual equilibrium that consistently reflects both travel behavior and congestion dynamics.

A key assumption underlying this framework is that travelers act as congestion-takers[7]. Each traveler chooses a mode given prevailing costs and travel durations, without internalizing the small effect of their own decision on overall congestion. In aggregate, however, these choices feed back into traffic density and road speeds. The resulting equilibrium can therefore be interpreted as the outcome of a large population of individuals making optimal mode choices given expected conditions, with the distribution of choices in turn determining those conditions. This assumption, standard in the transportation economics literature, provides a tractable way to capture system-wide congestion effects.

5.3 Empirical Results

5.3.1 Baseline Multinomial Logit Results

Table 2 presents the baseline estimates from the MNL model of mode choice, with bus as the reference category. Across all specifications, the estimated coefficients on both monetary cost and travel time are negative and statistically significant, confirming that higher generalized costs and longer travel times reduce the likelihood of choosing a mode. The implied VOT ranges between approximately \$13 and \$37 per hour depending on the specification, which falls within the range reported in the transportation economics literature for U.S. urban settings.

Turning to the alternative-specific constants, the results reveal systematic preferences across modes relative to bus travel. Driving a private vehicle exhibits the largest positive constant, indicating a strong baseline preference for car use after accounting for cost and time. In contrast, rail modes such as commuter rail have substantially smaller constants, suggesting they are less attractive options compared to bus travel. Subway constants fall in between: their magnitude decreases once geography and demographics are controlled for, indicating that subway attractiveness varies considerably across traveler types and locations. Overall, the pattern of constants highlights large heterogeneity in baseline preferences across modes, even after accounting for generalized costs.

Non-motorized modes also display distinct patterns. Walking has the largest constants among all modes, indicating a strong baseline attractiveness for short trips even after accounting for relative cost and travel time. By comparison, biking has much smaller constants across all specifications, consistent with the relatively limited adoption of cycling in NYC

[7] The term *congestion-taker* is directly analogous to the *price-taker* assumption in competitive markets. Each individual ignores their own marginal impact on congestion, just as a price-taker ignores their marginal impact on market prices.

despite the availability of CitiBike and similar programs.

FHV,including taxis and ride-hail services, also have relatively small constants, placing them below bus in terms of baseline attractiveness. This pattern is consistent with FHVs being used more selectively, often for convenience or in response to circumstances where other modes are less suitable.

Taken together, the baseline MNL results reveal a clear preference ordering: private vehicles and walking are the most attractive modes relative to bus, while rail, bike, and FHVs are substantially less attractive. Subway falls between these extremes, with its relative position depending on the specification. These findings provide an important benchmark before turning to more flexible nested logit (NL) models that allow for correlation across alternatives.

5.3.2 Interaction Multinomial Logit Results and Value of Time

Tables 3 and 4 report the interaction MNL estimates, where the coefficients on cost and duration are allowed to vary by income group, age group, and residential zone. Together, these results highlight meaningful heterogeneity in mode choice preferences across socioeconomic and geographic dimensions.

Income groups. Cost sensitivity is lowest among low-income travelers (coefficient -0.0086), while their time sensitivity remains high (-0.0539), producing a relatively low VOT of about \$9.6 per hour. VOT rises with income, reaching \$26.7 for upper-middle-income households (\$100k–\$249k) and \$21.5 for the highest-income group ($\geq$ \$250k). However, when normalized by estimated hourly wages, the pattern reverses: low-income travelers have the highest normalized VOT (0.77), while high-income travelers have the lowest (0.11). This suggests that congestion pricing imposes a relatively greater burden on lower-income groups, who give up a larger share of their income to save time, even though their absolute willingness to pay for time savings is smaller. These normalized VOT patterns underscore that using monetary tolls alone to allocate road space can be regressive, since lower-income groups bear a larger relative cost of time savings.

Age groups. Younger travelers (ages 18–34) exhibit the highest VOT at \$25.2 per hour, while older adults (55+) report the lowest at \$10.1 per hour. This pattern is consistent with life-cycle expectations: younger and middle-aged commuters may face tighter time constraints due to work or childcare responsibilities, making travel time more valuable. Older individuals, in contrast, appear less willing to pay for time savings, reflecting more flexible schedules or lower opportunity costs of time.

Residential zones. Residents of inner zones (Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, Inner Brooklyn) exhibit higher VOT (\$33.7 per hour) compared to denser zones (Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn,

Staten Island), at $31.8 per hour. Although both groups display high VOTs relative to income and age breakdowns, the slightly higher inner-zone VOT may reflect greater congestion intensity and stronger competition for time in denser areas.

Overall, the results show that VOT is lowest for low-income and older travelers, and highest for upper-middle-income households, younger commuters, and inner-zone residents. When normalized by wages, however, the estimates show a regressive pattern, with lower-income travelers sacrificing a larger share of income to save time. These differences matter for policy: congestion pricing is more likely to induce behavioral responses among groups with higher absolute VOTs, while low-income and older travelers may face higher burdens relative to their income and valuation of time savings.

5.3.3 Semi-Elasticities of Cost and Duration

The semi-elasticities reported in Tables 5 and 6 are derived from the coefficients of the MNL estimates in column (4) of Table 2. They summarize how mode choice probabilities respond to changes in monetary cost and travel time.

Cost elasticities. Table 5 shows that all modes have negative own-cost semi-elasticities, as expected. The largest magnitude is observed for FHVs (−0.31), followed by private vehicles (−0.06). Public transit modes (bus −0.04, subway −0.03, rail −0.09) display smaller elasticities, consistent with flat transit fares. Bike also shows a moderate elasticity (−0.03), while walking is unaffected by monetary cost by construction.

Cross-cost elasticities indicate modest substitution patterns. The most notable is that increases in private vehicle costs raise the probability of choosing FHVs (0.019) and public transit (0.016–0.023). In contrast, higher transit costs generate very small shifts toward other modes, underscoring the limited short-run substitutability for transit riders.

Duration elasticities. Table 6 confirms that all modes respond negatively to increases in their own duration. Subway (−0.48) and bus (−0.61) have the largest own-duration elasticities among transit modes, indicating strong sensitivity to service speed. FHVs (−0.71) and private vehicles (−0.50) also respond to travel time, though to a smaller extent. Among active modes, biking shows the strongest response (−0.56), indicating that even small increases in biking time substantially reduce the probability of biking, while walking is somewhat less sensitive (−0.32).

Cross-duration elasticities highlight several significant substitution relationships. Longer subway travel times increase the probability of choosing FHVs (0.205), private vehicles (0.151), and walking (0.427). Similarly, longer bus times shift riders toward subway (0.049) and walking (0.283). Biking and walking also act as close substitutes: a 1% increase in biking time raises the probability of walking by 0.621 percentage points. In other words, if the baseline walking probability were 20%, it would increase to about 20.6% following a

1% increase in biking time. By contrast, longer walking times increase biking probability by only 0.009 percentage points. These results suggest that substitution is strongest (i) within transit modes, especially between bus and subway, and (ii) between active modes, biking and walking.

Overall, the semi-elasticities indicate that monetary cost is most influential for FHVs and private vehicles, while travel time plays a more dominant role for transit and active modes. Substitution effects align with intuitive clusters: private vehicles and FHVs, bus and subway, and bike and walk.

Taken together, the MNL estimation results and the elasticity analysis provide a comprehensive view of travel mode sensitivities to cost and duration. These estimates form the behavioral foundation for the simulation exercises in Section 6, where I evaluate counterfactual scenarios such as congestion pricing and revenue allocation mechanisms. By embedding the estimated mode choice parameters into a congestion feedback model, the simulations quantify how policy interventions shift mode shares, alter traffic speeds, and redistribute welfare across income, age, and residential groups.

6 Counterfactual Simulations and Distributional Effects

Building on the estimation results, this section presents counterfactual simulations that combine the mode choice model with the congestion feedback mechanism. The objective is to evaluate how congestion pricing policies affect travel mode shares, traffic conditions, and welfare distribution across demographic and geographic groups.

6.1 Simulation Scenarios

The counterfactual simulations are based on the income interaction specification of the MNL model (column (1) of Table 3), which captures heterogeneity in cost and time sensitivities across income groups. This choice allows the analysis to highlight distributional consequences of congestion pricing.

Scenario 1 simulates the current congestion toll policy ($9 peak / $2.25 off-peak, applied once per day for trips crossing into the Manhattan Core) without embedding the congestion feedback loop. This initial case provides a baseline counterfactual in which only the direct demand response to higher vehicle costs is considered.

Scenario 2 incorporates the congestion feedback mechanism. Following the method outlined in Section 4.3, I compute the short-run equilibrium in which mode choice and traffic

conditions are jointly consistent. In this framework, higher vehicle costs reduce driving demand, which in turn alleviates congestion, raises average speeds, and updates travel times across modes. The resulting equilibrium reflects both the direct effect of the toll on mode utilities and the indirect effect operating through congestion relief.

Scenario 3 explores the redistribution of toll revenues through a uniform per-capita rebate to all travelers. This transfer does not affect mode choice directly but alters welfare distribution by offsetting part of the toll burden, particularly for lower-income and more price-sensitive groups.

Scenario 4 evaluates an alternative revenue allocation policy that subsidizes public transit fares. I test two designs: (i) a \$0.35 discount for all bus and subway riders, and (ii) a larger \$0.90 discount targeted at Manhattan Core travellers. For each case, I compare the total cost of the fare discount program—calculated using predicted mode shares under the toll—with the toll revenues collected. If revenues exceed the cost of the fare discount, the remaining funds are distributed as a uniform lump-sum rebate to all travelers, as in Scenario 3. This policy design therefore combines targeted fare reductions with broader redistribution, affecting both mode choice behavior and overall welfare.[8]

In addition to these policy experiments, I conduct two supplementary analyses. First, I calculate the marginal congestion cost (MCC) by adding one additional vehicle trip while holding other mode choices fixed. This provides a measure of the external cost imposed by an additional driver on the system. Second, I perform a sensitivity analysis by varying the toll level from \$0 to \$100 (holding FHV charges fixed). For each toll level, I simulate welfare outcomes under the congestion feedback loop, which traces how changes in toll level affect consumer surplus, revenues, and net welfare. These supplementary analyses provide benchmarks for interpreting the main policy experiments and allow comparison with theoretical measures of external costs.

Across all scenarios, I measure welfare changes relative to the baseline, aggregating results across groups defined by income, age, and residential zone. This framework makes it possible to evaluate both the efficiency and the equity implications of congestion pricing and alternative revenue allocation policies.

[8] I also tested other discount levels. For the citywide case, \$0.35 is close to the maximum feasible discount; values between \$0.40 and \$1.00 exceed predicted revenues and are therefore infeasible. For the Manhattan Core case, a \$0.90 discount is feasible but \$1.00 is not.

6.2 Mode Share Changes: All Trips vs. CBD-Trips

Tables 7 and 8 report the simulated mode shares for all NYC trips and for trips with origins or destinations in the Manhattan Core. [9] Across the city as a whole, the introduction of the toll produces only modest shifts in travel behavior. Vehicle trips decline by about 19,000 per day (–0.06 percentage points), while subway, walking, and bus trips rise by about 13,000, 6,000, and 3,000 trips, respectively.

All of this adjustment occurs among travelers whose trips involve the Manhattan Core, since the toll is geographically targeted, and trips that never cross the cordon are unaffected. For CBD-related trips, the impacts are much more pronounced. Out of CBD-trips, vehicle use falls by 0.72 percentage points (pp) per day, offset by increases of 0.38 pp subway trips, 0.24 pp walking trips, and 0.02 bus trips.

The congestion feedback loop moderates the amount of switching. Without feedback, the model predicts that about 19,000 daily trips shift away from vehicles. With congestion relief, this falls to about 16,000 per day, as faster driving speeds make car use more attractive. Most switchers move to the subway, with smaller flows into walking and biking.

To better understand these behavioral adjustments, it is useful to dig deeper into the simulation beyond the summary tables. The detailed results show that vehicle-to-subway switching is concentrated in transit-rich outer zones such as Northern Manhattan and Middle Queens, while vehicle-to-walk substitution is most common in Inner Brooklyn and Inner Queens. These patterns matter for policy design: transit fare subsidies directly benefit those most likely to switch into public transit, while lump-sum rebates compensate groups with fewer alternatives.

Transit fare discount scenarios generate additional, though still limited, changes. Under a $0.35 citywide discount, bus and subway ridership increase by about 10,000 (+0.03 pp) and 32,000 per day (+0.10 pp), while vehicle trips decline by about 22,500 per day (–0.07 pp). The targeted $0.90 discount for CBD-trips produces a similar pattern. Because these discounts directly reduce generalized costs for transit modes, they encourage modest additional substitution away from private vehicles, concentrated in CBD-related trips.

Taken together, these results show that congestion pricing generates limited citywide mode share changes, though stronger effects within the CBD. Transit fare discounts amplify these changes slightly by making buses and subways more attractive relative to private vehicles. The next section examines how these mode share shifts translate into welfare outcomes for different traveler groups.

[9] All results are scaled using the weighted average of 32,113,908 weekday trips reported in the 2022 NYCCMS. Mode-share percentage point changes from the discrete choice model are multiplied by this daily total to obtain estimates of the number of trips affected.

6.3 Welfare Analysis and Distributional Effects

Table 9 reports the average per-person, per-day welfare effects of the counterfactual simulations. All scenarios are evaluated relative to the current toll policy ($9 peak / $2.25 off-peak, applied once per day for trips entering the Manhattan Core). I present results both with and without accounting for congestion feedback, but place more emphasis on the feedback model since it provides an internally consistent equilibrium of mode choices and travel conditions. All Scenario 4 results are based on the feedback model.

Current toll policy. Without any recycling of toll revenues, the current toll reduces consumer surplus by about $0.39 per person per day without feedback (Scenario 1) and $0.32 with feedback (Scenario 2). These losses are concentrated among vehicle users, only partly offset by gains to those who switch into faster transit or walking trips. Because the charge applies to the Manhattan Core, most of the losses fall on CBD-related trips, though modest impacts are spread across the city.

Uniform rebates (Scenario 3). Recycling toll revenues as a uniform per-capita rebate more than offsets these losses. Net welfare rises by about $2.14 without feedback and $2.24 with feedback per person per day. While the rebate is modest in size, it is distributed broadly—including to non-travelers—so all groups gain on average. This scenario demonstrates how redistributing toll revenues widely can transform congestion pricing from a welfare-reducing to a welfare-improving policy.

Transit discounts (Scenario 4). Using revenues to subsidize public transit produces smaller average welfare gains than rebates but reinforces the behavioral impacts of congestion pricing. A $0.35 citywide fare reduction generates a net gain of about $0.30 per person per day, while a $0.90 discount for CBD-trips produces a similar gain of $0.28. Because the cost of the subsidy is linked to ridership, surplus revenues remain and are returned as lump-sum rebates. Unlike uniform rebates, however, these benefits are concentrated among transit riders—the groups most likely to substitute away from driving—so the policy both supports mode shifts and improves distributional outcomes.

Transit fare discounts generate smaller net welfare gains than rebates because the benefits are narrowly targeted to transit riders, while the largest toll burdens fall on drivers. Moreover, much of the subsidy goes to riders who would have used transit anyway, limiting its effect on mode switching. In contrast, a uniform rebate reaches the entire population, including drivers and non-travelers, which makes it a more efficient redistribution mechanism and delivers larger overall gains.

Distributional impacts. Table 10 breaks down welfare effects by income, age, residential zone, gender, race, and employment. Without revenue recycling, nearly all groups experience losses, with the largest dollar costs falling on high-income and employed individuals (since they are more likely to drive into the CBD). However, when normalized by income, the burden is clearly regressive, with low-income travelers losing the largest share of their daily wages.

Revenue recycling fundamentally reshapes this picture. Uniform rebates provide equal per-capita benefits, leaving all groups better off on net. Transit discounts, by contrast, direct benefits more narrowly. Younger commuters gain the most, reflecting their greater reliance on transit. Geographically, outer-zone residents benefit more from rebates, while CBD-related travelers gain more from transit subsidies.

In short, the current congestion toll without revenue recycling imposes welfare losses that are regressive when measured relative to income. Recycling toll revenues can reverse these effects: uniform rebates maximize redistribution and deliver the largest overall welfare gains, while transit subsidies yield smaller gains but strengthen behavioral responses by lowering the cost of alternatives to driving.

6.4 Marginal Congestion Cost and Sensitivity Analysis

In addition to the main counterfactual scenarios, I conduct two supplementary analyses to benchmark the efficiency of the congestion toll.

Marginal congestion cost (MCC). The first analysis estimates the MCC, defined as the daily welfare loss from adding one additional vehicle trip while holding all other mode choices fixed. This calculation yields an MCC of about \$61 per day across all travelers, of which roughly \$60 comes from delays imposed on private vehicle users and about \$1 from delays to FHV and bus passengers. Importantly, the MCC is calculated at the baseline—that is, without any toll and without accounting for congestion feedback. As such, it provides a benchmark for the externality imposed by an additional driver under current conditions, but it does not fully reflect equilibrium adjustments once tolls are in place.

Sensitivity analysis. The second analysis varies the peak toll from \$0 to \$100, setting the off-peak toll equal to one-quarter of the peak, and simulates welfare outcomes while accounting for congestion feedback. Figure 8 and Table 11 summarize these results. Net welfare rises steadily with the toll, peaking at around \$80 per person per day before declining at higher toll levels. The fact that the welfare-maximizing toll exceeds the MCC reflects the distinction between concepts: the MCC captures the external cost of a marginal trip under

baseline conditions, while the optimal toll balances this externality against demand responses and the recycling of revenues in an equilibrium setting.

The sensitivity analysis highlights the nonlinear relationship between toll size, consumer surplus, and net welfare. At low toll levels, consumer surplus falls quickly while rebates rise, producing steady net welfare gains. At very high toll levels, however, toll revenues increase only slowly, so the per-capita rebate is nearly constant. Meanwhile, consumer surplus losses continue to grow, causing net welfare to decline despite larger transfers. Overall, the results suggest that politically feasible toll levels such as the current $9 charge are far below the welfare-maximizing level, but that raising tolls too aggressively would lead to diminishing returns in net welfare gains.

7 Conclusion

This paper provides the first comprehensive distributional analysis of NYC's congestion toll using detailed trip diary data and a structural mode choice framework with congestion feedback. The analysis highlights three broad lessons for the design of congestion tolling policies.

First, congestion pricing without revenue recycling imposes broad welfare losses. At the current toll level ($9 peak / $2.25 off-peak), consumer surplus declines by about $0.32 per person per day once congestion feedback is taken into account. Losses fall disproportionately on drivers into the Manhattan Core, but are spread across income and demographic groups. While higher-income individuals bear larger losses in absolute terms, the burden is regressive when measured relative to income, as low-income travelers lose the largest share of their daily wages. This price-side effect reflects how income, vehicle ownership, and access to public transit shape households' exposure to the toll.

Second, the revenue-side effect, which depends on how toll revenues are used, plays a decisive role in shaping overall outcomes. Returning revenues to households as uniform rebates turns the toll into a welfare-improving policy, raising average net welfare by more than $2 per person per day and leaving all income and demographic groups better off. Transit fare subsidies generate smaller average welfare gains, but they strengthen the behavioral effects of tolling by lowering the cost of transit for the groups most likely to substitute away from driving. These findings show that the design of revenue recycling determines whether congestion pricing remains regressive or becomes progressive, emphasizing the need to connect pricing policy with fiscal decisions.

Third, the efficiency of congestion pricing depends on toll levels and behavioral responses. The estimated marginal congestion cost at baseline is about $61, but the welfare-maximizing toll is closer to $80 once feedback and revenue recycling are accounted for. The divergence

reflects that the MCC captures the external cost of a marginal trip under current conditions, while the optimal toll balances congestion externalities against induced mode shifts and redistribution. At very high toll levels, however, toll revenues increase only slowly while consumer surplus losses accelerate, limiting further welfare gains.

Taken together, the results show that congestion pricing can improve both efficiency and equity, but only when policymakers consider the two mechanisms jointly. The price-side effect depends on traveler characteristics such as income, car ownership, and transit accessibility, while the revenue-side effect depends on policy design and how revenues are redistributed. Overlooking either dimension risks misjudging both the fairness and effectiveness of the policy. With well-designed transfers, congestion pricing can reduce traffic congestion and yield broad welfare improvements across income and demographic groups.

Beyond New York City, these findings offer guidance for other U.S. cities considering congestion pricing, including San Francisco, Los Angeles, and Chicago. Cities with greater income inequality or limited public transit options are likely to face a more regressive price-side burden unless revenues are carefully recycled. In contrast, denser cities with extensive transit systems may experience stronger behavioral responses and smaller welfare disparities. The framework developed in this paper provides a way to evaluate both mechanisms together by linking the travel behavior that drives the price-side burden with the fiscal policies that determine the revenue-side outcome. Ultimately, the fairness and efficiency of congestion pricing depend not only on how much is charged, but also on how the resulting revenues are used.

References

Almagro, Milena et al. (2024). "Optimal Urban Transportation Policy: Evidence from Chicago". In: *NBER Working Paper* 32185. URL: https://www.nber.org/papers/w32185.

Arnott, Richard, André De Palma, and Robin Lindsey (1994). "The Welfare Effects of Congestion Tolls with Heterogeneous Commuters". In: *Journal of Transport Economics and Policy* 28.2, pp. 139–161. ISSN: 0022-5258. URL: https://www.jstor.org/stable/20053032.

Baghestani, Amirhossein et al. (2020). "Evaluating the Traffic and Emissions Impacts of Congestion Pricing in New York City". In: *Sustainability* 12.9. Publisher: Multidisciplinary Digital Publishing Institute, p. 3655. ISSN: 2071-1050. DOI: 10.3390/su12093655. URL: https://www.mdpi.com/2071-1050/12/9/3655.

Barwick, Panle Jia et al. (2024). "Efficiency and Equity Impacts of Urban Transportation Policies with Equilibrium Sorting". In: *American Economic Review* 114.10, pp. 3161–3205. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20220212.

Börjesson, Maria and Ida Kristoffersson (2014). "Assessing the welfare effects of congestion charges in a real world setting". In: *Transportation Research Part E: Logistics and Transportation Review* 70, pp. 339–355. DOI: 10.1016/j.tre.2014.07.006. URL: https://www.sciencedirect.com/science/article/abs/pii/S1366554514001276.

Brownstone, David and Kenneth A. Small (1989). "Efficient Estimation of Nested Logit Models". In: *Journal of Business & Economic Statistics* 7.1. Publisher: [American Statistical Association, Taylor & Francis, Ltd.], pp. 67–74. ISSN: 0735-0015. DOI: 10.2307/1391839. URL: https://www.jstor.org/stable/1391839.

Cain, Alasdair and Peter M. Jones (2008). "Does Urban Road Pricing Cause Hardship to Low-Income Car Drivers?: An Affordability-Based Approach". In: *Transportation Research Record* 2067.1. Publisher: SAGE Publications Inc, pp. 47–55. ISSN: 0361-1981. DOI: 10.3141/2067-06. URL: https://doi.org/10.3141/2067-06.

Chiroleu-Assouline, Mireille and Mouez Fodha (2014). "From regressive pollution taxes to progressive environmental tax reforms". In: *European Economic Review*. Sustainability and Climate Change: From Theory to Pragmatic Policy 69, pp. 126–142. ISSN: 0014-2921. DOI: 10.1016/j.euroecorev.2013.12.006. URL: https://www.sciencedirect.com/science/article/pii/S0014292113001712.

Cook, Cody et al. (2025). "The Short-Run Effects of Congestion Pricing in New York City". In: 33584. URL: http://www.nber.org/papers/w33584.

Ecola, Liisa and Thomas Light (2009). "Equity and Congestion Pricing: A Review of the Evidence". In: URL: https://trid.trb.org/View/891235.

Eliasson, Jonas (2009). "A cost–benefit analysis of the Stockholm congestion charging system". In: *Transportation Research Part A: Policy and Practice* 43.4, pp. 468–480. ISSN: 0965-8564. DOI: 10.1016/j.tra.2008.11.014. URL: https://www.sciencedirect.com/science/article/pii/S0965856408002140.

Eliasson, Jonas and Lars-Göran Mattsson (2006). "Equity effects of congestion pricing: Quantitative methodology and a case study for Stockholm". In: *Transportation Research Part A: Policy and Practice* 40.7, pp. 602–620. ISSN: 0965-8564. DOI: 10.1016/j.tra.2005.11.002. URL: https://www.sciencedirect.com/science/article/pii/S0965856405001618.

Fried, Stephie, Kevin Michael Novan, and William Peterman (2021). *Recycling Carbon Tax Revenue to Maximize Welfare*. SSRN Scholarly Paper. Rochester, NY. DOI: 10.17016/FEDS.2021.023. URL: https://papers.ssrn.com/abstract=3865416.

Grainger, Corbett A. and Charles D. Kolstad (2010). "Who Pays a Price on Carbon?" In: *Environmental and Resource Economics* 46.3, pp. 359–376. ISSN: 1573-1502. DOI: 10.1007/s10640-010-9345-x. URL: https://doi.org/10.1007/s10640-010-9345-x.

Green, Colin P., John S. Heywood, and Maria Navarro Paniagua (2020). "Did the London congestion charge reduce pollution?" In: *Regional Science and Urban Economics* 84. DOI: https://doi.org/10.1016/j.regsciurbeco.2020.103573.

Greenshields, Bruce D et al. (1935). "A study of traffic capacity". In: *Highway research board proceedings*. Vol. 14. 1. Washington, DC, pp. 448–477. URL: https://onlinepubs.trb.org/Onlinepubs/hrbproceedings/14/14P1-023.pdf.

He, Brian Y. et al. (2020). "Evaluation of city-scale built environment policies in New York City with an emerging-mobility-accessible synthetic population". In: *Transportation Research Part A: Policy and Practice* 141, pp. 444–467. ISSN: 0965-8564. DOI: 10.1016/j.tra.2020.10.006. URL: https://www.sciencedirect.com/science/article/pii/S096585642030745X (visited on 11/02/2024).

He, Brian Yueshuai et al. (2021). "A validated multi-agent simulation test bed to evaluate congestion pricing policies on population segments by time-of-day in New York City". In: *Transport Policy* 101, pp. 145–161. ISSN: 0967-070X. DOI: 10.1016/j.tranpol.2020.12.011. URL: https://www.sciencedirect.com/science/article/pii/S0967070X20309483.

Herzog, Ian (2024). "The City-Wide Effects of Tolling Downtown Drivers: Evidence from London's Congestion Charge". In: *Journal of Urban Economics* 144. DOI: https://doi.org/10.1016/j.jue.2024.103714.

Johnstone, Nick and Janaki R. R. Alavalapati (1998). *The Distributional Effects of Environmental Tax Reform*. Num Pages: 47 Series: Discussion Paper DP 98-01. DOI: 10.22004/ag.econ.24140. URL: https://ageconsearch.umn.edu/record/24140.

Kayser, Hilke A (2000). "Gasoline demand and car choice: estimating gasoline demand using household information". In: *Energy Economics* 22.3, pp. 331–348. ISSN: 0140-9883. DOI: 10.1016/S0140-9883(99)00043-2. URL: https://www.sciencedirect.com/science/article/pii/S0140988399000432.

Leape, Jonathan (2006). "The London Congestion Charge". In: *Journal of Economic Perspectives* 20.4, pp. 157–176. ISSN: 0895-3309. DOI: 10.1257/jep.20.4.157. URL: https://www.aeaweb.org/articles?id=10.1257/jep.20.4.157.

Levinson, David (2010). "Equity Effects of Road Pricing: A Review". In: *Transport Reviews* 30.1. Publisher: Routledge, pp. 33–57. ISSN: 0144-1647. DOI: 10.1080/01441640903189304. URL: https://doi.org/10.1080/01441640903189304 (visited on 11/02/2024).

Li, Sen, Kameshwar Poolla, and Pravin Varaiya (2021). "Impact of congestion charge and minimum wage on TNCs: A case study for San Francisco". In: *Transportation Research Part A: Policy and Practice* 148, pp. 237–261. ISSN: 0965-8564. DOI: 10.1016/j.tra.2021.02.008. URL: https://www.sciencedirect.com/science/article/pii/S0965856421000367.

Maheshwari, Chinmay et al. (2024). *Congestion Pricing for Efficiency and Equity: Theory and Applications to the San Francisco Bay Area*. arXiv:2401.16844 [cs]. DOI: 10.48550/arXiv.2401.16844. URL: http://arxiv.org/abs/2401.16844.

Manville, Michael and Emily Goldman (2018). "Would Congestion Pricing Harm the Poor? Do Free Roads Help the Poor?" In: *Journal of Planning Education and Research* 38.3. Publisher: SAGE Publications Inc, pp. 329–344. ISSN: 0739-456X. DOI: 10.1177/0739456X17696944. URL: https://doi.org/10.1177/0739456X17696944.

MTA (2025). *Six Months In, Governor Hochul Highlights Success of Congestion Pricing: Traffic Is Down, Business Is Up, and Critical Investments Are Being Made to Improve Transit — Governor Kathy Hochul*. URL: https://www.governor.ny.gov/news/six-months-governor-hochul-highlights-success-congestion-pricing-traffic-down-business-and.

Pigou, A. C (1920). *The Economics of Welfare*. Macmillan. URL: https://oll.libertyfund.org/titles/pigou-the-economics-of-welfare.

Schrank, David et al. (2024). "2023 Urban Mobility Report". In: *Texas A&M Transportation Institute, Tech. Rep.* URL: https://static.tti.tamu.edu/tti.tamu.edu/documents/mobility-report-2023-appx-a.pdf.

Simeonova, Emilia et al. (2021). "Congestion Pricing, Air Pollution, and Children's Health". In: *Journal of Human Resources* 56.4, pp. 971–996. DOI: 10.3368/jhr.56.4.0218-9363R2. URL: https://jhr.uwpress.org/content/56/4/971.

Sipes, Kristin N and Robert Mendelsohn (2001). "The effectiveness of gasoline taxation to manage air pollution". In: *Ecological Economics* 36.2, pp. 299–309. ISSN: 0921-8009. DOI: 10.1016/S0921-8009(00)00230-5. URL: https://www.sciencedirect.com/science/article/pii/S0921800900002305.

Small, Kenneth A, Erik T Verhoef, and Robin Lindsey (2004). *The Economics of Urban Transportation (3rd ed.)* 2nd ed. London: Routledge. ISBN: 978-0-203-64230-6. DOI: https://doi.org/10.4324/9781315157375.

Tarduno, Matthew (2024). "For whom the bridge tolls: Congestion, air pollution, and second-best road pricing". In: URL: https://matt-tarduno.github.io/assets/jmp.pdf.

U.S. Energy Information Administration (2024). *Weekly Retail Gasoline Prices: New York All Grades All Formulations*. Accessed: 2024-04-15. URL: https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=pet&s=emm_epm0_pte_sny_dpg&f=w.

Vickrey, William S. (1963). "Pricing in Urban and Suburban Transport". In: *The American Economic Review* 53.2. Publisher: American Economic Association, pp. 452–465. ISSN: 0002-8282. URL: https://www.jstor.org/stable/1823886 (visited on 11/02/2024).

Wang, Binglin (2023). *Public Transit Provision and Fare Structure in U.S. Cities*. URL: https://docsend.com/view/iahctpp37e9k7epd.

West, Sarah E. (2004). "Distributional effects of alternative vehicle pollution control policies". In: *Journal of Public Economics* 88.3, pp. 735–757. ISSN: 0047-2727. DOI: 10.1016/S0047-2727(02)00186-X. URL: https://www.sciencedirect.com/science/article/pii/S004727270200186X.

Yang, Jun, Avralt-Od Purevjav, and Shanjun Li (2020). "The marginal cost of traffic congestion and road pricing: Evidence from a natural experiment in Beijing". In: *American Economic Journal: Economic Policy* 12.1, pp. 418–453. URL: https://www.aeaweb.org/articles?id=10.1257/pol.20170195.

Table 1: Summary Statistics of Trip Cost and Duration by Travel Mode

Mode	Variable	Obs	Mean	Std. Dev.	Min	Max
FHV	Cost	51,976	25.79	24.05	5.30	199.62
	Duration (min)	51,976	15.35	16.93	1	424
Private Vehicle	Cost	22,767	11.88	9.94	1.25	30.66
	Duration (min)	22,767	21.52	73.49	0	4,196
Bus	Cost	56,274	2.75	0.00	2.75	2.75
	Duration (min)	56,274	19.85	72.14	0	4,196
Subway	Cost	56,274	2.75	0.00	2.75	2.75
	Duration (min)	56,274	19.85	72.14	0	4,196
Rail	Cost	54,740	5.37	1.27	5.00	15.00
	Duration (min)	54,740	19.32	72.49	0	4,196
Bike	Cost	52,125	2.68	2.27	0.00	14.06
	Duration (min)	52,125	20.03	72.78	0	4,196
Walk	Cost	54,883	0.00	0.00	0.00	0.00
	Duration (min)	54,883	19.82	71.51	0	4,196

Notes: Statistics are from the 2022 NYCCMS matched to Google Maps travel times. Costs are in 2025 USD. Private vehicle statistics are reported only for households that own a vehicle; for non-owners, driving cost is set to 1000 to exclude it from the choice set. FHV costs include base fare, per-mile and per-minute charges, and congestion surcharges. Commuter rail costs reflect published fares. Bus and subway costs are fixed at \$2.75. Bike costs are zero for private bicycle owners and follow Citi Bike charges otherwise. Walking has zero monetary cost. Durations are reported in minutes; large standard deviations reflect trip length heterogeneity and congestion. Differences in the number of observations across modes arise because the choice set is not uniform: the minimum available set is 3 modes and the maximum is 7 modes per trip.



Figure 1: NYC Congestion Relief Zone

Note: Vehicles entering the area south of 60th Street in Manhattan are subject to the congestion charge. The West Side Highway, FDR Drive, and Hugh L. Carey Tunnel are excluded from the tolling zone.



Figure 2: Survey zones in the NYCCMS

Note: The figure shows the ten survey zones defined in the NYCCMS. Each zone aggregates several adjacent neighborhoods into larger geographic units, such as the Manhattan Core, Northern Bronx, and Outer Brooklyn. Nearly the entire Manhattan Core zone falls within the CBD congestion tolling area, which is the focus of NYC's pricing policy. Trips that originate or end outside NYC are grouped into an additional "Outside Region" category, not shown on the map. These zones serve as the spatial framework for constructing origin–destination pairs and mode attributes in the analysis.

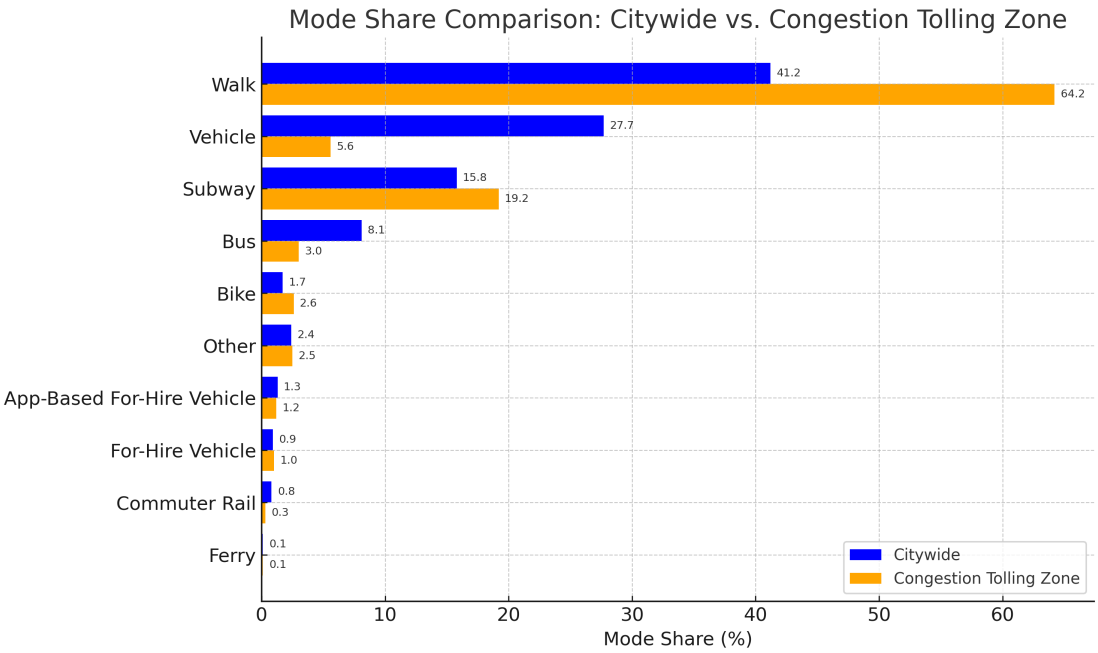


Figure 3: Mode shares: Citywide vs. Congestion Tolling Zone (CTZ).

Notes: Shares are based on trips recorded in the 2022 NYCCMS. Citywide shares reflect all trips across the ten survey zones; CTZ shares are restricted to trips starting or ending in the Manhattan Core, which overlaps almost entirely with the congestion pricing area.

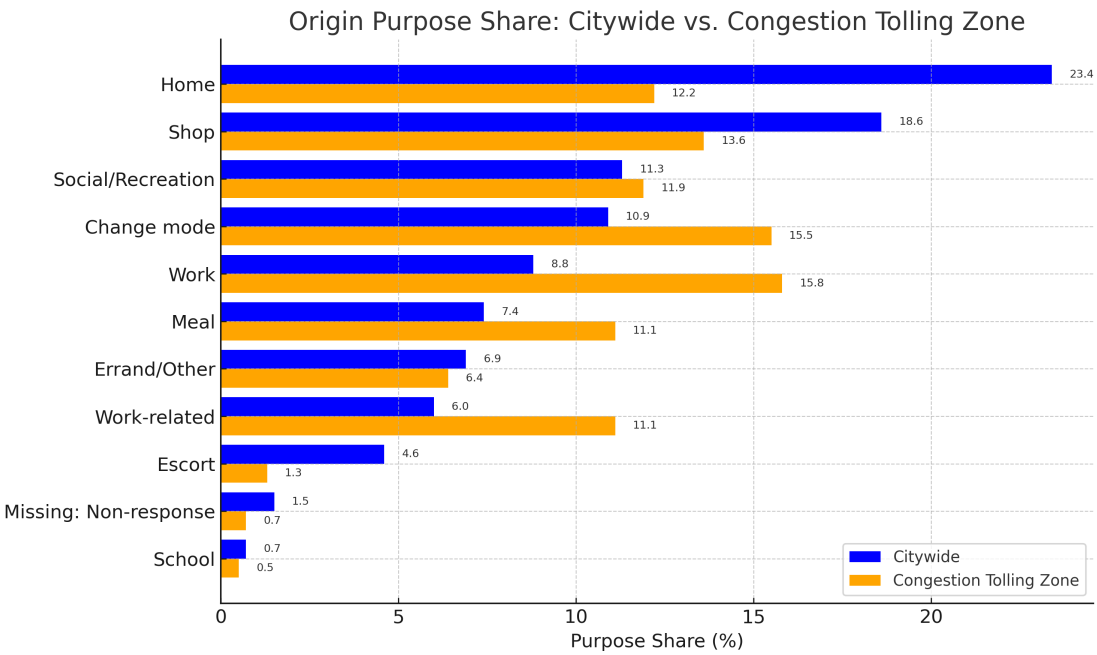


Figure 4: Origin purpose shares: Citywide vs. CTZ.

Notes: Trip purposes are drawn from the 2022 NYCCMS trip diary. Shares represent the percentage of trips by reported origin category. CTZ origins are disproportionately work- and transfer-related compared to citywide origins.

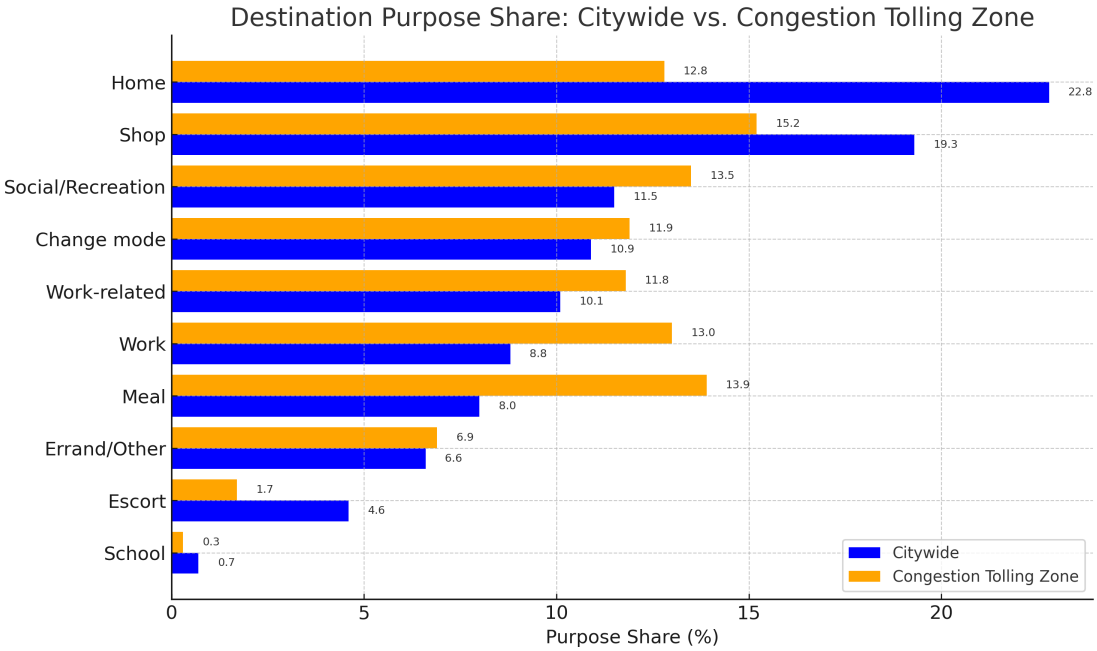


Figure 5: Destination purpose shares: Citywide vs. CTZ.

Notes: Trip purposes are based on the reported destination category in the 2022 NYCCMS. Work and work-related activities dominate CTZ-bound trips, while home-based trips are more common citywide.

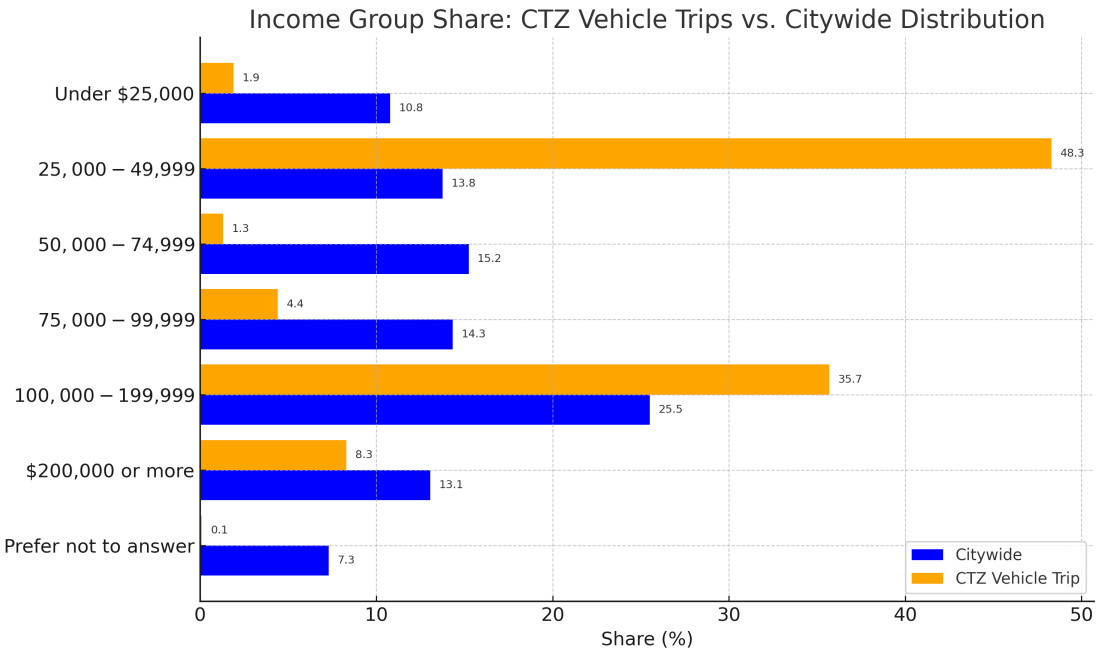


Figure 6: Income distribution of vehicle trips: Citywide vs. CTZ.

Notes: Household income groups are taken from the 2022 NYCCMS. Shares show the distribution of private vehicle trips by income. CTZ-bound vehicle trips are disproportionately concentrated among households earning $25,000–$49,999.

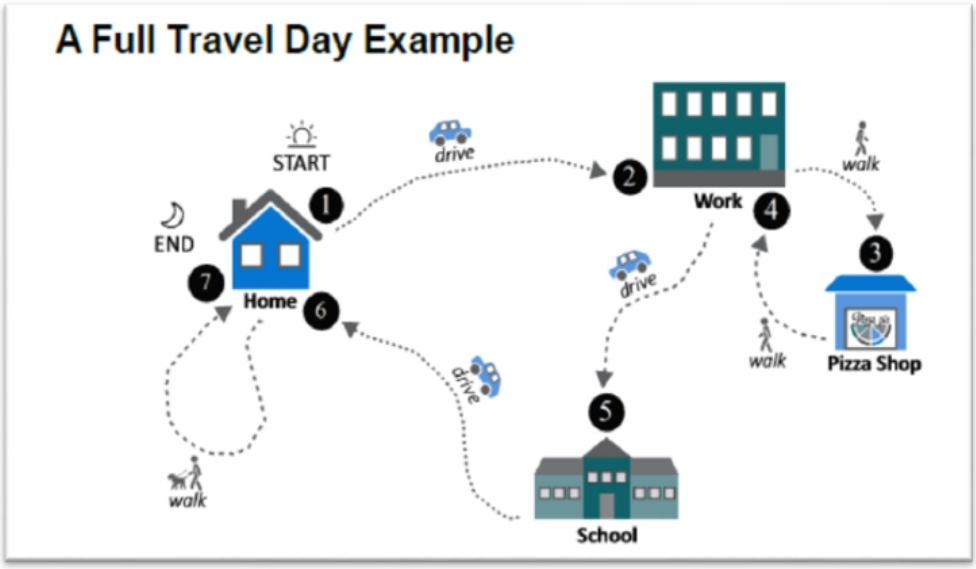


Figure 7: Travel Day Example (Source: NHTS 2017)

Notes: The figure illustrates a home-based tour consisting of seven trips. Trips 1–2–6–7 form the main tour between home, work, and school. Trips 3 and 4 represent a sub-tour where the traveler leaves work to visit a restaurant and then returns. The example highlights how tours are defined as closed loops, and how sub-tours may occur within them.

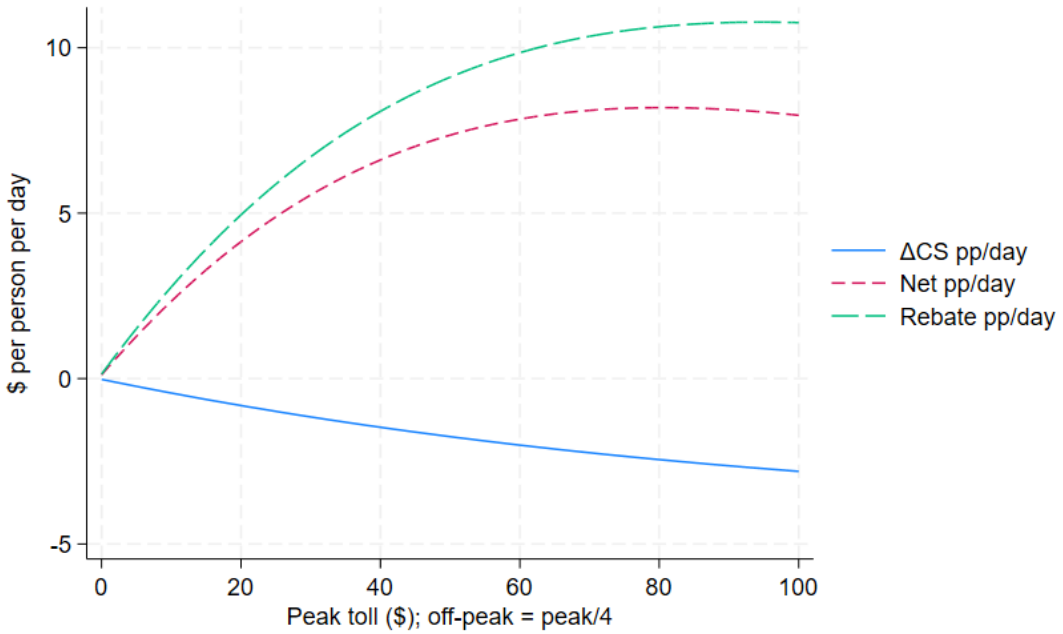


Figure 8: Per-Person Per-Day Welfare Effects under Alternative Toll Levels

Notes: This figure shows the per-person per-day welfare effects of varying the peak toll from \$0 to \$100 (with the off-peak toll set to one-quarter of the peak). The solid blue line represents the change in consumer surplus (ΔCS), the dashed red line shows the net welfare effect (ΔCS plus rebates), and the dashed green line shows the rebate amount. Results are estimated using the congestion feedback model and weighted by survey person weights. Net welfare increases with the toll level and reaches a maximum at approximately \$80 before declining, reflecting the balance between rising revenues and growing efficiency losses from higher tolls.

Table 2: MNL Model Results for Travel Mode Choices

	(1)	(2)	(3)	(4)
Cost ($)	-0.0339***	-0.0261***	-0.0166**	-0.0125**
	(0.00687)	(0.00708)	(0.00751)	(0.00632)
Duration (min)	-0.0551***	-0.0561***	-0.0559***	-0.0570***
	(0.00258)	(0.00265)	(0.00268)	(0.00269)
Mode Constants (Bus as baseline)				
FHV	-0.646***	-1.691***	-0.957***	-1.875***
	(0.0851)	(0.633)	(0.203)	(0.698)
Rail	-2.666***	-4.415***	-2.679***	-4.183***
	(0.157)	(0.838)	(0.589)	(0.936)
Subway	0.577***	-0.425*	0.0375	-0.849***
	(0.0537)	(0.243)	(0.129)	(0.252)
Vehicle	2.333***	1.100***	2.011***	1.103***
	(0.0664)	(0.281)	(0.147)	(0.302)
Bike	-1.558***	-3.361***	-3.319***	-5.597***
	(0.103)	(0.604)	(0.364)	(0.714)
Walk	2.232***	1.693***	1.974***	1.607***
	(0.0509)	(0.208)	(0.103)	(0.208)
Individual Demographics	No	Yes	No	Yes
Resident Location FE	No	No	Yes	Yes
Observations	379,877	379,877	379,877	379,877
Cases	56,274	56,274	56,274	56,274
VOT ($/hr)	36.89	27.91	17.81	13.20
Log-likelihood	-23570368.8	-21704565.0	-22166664.0	-21482592.9

Notes: Bus is the reference category. Dependent variable is travel mode choice, estimated using person-day weights. Standard errors are clustered at the trip level and reported in parentheses. VOT (value of time) is calculated as the ratio of the travel time coefficient to the cost coefficient, converted to dollars per hour. Demographics include household income, age, gender, race, education, household size. Also include weekend FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: MNL Estimates with Income Interactions

	Income			
	Cost	Duration	VOT	VOT(normalize)
Low	-0.0086***	-0.0539***	9.57	0.7658
	(0.0006)	(0.0029)		
Lower-middle	-0.0177	-0.0575***	18.51	0.5001
	(0.0131)	(0.0060)		
Upper-middle	-0.0255***	-0.0573***	26.66	0.3508
	(0.0081)	(0.0033)		
High	-0.0217*	-0.0608***	21.47	0.1113
	(0.0120)	(0.0071)		
Observations	379,877			
Case	56,274			

Notes: All specifications include alternative-specific (mode) fixed effects, individual demographic variables, residential location fixed effects, and a weekend indicator. Models are estimated with person-day weights; standard errors are clustered at the trip level. Income groups are defined as: Low <\$50k, Lower-middle = \$50–99k, Upper-middle = \$100–249k, High $\geq$ \$250k. VOT computed as the ratio of the travel-time coefficient to the cost coefficient (in \$ per hour) from the interaction MNL. VOT (normalize) is calculated based on estimated hourly wages, using median household income for each income group from the 2022 IPUMS ACS sample of NYC residents, assuming 240 workdays per year and 8 work hours per day. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: MNL Estimates with Age & Geographic Groups

	Age			Zone		
	Cost	Duration	VOT	Cost	Duration	VOT
18–34	-0.0193**	-0.0565***	25.24			
	(0.0080)	(0.0034)				
35–54	-0.0174***	-0.0593***	18.31			
	(0.0058)	(0.0027)				
55+	-0.0082***	-0.0517***	10.12			
	(0.0020)	(0.0057)				
Outer zones				-0.0289***	-0.0596***	31.75
				(0.0056)	(0.0027)	
Denser zones				-0.0241***	-0.0535***	33.66
				(0.0094)	(0.0038)	
Observations			379,877			
Case			56,274			

Notes: All specifications include alternative-specific (mode) fixed effects, individual demographic variables, residential location fixed effects, and a weekend indicator. Models are estimated with person-day weights; standard errors are clustered at the trip level. VOT computed as the ratio of the travel-time coefficient to the cost coefficient (in $ per hour) from the interaction MNL. Estimation uses person-day weights; standard errors are delta-method estimates with trip-level clustering. Zone groups follow the same definition as in the parking cost construction: Outer zones include Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn, and Staten Island, while Denser zones include Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, and Inner Brooklyn.

Table 5: Semi-Elasticities of Mode Choice with Respect to a 1% Increase in Travel Cost

	FHV	Rail	Bus	Subway	Vehicle	Bike	Walk
FHV	-0.3135	0.0005	0.0019	0.0073	0.0194	0.0017	0.0000
Rail	0.0087	-0.0864	0.0031	0.0079	0.0304	0.0011	0.0000
Bus	0.0043	0.0006	-0.0366	0.0053	0.0163	0.0006	0.0000
Subway	0.0089	0.0016	-0.0181	-0.0299	0.0225	0.0014	0.0000
Vehicle	0.0032	0.0004	0.0015	0.0037	-0.0592	0.0005	0.0000
Bike	0.0155	0.0004	0.0011	0.0089	0.0269	-0.0284	0.0000
Walk	0.0054	0.0003	0.0024	0.0060	0.0173	0.0008	0.0000

Notes: Each entry reports a semi-elasticity, defined as the percent change in the predicted probability of choosing a given mode (row) in response to a 1% increase in the travel cost of another mode (column). Diagonal values represent own-cost semi-elasticities and are negative, consistent with the disutility of higher monetary costs. Off-diagonal values capture cross-cost substitution patterns across modes. Elasticities are derived from the coefficients of the MNL estimates in column (4) of Table 2, applying person-day weights and clustering standard errors at the trip level.

Table 6: Semi-Elasticities of Mode Choice with Respect to a 1% Increase in Duration

Mode	FHV	Rail	Bus	Subway	Vehicle	Bike	Walk
FHV	-0.7100	0.0064	0.0441	0.2054	0.1265	0.0461	0.5335
Rail	0.0216	-0.7716	0.0856	0.2685	0.2345	0.0282	0.4030
Bus	0.0092	0.0035	-0.6111	0.0687	0.0709	0.0061	0.2831
Subway	0.0206	0.0075	0.0483	-0.4770	0.1506	0.0334	0.4549
Vehicle	0.0075	0.0041	0.0294	0.0880	-0.2585	0.0138	0.1947
Bike	0.0397	0.0083	0.0416	0.3277	0.2408	-0.5076	0.6207
Walk	0.0117	0.0017	0.0303	0.0688	0.0502	0.0091	-0.3195

Notes: Each entry reports a semi-elasticity, defined as the percent change in the predicted probability of choosing a given mode (row) in response to a 1% increase in the travel duration of another mode (column). Diagonal values represent own-duration semi-elasticities and are negative, consistent with the disutility of longer travel times. Off-diagonal values reflect substitution patterns between modes. For example, the entry of 0.2054 in row "FHV" and column "Subway" implies that a 1% increase in subway travel time is associated with a 0.205% increase in the predicted probability of choosing FHV. All elasticities are calculated based on the coefficients from the MNL estimates in column (4) of Table 2, using person-day weights and clustering standard errors at the trip level.

Table 7: Simulated Mode Shares for All Trips

	Baseline	Toll Only		Toll + Transit Discount	
Mode		No Feedback	With feedback	$0.35 Citywide	$0.90 CBD-trip
FHV	1.86%	1.86%	1.86%	1.86%	1.85%
		(0.00)	(0.00)	(-0.01)	(-0.01)
Rail	0.51%	0.51%	0.51%	0.51%	0.51%
		(0.00)	(0.00)	(0.00)	(0.00)
Bus	5.09%	5.10%	5.10%	5.12%	5.11%
		(+0.01)	(+0.01)	(+0.03)	(+0.02)
Subway	15.58%	15.62%	15.61%	15.68%	15.68%
		(+0.04)	(+0.03)	(+0.10)	(+0.10)
Vehicle	23.22%	23.16%	23.17%	23.15%	23.16%
		(-0.06)	(-0.05)	(-0.07)	(-0.06)
Bike	2.99%	3.00%	3.00%	3.00%	2.99%
		(+0.01)	(+0.01)	(0.00)	(0.00)
Walk	50.74%	50.76%	50.75%	50.69%	50.69%
		(+0.02)	(+0.01)	(-0.04)	(-0.05)

Notes: The sample includes all trips in NYC. The first row for each mode reports simulated mode shares, and the second row shows the percentage-point change relative to the baseline. *Toll Only (No Feedback)* applies the congestion toll without adjusting for changes in traffic speeds, while *Toll Only (With Feedback)* incorporates the congestion feedback loop. The *Toll + Transit Discount* columns allocate toll revenues to reduce public transit fares, with either a $0.35 citywide discount or a $0.90 discount for bus and subway trips to or from the Manhattan Core. If revenues exceed the cost of the subsidy, the surplus is distributed as a uniform per-capita rebate. Standard errors are omitted as all results are deterministic model outputs.

Table 8: Simulated Mode Shares for CBD-Trips

	Baseline	Toll Only		Toll + Transit Discount	
Mode		No Feedback	With feedback	$0.35 Citywide	$0.90 CBD-trip
FHV	3.12%	3.08%	3.11%	3.08%	3.07%
		(-0.04)	(-0.01)	(-0.04)	(-0.05)
Rail	0.96%	0.97%	0.97%	0.96%	0.96%
		(+0.01)	(+0.01)	(0.01)	(0.00)
Bus	2.76%	2.78%	2.77%	2.82%	2.83%
		(+0.02)	(+0.01)	(+0.06)	(+0.07)
Subway	35.55%	35.93%	35.83%	35.93%	36.10%
		(+0.38)	(+0.28)	(+0.39)	(+0.55)
Vehicle	22.82%	22.10%	22.26%	22.22%	22.18%
		(-0.72)	(-0.56)	(-0.86)	(-0.64)
Bike	7.82%	7.91%	7.88%	7.87%	7.84%
		(+0.09)	(+0.06)	(+0.05)	(+0.02)
Walk	26.98%	27.22%	27.17%	27.12%	27.02%
		(+0.24)	(+0.19)	(+0.14)	(+0.04)

Notes: The sample includes trips to and from the Manhattan Core (CBD), excluding within-zone trips inside the CBD. The first row for each mode reports simulated mode shares, and the second row shows the percentage-point change relative to the baseline. *Toll Only (No Feedback)* applies the congestion toll without adjusting for changes in traffic speeds, while *Toll Only (With Feedback)* incorporates the congestion feedback loop. The *Toll + Transit Discount* columns allocate toll revenues to reduce public transit fares, with either a $0.35 citywide discount or a $0.90 discount for bus and subway trips to or from the Manhattan Core. If revenues exceed the cost of the subsidy, the surplus is distributed as a uniform per-capita rebate. Standard errors are omitted as all results are deterministic model outputs.

Table 9: Per Person Per Day Welfare Effects of Counterfactual Simulations

	Toll Only		Toll + Rebate		Toll + Transit Discount	
	No Feedback	With feedback	No Feedback	With feedback	$0.35 Citywide	$0.9 CBD-trip
DCS ($)	-0.394	-0.324	-0.394	-0.324	0.023	-0.027
Rebate ($)			2.53	2.57	0.28	0.26
Net welfare ($)	-0.394	-0.324	2.137	2.244	0.301	0.279

Notes: No Feedback applies the congestion toll without adjusting for traffic speed changes, while *With Feedback* incorporates the congestion feedback loop. *Toll + Rebate* allocates toll revenues as a uniform per-capita transfer. *Toll + Transit Discount* reduces transit fares with either a $0.35 citywide discount or a $0.90 discount for bus and subway trips to or from the Manhattan Core; any surplus is rebated uniformly. DCS is the per person per day change in consumer surplus under the toll and any fare discount. Rebate is the per person per day transfer. Net welfare is the sum of DCS and rebate.

Table 10: Per Person Per Day Welfare Changes by Group ($)

	Toll Only		Toll + Rebate		Toll + Transit Discount	
	No Feedback	With feedback	No Feedback	With feedback	$0.35 Citywide	$0.9 CBD-trip
By Income						
Low	-0.1966	-0.1709	2.3334	2.3991	0.4467	0.3228
Lower-middle	-0.3975	-0.3285	2.1325	2.2415	0.3261	0.2944
Upper-middle	-0.5334	-0.4337	1.9966	2.1363	0.1809	0.2040
High	-0.5431	-0.4309	1.9869	2.1391	0.1898	0.4076
By Income (normalized)						
Low	-0.0020	-0.0017	0.0233	0.0240	0.0045	0.0032
Lower-middle	-0.0013	-0.0011	0.0072	0.0076	0.0011	0.0010
Upper-middle	-0.0009	-0.0007	0.0033	0.0035	0.0003	0.0003
High	-0.0004	-0.0003	0.0013	0.0014	0.0001	0.0003
By Age						
18–34	-0.3339	-0.2764	2.1961	2.2936	0.4315	0.4141
35–54	-0.4774	-0.3813	2.0526	2.1887	0.2519	0.2242
55+	-0.3420	-0.2955	2.1880	2.2745	0.2448	0.2279
By Resident Zone						
Outer zones	-0.3753	-0.3166	2.1547	2.2534	0.3279	0.3964
Denser zones	-0.4099	-0.3308	2.1201	2.2392	0.2807	0.1806
By Gender						
Female	-0.3255	-0.2622	2.2045	2.3078	0.3794	0.3409
Male	-0.4277	-0.3602	2.1023	2.2098	0.2514	0.2448
By Race						
Asian	-0.5392	-0.4313	1.9908	2.1387	0.2576	0.2426
Black	-0.3215	-0.2717	2.2085	2.2983	0.3828	0.2278
White	-0.3409	-0.2837	2.1891	2.2863	0.2888	0.3527
By Employment						
Unemployed	-0.1821	-0.1510	2.3479	2.4190	0.3678	0.2955
Employed	-0.4820	-0.3963	2.0480	2.1737	0.2757	0.2755

Notes: Values are daily per-person welfare changes, weighted by survey person weights. Groups are defined by reported income, age, home location, and other demographics. Income groups are defined as: Low $< \$50k$, Lower-middle = \$50–99k, Upper-middle = \$100–249k, High $\geq \$250k$. For Income (normalized), welfare changes are scaled by estimated daily wages, computed using median household income for each income group from the 2022 IPUMS ACS sample of NYC residents, assuming 240 workdays per year. Zone groups follow the same definition as in the parking cost construction: Outer zones include Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn, and Staten Island; Denser zones include Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, and Inner Brooklyn.

Table 11: Per Person Per Day Welfare under Selected Toll Levels

Peak	**Off-Peak**	**Rebate**	**DCS**	**Net**
0	0.00	0.13884632	-0.02071	0.118134
1	0.25	0.4234312	-0.06359	0.359845
5	1.25	1.5144136	-0.23151	1.2829
9	2.25	2.5305431	-0.39377	2.136769
15	3.75	3.9177193	-0.62669	3.29103
25	6.25	5.8804957	-0.98767	4.892827
50	12.51	9.1122212	-1.75171	7.360508
80	20.00	10.636439	-2.44524	8.191199
81	20.25	10.655865	-2.46481	8.191055
100	25.01	10.760899	-2.79995	7.960948

Notes: Peak and Off-Peak denote the toll levels (in dollars). Rebate ($) is the per person per day lump-sum transfer distributed uniformly to all travelers, calculated from toll revenues including FHV charges. DCS is the per person per day change in consumer surplus from applying the toll, relative to the no-toll baseline. Net equals the sum of Rebate and DCS, representing the per person per day net welfare change. The net welfare effect rises with the toll and reaches a maximum at a peak toll of about $80, after which it declines. All results are estimated using the congestion feedback model, weighted by survey person weights.

A Theoretical First-Best and Second-Best Congestion Pricing

The theoretical benchmark for congestion pricing is a *first-best* toll equal to the marginal external cost of road use. In a static model, let q denote traffic volume and $C(q)$ the average cost of travel (in time or money). Each additional driver increases congestion and imposes an external cost on all other drivers. The efficient outcome requires setting a toll $\tau(q)$ such that:

$$\tau(q) = q \cdot C'(q), \tag{9}$$

where $C'(q)$ is the marginal increase in average cost caused by one more driver. This toll aligns private and social costs, eliminating the wedge created by the congestion externality. In principle, the toll would vary continuously by road segment, time of day, and prevailing traffic density.

In practice, such fine-tuned pricing is infeasible. Implementing the first-best toll requires real-time monitoring of traffic conditions, continuous price adjustments, and universal compliance. Even if technologically possible, political acceptability poses a binding constraint: minute-by-minute tolls are difficult to communicate, enforce, and justify to the public.

Consequently, most real-world congestion pricing schemes adopt a *second-best* design. Instead of marginal-cost pricing, policymakers use simpler rules such as fixed cordon tolls (London, Stockholm, New York), time-of-day charges (Singapore), or distance-based fees (Milan). These second-best instruments approximate the efficiency of first-best pricing but sacrifice precision for feasibility and political support.

NYC's cordon charge is an example of such a second-best policy. Drivers pay a fixed fee when entering the Manhattan CBD, regardless of time of day or current traffic conditions. While this design cannot perfectly internalize the congestion externality, it captures part of the efficiency gains, generates substantial revenue, and remains administratively manageable.

B Construction of Mode Attributes from Google Maps

To construct mode-specific attributes, I start by defining OD pairs based on the 2022 NY-CCMS. The survey records trips at the level of ten zones within
acnyc, plus an "Outside Region" category. I use these zones to form OD pairs and select representative points within each zone.

Representative OD points are chosen to approximate the average trip distance and travel time reported in the survey. For each zone pair and mode, I calculated average observed trip

distances from the NYCCMS. I then selected OD pairs in Google Maps that closely matched these averages. For transit modes, the chosen OD points were anchored to stations or hubs (e.g., Penn Station, Grand Central) to reflect actual traveler access. For walking and biking, OD points were restricted to locations with safe pedestrian and bicycle connections. This ensures that constructed routes are contextually realistic.

Travel time data were collected using the Google Maps API over a one-week period from February 4 to February 10, 2025. Each day, I requested the best route and estimated travel time for each OD pair corresponding to the previous day (for example, on February 8 I retrieved estimates for February 7). For each day, I recorded travel times for four distinct periods:

- Morning rush hour (6:00 – 10:00 AM, using 8:00 AM as the representative time, which typically captures inbound congestion).
- Midday (10:00 AM – 3:00 PM, using 1:00 PM).
- Evening rush hour (3:00 – 8:00 PM, using 6:00 PM, which typically captures outbound congestion).
- Nighttime to early morning (8:00 PM – 6:00 AM, using 11:00 PM).

This process produced a comprehensive dataset of travel time estimates by day of week, time of day, OD pair, and travel mode. Transit times were further decomposed into in-vehicle time, walking access and egress time, and wait time. FHV durations were not directly available; instead, I adjusted driving times using a regression of observed FHV trip durations on private vehicle durations in the NYCCMS, which accounts for pickup and wait times.

The resulting attributes—mode-specific travel times by OD pair and time-of-day—are merged with the 2022 NYCCMS trip data to create an individual-trip-mode choice set for estimation. Monetary costs (tolls, fares, fuel, depreciation, and wear-and-tear) are added separately, as described in Section 3.2.3. For the "Outside Region" category, which includes flows from New Jersey, Connecticut, and other neighboring areas, I later supplement the NYCCMS with trips from the NHTS to ensure that cross-regional commuting into NYC is properly represented in the mode choice framework.

C Cost Construction

Driving. The monetary cost of driving includes fuel, parking, tolls, and wear-and-tear. For a trip of distance d_{ij} miles between zones i and j, the cost is defined as:

$$Cost_{ij}^{driving} = FuelCost_{ij} + Parking_j + Tolls_{ij} + Wear_{ij}. \tag{10}$$

Fuel cost is calculated as:

$$FuelCost_{ij} = \frac{d_{ij}}{MPG_h} \times P_{gas}, \tag{11}$$

where MPG_h is the household-specific fuel economy and P_{gas} is the average gasoline price in New York State in 2025. Because the NYCCMS does not report household vehicle fuel economy, I impute MPG_h using regression coefficients from the 2017 NHTS , based on vehicle year and fuel type. The estimated coefficients from this regression (Table A13) are then used to impute household vehicle fuel economy in the NYCCMS sample. The gasoline prices used are obtained from the U.S. Energy Information Administration's (EIA) weekly series for New York State (U.S. Energy Information Administration, 2024). Wear-and-tear is applied on a per-mile basis using American Automobile Association (AAA) estimates.

Parking costs are assigned by zone (see Table A12). In outlying areas (Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn, Staten Island), I assume greater access to metered on-street parking. In denser areas (Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, Inner Brooklyn), I use commercial garage rates. Metered rates are sourced from NYC DOT, and garage rates are obtained from ParkMe by averaging the ten closest garages to each zone's centroid.

For-Hire Vehicles (FHVs). For-hire vehicle (FHV) costs are calculated using the fare structure of the NYC Taxi and Limousine Commission (TLC). The generalized cost includes a base fare, per-mile and per-minute charges, tips, taxes, and surcharges:

$$Cost_{ij}^{FHV} = BaseFare + \alpha \cdot d_{ij} + \beta \cdot t_{ij} + Tip + Tax + Surcharge_{CBD}, \tag{12}$$

where d_{ij} is distance, t_{ij} is travel time, α and β are the per-mile and per-minute charges, and $Surcharge_{CBD} = \$1.50$ for trips entering the CBD.

Public Transit. Public transit fares are based on published schedules: \$2.90 for subways and local buses, \$7.00 for express buses, and distance-based fares for commuter rail. For commuter rail trips, fares vary by travel distance: \$5 for intra-city trips (under 5 miles), \$9 for short trips (5–15 miles), \$15 for medium trips (15–30 miles), and \$22 for long trips (over 30 miles).

Bike and Walk. Walking is assumed to have zero monetary cost. For biking, the cost is zero if the household reports owning a personal bicycle. Otherwise, costs are based on the Citi Bike pricing structure, based on trip duration.

D Extended Nested Logit Models

This appendix presents the extended nested logit specifications that complement the benchmark model in Section 4. These alternative formulations are not included in the main text

for brevity, but they provide robustness checks and illustrate richer substitution patterns across travel modes.

D.1 Baseline Nested Logit Model

Building on the consistency-in-driving assumption, and following Brian Y. He et al. (2020) and Wang (2023), I apply a two-level NL model to estimate travelers' mode decisions at the tour level. The structure of the model is illustrated in Figure 9.

In the first stage, the individual chooses between driving a personal vehicle and non-driving. If the non-driving option is selected, the second stage involves choosing among public transit (subway, commuter rail, local bus, etc.), for-hire vehicles (taxis, ride-hail, car services), walking, biking (personal or shared), and other modes.

The expected utility for the non-driving nest is defined as

$$E(V_{\text{non-driving},it}) = \frac{1}{\theta} \sum_{k=1}^{K_t} \ln \left(\sum_{j \in C_{it}} \exp(\theta V_{itj}) \right), \tag{13}$$

where k indexes trips, K_t is the number of trips in tour t, and $1/\theta$ is the inclusive-value scale parameter that links the upper and lower nests.

For the driving nest, there is only one alternative for the entire tour. The utility is

$$V_{\text{driving},it} = \sum_{k=1}^{K_t} \left(\alpha_i \cdot Cost_{\text{driving},it} + \beta_i \cdot Time_{\text{driving},it} + \delta_{\text{driving},i} + \mathbf{x}'_{\text{driving},it} \kappa \right). \tag{14}$$

For non-driving modes (excluding public transit), the utility function is

$$V_{itj} = \alpha_i \cdot Cost_{itj} + \beta_i \cdot Time_{itj} + \delta_{ij} + \mathbf{x}'_{itj} \kappa. \tag{15}$$

For public transit, I extend the specification to account for transfers:

$$V_{\text{PT},it} = \alpha_i \cdot Cost_{\text{PT},it} + \beta_i \cdot TransitTime_{\text{PT},it} + \phi_i \cdot F_{it} + \delta_{ij} + \mathbf{x}'_{itj} \kappa, \tag{16}$$

where $TransitTime_{\text{PT},it}$ includes schedule delay, access and egress time, waiting time, and in-vehicle ride time. The variable F_{it} is a dummy that equals one if the tour involves at least one transfer.

This nested structure relaxes the independence of irrelevant alternatives (IIA) assumption across the broad categories of driving and non-driving while still retaining computational tractability.

D.2 Alternative Nested Logit Structures

In addition to the baseline specification that separates driving from non-driving, more detailed nesting structures can be considered. These specifications group alternatives that are likely to share unobserved factors or to serve as closer substitutes.

Figure 10 illustrates one such alternative, in which travel modes are divided into five broad nests: *Driving Personal Car*, *Other Driving*, *Public Transit*, *Outdoor Modes*, and *Other Modes*. Within the driving category, personal car use is separated from carpools, carshare, and for-hire vehicles, reflecting different costs and availability. The public transit nest includes rail, bus, and subway, which are all operated under coordinated networks and may share correlated attributes such as service frequency, reliability, or crowding. The outdoor modes nest includes walking and biking, which are both active modes and often act as substitutes for short-distance trips. Figure **??** shows another potential nested structure.

The rationale for adopting such structures is that standard MNL models impose the independence of irrelevant alternatives (IIA) across all choices. In reality, some modes are closer substitutes than others: for example, bus and subway services may compete for similar riders, and walking and biking often represent alternative active travel options. By allowing these correlations within nests, the NL model can better capture realistic substitution patterns and improve the credibility of policy simulations.

D.3 Advanced Model

The baseline NL frameworks model travel behavior as a sequence of mode choices conditional on fixed origins and destinations. While this captures substitution across travel modes, it abstracts from the fact that, in many cases, travelers make joint decisions over both *mode* and *destination*.

For home-based tours, the origin and final destination are fixed at home, so individuals choose only among travel modes. In contrast, for non-home-based tours, individuals may simultaneously decide *where* to travel and *how* to get there. For example, a traveler leaving work for a meal might consider both alternative restaurants (destinations) and alternative travel modes, where destination attributes influence the available set of feasible modes.

This extension allows the utility specification to incorporate both mode j and destination $dest$. The indirect utility for individual i on trip t can be written as:

$$\max_{(j,dest)\in C_{it}} U_{itj} = \theta_{ij} + \alpha_i \cdot Cost_{itj}(dest) + \beta_i \cdot Time_{itj}(dest) + \delta_{jc} + \mathbf{x}'_{itj}\kappa + \varepsilon_{itj}. \tag{17}$$

Because travel costs and times are functions of traffic conditions, the utility can also be expressed with the congestion feedback model:

$$U_{itj} = \theta_{ij} + \alpha_i \cdot Cost_{itj}(v, d(dest), \text{etc.}) + \beta_i \cdot Time_{itj}(v(dest)) + \delta_{ij} + \mathbf{x}'_{itj}\kappa + \varepsilon_{itj}. \tag{18}$$

This advanced specification integrates destination choice with mode choice, so that congestion pricing policies may shift not only *how* people travel but also *where* they go. Moreover, the feasible choice set varies systematically across regions and hours of the day. For example, local bus service may be unavailable at night while subway service continues, altering the set of available options.

D.4 Estimation of the Nested Logit Model

Following the notation from (Brownstone and Kenneth A. Small, 1989), the probability of individual i choosing travel mode j within the nest s (driving a personal car, non-driving; or driving a personal car, other driving, public transit, outdoor mode, others) for tour t is

$$P_{itj} = P_{it}(j|s)P_{it}(s), \tag{19}$$

where

$$P_{itj}(j|s) = \frac{exp(V_{itj}/\rho_s)}{\sum_{b\in B_s} exp(V_{itb}/\rho_s)} = \frac{exp(V_{itj}/\rho_s)}{exp(I_s)}, \tag{20}$$

$$P(s) = \frac{exp(\rho_s I_s)}{\sum_r exp(\rho_r I_r)}, \tag{21}$$

$$I_r = log \sum_{b\in B_s} exp(V_{itb}/\rho_r), \tag{22}$$

$$V_{itj} = U_{itj} - \epsilon_{itj} \tag{23}$$

where r represents the different nests (e.g., driving alone, other driving, public transit, outdoor modes, others), and b includes the specific travel modes within each nest, such as subway, for-hire vehicles, local bus, walking, biking, and other modes.

One approach is to estimate the parameters in two stages: first, estimate parameters by treating the choice within nests as a simple MNL model, compute the inclusive value (IV)s for all nests, then estimate the parameters by treating the travel mode among nests as a MNL model. The second approach is applying the full information maximum likelihood (FIML) in a single phase. While FIML is more computationally intensive, it also captures potential dependencies across nests more effectively. I will try both approaches. I also plan to simulate the choice probabilities with Monte Carlo integration, first assuming random taste. Then a heterogeneous taste of commuting times, represented by β_j can be included

D.4.1 Nested Logit Results

Table A14 reports the estimation results from several NL specifications. Columns (1) and (2) adopt the baseline nesting structure, where the first-level decision is between driving a

personal vehicle and all other non-driving options. Column (3) implements the alternative specification shown in Figure 11, where the first-level nests are defined as *Driving* (private vehicles and FHVs), *Transit* (bus, rail, subway), and *Active* (walking and biking).

Across all specifications, the estimated cost coefficients are negative and significant, consistent with theoretical expectations. Duration terms are also negative, though their magnitudes vary by nest. In the three-nest specification (column 3), the sensitivity to duration is lowest for driving modes (-0.0117), somewhat larger for transit (-0.0191), and highest for active modes (-0.122), a pattern that aligns with intuition about mode substitutability. Travelers are relatively tolerant of driving time, less tolerant of transit time, and least tolerant of walking or biking time.

While the NL models provide internally consistent results, the predicted mode shares from these specifications do not perform well. In practice, either private vehicle or walking dominates the predicted probabilities, while other modes are substantially under-predicted. This problem appears in both the baseline and the alternative nested structures, suggesting that the nesting designs may not adequately reflect how travelers actually make mode choice decisions in NYC. In particular, the dichotomy between driving and non-driving or the grouping of all public transit modes into a single nest may be too coarse to capture meaningful substitution patterns.

For these reasons, I use the NL results primarily for robustness checks and interpretation of substitution patterns across related modes, but not for the main counterfactual simulations. The MNL framework, while more restrictive in terms of the IIA assumption, delivers more accurate in-sample predictions of mode shares and therefore provides a more reliable basis for policy simulations.

E Additional Tables and Figures

Table A12: Average Hourly Parking Rates by Survey Zone

Survey Zone	**Public (Meter)**	**Private (Garage)**
1. Northern Bronx	\$1.50	–
2. Southern Bronx	–	\$17.17
3. Northern Manhattan	–	\$13.20
4. Manhattan Core	–	\$25.63
5. Inner Queens	–	\$11.50
6. Middle Queens	\$1.25	–
7. Outer Queens	\$1.25	–
8. Inner Brooklyn	–	\$21.00
9. Outer Brooklyn	\$1.25	–
10. Staten Island	\$1.25	–

Notes: Public (meter) rates are sourced from the NYC DOT. Private garage rates are computed using ParkMe, averaging the ten closest garages to each zone centroid. Zones are defined by the 2022 NYCCMS. In outer zones (Northern Bronx, Middle Queens, Outer Queens, Outer Brooklyn, Staten Island), metered street parking is assumed to dominate, while in denser zones (Southern Bronx, Northern Manhattan, Manhattan Core, Inner Queens, Inner Brooklyn), private garage parking is assumed to dominate. These values represent average hourly rates in 2025 dollars.

Table A13: Regression Results for Predicting Household Vehicle Fuel Economy

VARIABLES	(1) **All Observations**	(2) **NYC Urban Sample**
Vehicle Year		
2000–2009	0.406***	-1.138*
	(0.115)	(0.600)
2010–2014	3.058***	1.494**
	(0.120)	(0.614)
2015–2019	4.051***	2.152***
	(0.137)	(0.631)
Fuel Type		
Gasoline	-2.995***	-0.519
	(1.086)	(1.867)
Electric & Hybrid	19.38***	20.72***
	(1.176)	(2.853)
Diesel	-7.530***	-0.196
	(1.106)	(2.667)
Constant	23.70***	23.02***
	(1.084)	(1.927)
Observations	256,115	7,137
R-squared	0.201	0.228

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is household vehicle fuel economy (miles per gallon). Column (1) uses all 2017 NHTS observations; Column (2) restricts to urban households in the NYC metropolitan area. All regressions are weighted using the NHTS household weight.

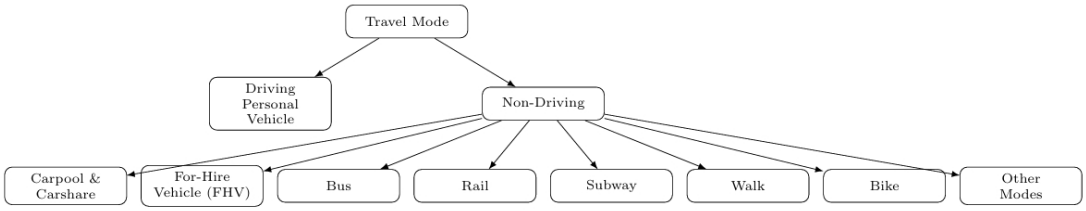


Figure 1: Updated Nested Structure of Travel Mode Choice

Figure 9: Structure of Baseline NL Model

Notes: The figure illustrates the two-level NL structure. At the upper level, travelers first decide between driving a personal vehicle and non-driving modes. Conditional on selecting the non-driving nest, they then choose among for-hire vehicles (FHVs), carpool and carshare, bus, rail, subway, walk, bike, or other modes. This structure relaxes the independence of irrelevant alternatives (IIA) across the broad categories of driving versus non-driving, while maintaining tractability within each nest.

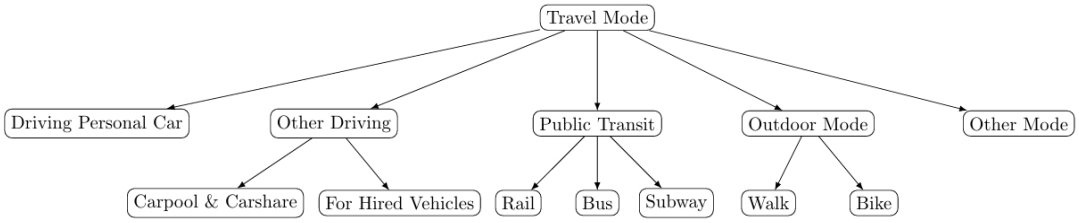


Figure 10: Alternative NL structure of travel mode choice

Notes: Compared to the baseline driving vs. non-driving division, this specification groups related modes such as bus and subway within public transit, and walking and biking within outdoor (active) modes, allowing substitution patterns to better reflect observed behavior.

Table A14: NL Estimation Results for Travel Mode Choices

Variables	**(1)**	**(2)**	**(3)**
Top-level cost and duration terms			
Cost ($)	-0.0030***	-0.0030***	-0.0683***
	(0.0000)	(0.0000)	(0.0023)
Duration (min)	-0.0450***	-0.0541***	
	(0.0015)	(0.0015)	
First-nest duration (Figure 11)			
Driving: Duration (min)			-0.0117***
			(0.0022)
Transit: Duration (min)			-0.0191***
			(0.0015)
Active: Duration (min)			-0.1220***
			(0.0018)
Mode FE	Yes	Yes	Yes
Resident Location FE	Yes	Yes	No
Individual Demographics	Yes	Yes	Yes
Weekend FE (first level)	No	Yes	Yes
Observations	301,035	103,364	103,364
Cases	43,005	43,005	43,005

Notes: Columns (1) and (2) estimate a baseline NL with two top nests (Driving vs. Non-driving). Column (3) estimates the alternative structure in Figure 11, with top nests *Driving* (private vehicle and FHV), *Transit* (bus, rail, subway), and *Active* (bike, walk). Duration and weekend fixed effects are estimated at the first (nest) level where indicated. All specifications include alternative-specific (mode) fixed effects, individual demographics (household income, age, gender, race, education, household size), and—except in column (3)—residential location fixed effects. Models are estimated with person-day weights; standard errors (in parentheses) are clustered at the trip level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

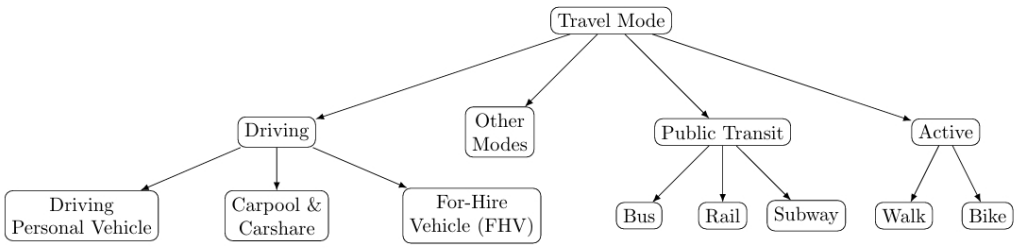


Figure 11: Alternative NL structure of travel mode choice

Notes: This nesting structure allows for correlations among related alternatives. For example, bus, rail, and subway services are grouped in the Public Transit nest, reflecting shared service characteristics and infrastructure. Walking and biking are grouped as Active modes, which are often substitutes for short-distance trips. Compared to a simple driving vs. non-driving specification, this design relaxes the IIA assumption and provides a richer substitution pattern across mode groups.