The Short-Run Effects of Congestion Pricing in New York City*

Cody Cook[*3] Aboudy Kreidieh[*1] Shoshana Vasserman[*124]

Hunt Allcott[24] Neha Arora[1] Freek van Sambeek[2] Andrew Tomkins[1] Eray Turkel[1]

[1]Google Research
[2]Stanford University
[3]Yale University
[4]NBER

September 15, 2025

Preliminary and incomplete; please do not distribute

Abstract

We study the impacts of New York City's 2025 Central Business District (CBD) Tolling Program, the first cordon-based congestion pricing scheme in the United States. Using a synthetic controls approach, we find that the policy increased speeds on CBD roads by 10%, with little-to-no effect on air quality, transactions at shops and restaurants, or overall foot traffic in the CBD. Speeds also increased *outside* the CBD, especially on roads commonly traversed by drivers headed to the CBD. These spillovers lead to faster trips throughout the city, including for many unpriced trips. We develop a simple model to bound driver welfare gains. If drivers have a Value of Travel Time (VOTT) of $40/hour, then we estimate that driver welfare increased by at least $1.0 million/day, before any revenue recycling. Drivers to the CBD are only better off if their VOTT exceeds $176/hour, but the modest speed improvements for the many unpriced trips bring the overall break-even VOTT down to $29/hour. Finally, we show how characteristics of the local travel patterns and road networks in other cities can inform the potential impacts of introducing cordon-based congestion pricing.

JEL codes: **R41, R48, R5**

Keywords: congestion pricing, traffic, externalities, Pigouvian taxation

*Cook, Kreidieh and Vasserman are joint first authors (denoted with *). Shoshana Vasserman is the corresponding author (svass@stanford.edu). We are grateful to Andrew Stober, Carolina Osorio, Ivan Kuznetsov, Prem Ramaswami, Preston McAfee and Willa Ng for their support and helpful comments throughout the writing of this paper.

1 Introduction

Congestion pricing is a textbook application of Pigouvian taxation (Pigou, 1920), and, in theory, can achieve first-best allocations by internalizing congestion externalities (Vickrey, 1963, 1969). In practice, however, real-world programs are second-best: they price only a subset of roads, trips, or times. Whether such programs improve welfare depends on how pricing reshapes flows across the road network, and on how changes in trip flows affect the speeds of the roads they traverse. These spillovers are theoretically ambiguous, as traffic may reallocate toward untolled roads and worsen their congestion, or decrease in total volume and reduce congestion citywide (Verhoef, Nijkamp and Rietveld, 1996; De Palma and Lindsey, 2000; Verhoef, 2002).[1] Despite their potential importance, empirical evidence on the spillovers of second-best congestion pricing is scarce, in part because the effects may be diffuse and causal identification is challenging.

We contribute new evidence using the introduction of congestion pricing in New York City, the first cordon-based congestion pricing program in the United States. Under the policy, passenger vehicles must pay $9 to enter NYC's central business district (CBD) during peak hours. We measure the policy's impact using a generalized synthetic control (GSC) design (Xu, 2017) combined with data from NYC and five control cities. We find that the policy increased CBD speeds by 10%, with little-to-no effect on air quality, restaurant/retail transactions, or foot traffic. Moreover, the policy had positive spillovers on speeds throughout the NYC metropolitan area, with the largest gains on highways commonly traversed by drivers headed to the CBD. These spillovers are first-order for aggregate driver welfare, as the majority of the travel time savings accrue to drivers on unpriced trips. Using a simple model and valuing travel time at $40/hour, we estimate that total driver welfare rose by at least $1.0 million per day, before accounting for any revenue recycling and despite the fact that drivers entering the CBD pay tolls large enough that their own trips may be privately welfare-reducing unless they value travel time above $176/hour.

Our primary data include anonymized and aggregated statistics from trips taken with Google Maps, which we complement with information on air quality from PurpleAir sensors, spending at shops and restaurants from a panel of credit/debit cards, and foot traffic inferred from a sample of GPS devices. The sample covers the metropolitan areas corresponding to New York City, Philadelphia, Boston, Chicago, Atlanta, and Baltimore. For each outcome, we use the GSC approach to estimate the treatment effect of the policy by comparing the realized outcomes in NYC after the policy was introduced to a prediction of counterfactual outcomes, which is based on concurrent observations in other major US cities and a set of fixed effects and latent factors estimated using pre-policy data.

We begin by evaluating the effects on the now-priced region itself, where we find that the introduction of congestion pricing led to an immediate and persistent increase in speeds. Raw average speeds in the CBD during weekday peak hours rose from 7.1 miles per hour (mph) pre-policy to 8.2 mph in the six months following implementation. However, average speeds in most cities rose during the same period, so comparing raw average speeds conflates the treatment effect of the policy with any

[1] In a survey of economists before the policy's implementation, 90% of respondents agreed that congestion pricing in NYC would "lead to a substantial reduction in traffic congestion in the targeted area," but the majority were uncertain as to whether traffic would increase on roads just outside the CBD (Clark Center Forum, 2024).

concurrent trends in traffic conditions. Using the GSC approach, we estimate that the policy increased speeds in NYC's CBD by 10% relative to a synthetic control formed from other cities, suggesting that three-fourths of the observed change in speeds is attributable to the policy. The effects on speeds are larger during the afternoon—historically the most congested time of day—and persist even after peak-hours pricing ends at 9pm. Within a given hour, the effects are largest for the slowest roads in the CBD. Introducing congestion pricing also reduced the dispersion in speeds across roads, primarily by compressing the tail of extreme delays, which has potentially important welfare implications given driver preferences for reliability (Small, Winston and Yan, 2005; Bento, Roth and Waxman, 2020*a*).

Beyond road speeds, the policy had little-to-no effect on CBD air quality, visits to restaurants/shops, or overall foot traffic. Estimated fuel efficiency[2] improved by 2-3% for cars traveling in the CBD, but these improvements have yet to have a detectable treatment effect on ambient concentrations of fine particulates ($PM_{2.5}$). We also find no effect on the number of credit/debit card transactions at shops and restaurants in the CBD. Overall foot traffic in the CBD—inferred using location data from GPS devices—has also remained steady, suggesting that any reduction in car trips to the CBD has been offset by substitution toward other travel modes.

To evaluate effects beyond the CBD, we first develop a segment-level measure of exposure to congestion pricing based on its *co-occurrence* with the CBD, which we define as the share of drivers crossing a segment who are bound for the CBD. A road segment has a co-occurrence of 50% if half of the cars that typically traversed it prior to the policy eventually entered the CBD. If the policy's main effect were to reduce CBD trips by 10% uniformly, a segment with 50% co-occurrence would see a proportional 5% decrease in cars. Many drivers not bound for the CBD will still traverse segments exposed to the policy, and, through this exposure to CBD-bound drivers, may experience some changes in speed after the policy's implementation.

We find that the policy increased speeds throughout the NYC metropolitan area, with larger effects on roads with higher levels of co-occurrence with the CBD. For road segments with 80-100% co-occurrence—which include the Lincoln Tunnel, Holland Tunnel, and other main entries to the CBD—the policy's implementation increased speeds by an average of 15%. The effects then decline monotonically in the level of co-occurrence, but even speeds on the average road with 0-20% co-occurrence experienced a smaller 4% increase in speeds. Across road types, highways experienced larger increases in speeds—especially at lower levels of co-occurrence—while the effects on local and arterial roads have been more limited. The smaller effect on roads within the CBD than roads with 80-100% co-occurrence is likely largely due to the composition of road types. The CBD is comprised of local and arterial roads with auxiliary congestion from pedestrians and stopped vehicles—even at 3am, the average CBD speed is less than 15 mph—while roads leading into the CBD tend to be larger highways.

The changes in road speeds add up to meaningful differences in travel times between different parts of the city. For trips to the CBD, the implementation of congestion pricing increased average trip speeds by about 9%. While drivers on trips to the CBD must now pay a toll for these time savings,

[2]We infer fuel efficiency on road segments using an engineering model that transforms speed profiles and segment characteristics into an estimate of liters consumed per 100 kilometers traveled, similar to the approach in Brooker et al. (2015).

drivers on other trips within, from, and outside the CBD benefit from increased speeds without changes in price. Speeds on trips within the CBD—which, for passenger cars, do not pay any toll—increased by 6%, while speeds on trips leaving the CBD increased by 7%. Even trips that took place entirely outside the CBD experienced modest gains of 0.5-2%, depending on the distance of their origin and destination to the CBD.

Distributionally, the policy had similarly positive effects on speeds for Census tracts of different income levels and in specific areas, such as the Bronx, New Jersey, and Staten Island. We group Census tracts by the quintile of median household income and estimate the effects for each group of tracts separately. We find that the policy has increased average speeds on segments within tracts across the household income distribution. Tracts in the bottom income quintile saw speed increases of 3.2%, with smaller effects of 1-2% in the other income quintiles. Similarly, trips to the CBD originating in each income quintile have all seen speed gains of 7-9%. For New Jersey residents, the policy increased speeds on road segments within Hudson and Bergen counties by 4.7%, and increased speeds on trips to (from) the CBD by 9.8% (9.6%). For residents of the Bronx and Long Island, the policy also increased speeds on road segments within each area by at least 3% and increased speeds on trips to/from the CBD by 5-8%.

Next, we bound the welfare gains for drivers traveling throughout the metro area using a simple model that weighs the time savings for drivers against the new congestion fee. On the average weekday during peak hours, we estimate that the 274,000 passenger vehicles traveling to the CBD save 2.7 minutes and pay an average of \$7.9.[3] For these drivers, the policy is only welfare-improving if their Value of Travel Time (VOTT) is over \$176/hour. However, because of the spillovers in road speeds, many *unpriced* trips are also now faster. While we estimate that the average unpriced passenger vehicle trip saves just 8 seconds, there are far more unpriced trips than priced trips. Including unpriced trips, the 'break-even' VOTT of drivers for the policy is at most \$29/hour.[4] If all drivers have a VOTT of \$40/hour, then we estimate total daily welfare gains—before any revenue recycling—of at least \$1.0 million, which come almost entirely from unpriced trips.

Finally, we investigate the mechanisms behind the large effects in NYC, and whether similar results are plausible in other congested cities. The impact of cordon-based congestion pricing depends on three underlying factors: 1) how pricing changes traffic volumes; 2) how much other drivers are exposed to those volume changes; and 3) the shape of the *congestion functions* that describe the relationship between traffic density and speed on different roads. Intuitively, if most CBD-bound trips begin just outside of the cordon area, or if the roads they traverse are already at free-flow speeds, then few other drivers will be affected by any change in the number of CBD-bound trips. If, instead, CBD-bound trips originate from the outskirts of the metro area and travel on congested roads, then changes in CBD trip volumes will have larger spillovers on speeds for other drivers. To make the problem tractable in the absence of a demand model, we assume each city sets some price that reduces CBD-bound trips uniformly by 10%, and compute the effects on other drivers' travel times and the 'break-even' price that would make such a policy welfare-improving if drivers have a VOTT of \$40/hour.

[3]The average price paid is lower than the advertised \$9 toll due to crossing credits for entering the CBD through an already-tolled entrance, such as the Lincoln Tunnel.

[4]We repeat a similar exercise for taxi/FHV trips, who save similar amounts of time but at a lower per-trip fee.

To measure the exposure of drivers to CBD-bound trips, we use our segment-level co-occurrence measures. For each origin-destination (OD) pair in each city, we measure exposure as the average duration-weighted co-occurrence of segments traversed for trips. Trips to/from the CBD have average exposures of 55-75% in each city, with the largest values in NYC and Chicago and the smallest values in Atlanta and Baltimore. Trips outside the CBD are naturally less exposed on average ($< 10\%$ in all cities), but the values range from a lower end of 3% in Philadelphia and 4.8% in NYC up to 9.7% in Chicago and 9.3% in Boston. This suggests that the spillovers of cordon-based congestion pricing on non-CBD drivers may be smaller than NYC's spillovers in Philadelphia, but larger in Boston and Chicago.

To link changes to volumes to changes in travel times, we estimate congestion functions for groups of road segments using pre-policy data on speeds and corresponding density. We group road segments by their co-occurrence bin (i.e. 0-20%, 20-40%, ...) and road type (i.e. local, arterial, and highway) and estimate separate congestion functions for each set of roads in each city. Consistent with the traffic engineering literature, these functions are typically convex, especially on highways (Seo et al., 2017). As a result, the effects on speeds of a constant change in density depend on *where* along the congestion function pre-policy traffic typically operated. If the road is operating near a steep part of its congestion function, removing even a few cars can substantially increase speed.

We combine these ingredients to simulate the time savings and break-even price for a 10% decrease in CBD-bound volumes in each city. For each OD pair, we compute the average duration and average speed for segments of each co-occurrence bin and road type using pre-period trips. Then, for each group of segments, we use the estimated congestion function to evaluate the change in speeds that would occur following a 10% decrease in CBD-bound trips, scaled by the average co-occurrence of those segments (i.e. a 10% decrease for segments with an average co-occurrence of 50% would be a 5% decrease in volumes on those segments). We can then aggregate across all groups of segments and across OD pairs to compute the city-wide time savings, weighting each OD pair by its pre-period volumes. We find that... [Results pending.]

Our work most directly contributes to studies of cordon-based congestion pricing programs, including early evaluations of the New York City policy and a larger body of work on older policies in London, Stockholm, Milan, and elsewhere.[5] Since its launch, year-to-year comparisons in speeds and travel times in NYC's CBD have been reported by journalists (Gordon et al., 2025; Ley, Hu and Collins, 2025), the Congestion Pricing Tracker website (Moshes and Moshes, 2025), and the Metropolitan Transportation Authority (MTA, 2025*a*).[6] Perhaps due to having access to limited data covering subsets of NYC roads, however, these analyses occasionally disagree on whether the program decreased or increased travel times (Hu, Ley and Schweber, 2025). In contrast, we evaluate the causal effects of congestion pricing using a consistent econometric strategy and large-scale data covering roads in NYC and a set of comparison cities. Outside of NYC, the 2003 implementation of a congestion charge in

[5] A related set of studies conduct ex-ante simulations of potential congestion pricing policies in cities without any existing policy. See, for example, Almagro et al. (2024); Barwick et al. (2024); Durrmeyer and Martínez (2024); Hierons (2024); Ater et al. (2025); Montero, Sepúlveda and Basso (2025).

[6] Even before the policy's implementation, MTA (2022) and Hierons (2024) attempted to anticipate the policy's potential impacts on congestion, air pollution, and mode choice.

downtown London lead to reduced traffic, lower pollution, and fewer accidents (Leape, 2006; Green, Heywood and Paniagua, 2020; Tang and van Ommeren, 2022). Similar effects on downtown traffic conditions have been documented for congestion pricing in Stockholm (Eliasson et al., 2009; Simeonova et al., 2021) and Milan (Gibson and Carnovale, 2015).[7]

Evidence on the spillovers of congestion pricing is more mixed. Early theoretical work focused on the welfare effects of first-best pricing in a world with just a single road, setting aside any spillovers (Vickrey, 1963, 1969). Later, studies of second-best congestion pricing—where only a subset of roads are priced—incorporated the effects of spillovers, but were primarily focused on the *negative* impact on speeds of roads that were substitutes for a tolled road (Verhoef, Nijkamp and Rietveld, 1996; De Palma and Lindsey, 2000).[8] Empirically, Gibson and Carnovale (2015) finds that congestion pricing increased traffic on radial roads that were substitutes to traveling through downtown Milan, while Herzog (2024) finds that London's policy reduced traffic on certain untolled roads leading into the downtown area. For these two papers—which study programs introduced before the prevalence of Google Maps or other large data sources—the authors are limited to a handful of speed observations from only the focal city. Our more detailed data lets us detect diffuse effects on speeds throughout the metro area, which we show are first-order for welfare accounting, as the small speed improvements for each unpriced trip add up to large effects overall.

2 Congestion Pricing in NYC

New York City's congestion pricing policy, officially known as the Central Business District Tolling Program, was implemented on January 5th, 2025. The initiative imposes a fee on vehicles entering Manhattan south of 60th Street, excluding the FDR Drive and West Side Highway. The tolls are collected electronically and vary based on the time of day, vehicle type, and whether the vehicle is equipped with an E-ZPass transponder. Since the policy's launch, nearly 248,000 passenger cars and 143,000 taxis/FHVs have entered the CBD on the average weekday during peak hours.[9] The Metropolitan Transportation Authority (MTA) anticipates that the tolls will generate approximately $15 billion, which will be allocated to fund repairs and enhancements to the city's subway, bus, and commuter rail systems.

Toll rates are set at $9 per day for passenger cars and small commercial vehicles if paid by E-ZPass. Motorcycles pay $4.50 per day, while trucks and buses pay between $14.40 and $21.60 per day, depending on their size. These rates are reduced by 75% overnight, and are up to 50% higher if

[7] In lieu of congestion pricing, some cities impose driving restrictions and car-free zones, which can reduce downtown congestion and pollution but are inefficiently targeted and do not raise any revenue (Davis, 2008; Gallego, Montero and Salas, 2013; Yang, Purevjav and Li, 2020; Galdon-Sanchez et al., 2022; Sleiman, 2024).

[8] The classic example is a tolled highway running parallel to a free highway. In the US, many roads now have 'express lanes' that are dynamically priced in response to current traffic conditions. These lanes offer drivers the option to pay for faster speeds when needed, but can increase congestion in the remaining untolled lanes (Hall, 2018; Bento, Roth and Waxman, 2020*b*; Cook and Li, 2024).

[9] This statistic is computed based on the publicly-available "MTA Congestion Relief Zone Vehicle Entries" dataset (source) for weekday peak hour entries by "Cars, Pickups and Vans" or "TLC Taxi/FHV" between January 5th and June 30th, 2025. These two categories account for about 85% of all entries. See Table B.1 for counts by other vehicle types.

drivers do not have E-ZPass and instead pay by mail. Taxis and ridesharing vehicles pay a per trip rate of $0.75 for taxis and $1.50 for ridesharing vehicles for trips that start, end, or pass through the CBD.[10] There are a few exempted vehicles (e.g., emergency vehicles), and vehicles entering via certain bridges or tunnels that are already tolled receive a partial credit. Low-income residents can also apply to receive 50% off after their first ten trips in a month. The passenger car rate is set to increase to $12 in 2028 and $15 in 2031.

The concept of congestion pricing in New York City has a history spanning several decades. In the early 1970s, Mayor John Lindsay proposed tolling East River crossings to mitigate traffic congestion in Lower Manhattan, but the plan was eventually withdrawn after facing stiff opposition from the trucking industry, taxi drivers, and local businesses concerned about the financial impact of tolls. Subsequent mayors attempted to implement similar measures. In 2007, Mayor Michael Bloomberg's administration proposed the PlaNYC initiative, which aimed to introduce congestion pricing to promote sustainability and fund transportation infrastructure. While this proposal stalled in the state assembly, it resurfaced in 2017 when Governor Andrew Cuomo advocated for congestion pricing to raise funds for improving transit. This led to its eventual inclusion in the 2019 state budget and, after many delays, its implementation in 2025.

The policy faced criticism and legal challenges during its development. Critics argue that this approach imposes an additional financial burden on residents, may have adverse economic impacts on business, and will shift traffic and pollution to other parts of the city (Ley, 2022). At least ten lawsuits were filed against the MTA and state officials by business coalitions, elected officials from New Jersey, and others (Hu and Ley, 2024).

3 Data & Empirical Strategy

We build data covering traffic conditions, air quality, consumer spending, and foot traffic for New York City and five other cities, which we use to form synthetic controls. The five control cities are Philadelphia, Boston, Chicago, Atlanta, and Baltimore. For each city, we define its boundaries according to the corresponding Core Based Statistical Area (CBSA) and its CBD using the most prominent version of a CBD or 'downtown' drawn by a city government-affiliated organization.[11] In this section, we provide an overview of each data source; Appendix A contains additional details.

3.1 Google Maps Traffic Trends

Our primary data on traffic conditions come from anonymized and aggregated statistics from Google Maps Traffic Trends, which are based on trips taken with Google Maps during the period of the study. The data include two primary sets of statistics: 1) road segment-level outcomes, which we then

[10] Ostrovsky and Yang (2024) evaluate the pricing by vehicle type and argue that the small per-trip charge on taxis and ridesharing companies is too low, as a single trip by a taxi likely contributes as much congestion as a trip by a private vehicle.

[11] In NYC, the congestion pricing cordon area aligns with the NYC CBD boundaries defined by the City of New York before it released any plans for a lower-Manhattan congestion pricing policy (NYC Department of City Planning, 2011). Appendix A.1 documents the official sources we use to define the CBD shapes in control cities.

further aggregate in time and space across sets of road segments; and 2) origin-destination (OD) level outcomes, which are aggregated across trips based on OD Census tract characteristics. The sample covers traffic conditions from September 2024 through June 2025. Except where otherwise noted, all analyses focus on data from weekday priced hours (5am – 9pm).

Segment-level outcomes. For all segment-level analyses, we group segments by either geographic location, such as Census tracts, or by shared characteristics, such as co-occurrence with the CBD. Road segments vary in length, with an average of approximately 50 meters. For each segment-level outcome, we consider aggregates composed of the harmonic mean weighted by traversal distance. In order to maintain a stable composition of road segments within a group, we further weight each segment by its average share of traversals within hour t between September and December 2024. That is, for a given outcome y measured across segment group j in hour t, we consider the hourly distance-weighted average outcome:

$$\bar{y}_{j,t} = \frac{\sum_{s \in S_j} \sigma_{s,t} \times d_{s,t}}{\sum_{s \in S_j} \sigma_{s,t} \left(d_{s,t} / y_{s,t} \right)} \tag{1}$$

where S_j is the set of traversals through segments s in segment group j, $\sigma_{s,t}$ is the pre-policy segment weight, and $d_{s,t}$ is the traversal distance (equivalent to the length of each segment). We consider three types of outcomes: traversal speed, normalized traversal speed relative to the segment's speed limit, and the estimated fuel consumption rate. For fuel consumption, we collaborate with the *National Renewable Energy Laboratory* (NREL) to model how fuel consumption rate varies based on speed and the segment characteristics for a typical sedan. Appendix A.2 provides additional details on how we model fuel consumption.

Origin-destination (OD) level outcomes. Our second set of data cover full trips throughout the metro area, rather than individual road segments. Because of the more granular definition of these data, we aggregate trips to Census PUMAs rather than Census Tracts and to two-hour and four-hour bins. We define origins and destinations based on groups of Census PUMAs, and whether the trip starts and/or ends within the CBD. For each OD, we measure the realized travel times and trip speeds (e.g. travel time divided by trip length), aggregated across all trips within a given OD-hour-bin. In order to maintain a stable composition of trips between different PUMAs within a group, we weight each trip by the average trip volume in the same hour bin between September and December of 2024.

3.2 Other Outcomes

Air quality. We use air quality data from PurpleAir to estimate the effects of the policy on pollutant levels in New York City. PurpleAir is a company that sells $PM_{2.5}$ sensors as a consumer product. Its customers can opt to share the data their sensors gather with the public, and the many "community scientists" form a network of $PM_{2.5}$ monitoring with large coverage. We focus on the effects on $PM_{2.5}$ over other pollutants as this pollutant accounts for most of the adverse health effects of air pollution in the U.S. (Tschofen, Azevedo and Muller, 2019; National Institute of Environmental Health Sciences,

2024). For each city, we query the data of all outdoor PurpleAir sensors in the corresponding CBSA. The data include 22 unique outdoor PurpleAir sensors in the NYC CBD and 1,067 sensors throughout NYC and our five control cities. We calibrate the PurpleAir $PM_{2.5}$ data as recommended by the Environmental Protection Agency (EPA) (Barkjohn et al., 2022); we describe this process in Appendix Section A.3.

Consumer spending. To measure consumer spending, we use proprietary credit and debit card transaction data compiled by MBHS3 and provided through Yale University. The underlying data contain approximately 35 billion annual transactions by 180 million individuals. We receive data aggregated at the level of zipcode, day, and 3-digit North American Industry Classification System (NAICS) code for transactions between January 2024 and April 2025. We focus on transactions at restaurants or retail establishments, which we identify using their 3-digit NAICS code. For retail establishments, we restrict attention to five NAICS categories that had over a million observed transactions in the NYC CBD during 2024.[12] In NYC, the data include 165 million restaurant transactions ($4.6 billion total spend) and 286 million retail transactions ($13.2 billion total spend) during the sample period (Appendix A.4). The boundaries of NYC zipcodes almost exactly align with the CBD boundaries (Figure A.2), so we define a zipcode as treated by the policy if its centroid lies within NYC CBD.

Foot traffic. We use Advan Research's Neighborhood Patterns to measure the amount of foot traffic in the CBD over time.[13] The underlying data are passively collected from mobile applications on GPS-enabled devices (e.g., smartphones) that record a device's coordinates and timestamps whenever connected to the GPS. These data, which cover approximately 7% of the overall population (Li et al., 2024), are then stripped of personal identifiers and aggregated for research purposes. The aggregate data we receive include the number of visits to each Census tract in each hour. We further aggregate to the tract-day level by summing the number of visits to a tract during peak hours (5am to 9pm). We use data from January 2024 through June 2025, and define a tract as treated by the policy if its centroid lies within the NYC CBD (see Figure A.2 for a map).

3.3 Empirical Strategy: Generalized Synthetic Controls (GSC)

Assessing the effects of New York City's congestion pricing policy is complicated by numerous confounding factors. A straightforward comparison of average outcomes before and after the policy combines the policy's effects with any other time-varying factors influencing traffic conditions. Existing evidence generally compares average traffic patterns in NYC after implementation to earlier months or to data from the prior year. Interpreting the results from such approaches as causal requires an assumption that, absent the policy, NYC traffic conditions after January 5th would have been similar to earlier periods. We introduce a method that instead compares changes in traffic conditions in NYC to contemporaneous changes in a set of comparison cities.

[12] The categories are Food & Beverage Stores, Health & Personal Care Stores, Clothing & Clothing Accessories Stores, General Merchandise Stores, and Miscellaneous Store Retailers.

[13] The data was previously distributed under the name SafeGraph Patterns.

Specifically, we adopt the generalized synthetic control (GSC) methodology introduced by Xu (2017). The strategy relies on comparing changes in NYC before and after the policy to changes in outcomes observed in other cities, which are combined into a "synthetic control" for NYC using pre-policy data to estimate weights. The GSC method extends the classical synthetic control approach (Abadie, Diamond and Hainmueller, 2010) by incorporating interactive fixed effects, such that unobserved confounders may vary both across units (e.g., cities, Census tracts) and over time. This approach identifies the causal effect of the policy under a weaker assumption that, absent the policy, NYC's outcomes would have tracked the estimated 'counterfactual' outcomes, which are constructed from contemporaneous data in control cities weighted to match pre-policy patterns.

More formally, let $Y_{it}(0)$ denote the potential outcome for unit i at time t if untreated and $Y_{it}(1)$ denote the potential outcome if treated. We only observe

$$Y_{it} = D_{it}Y_{it}(1) + (1 - D_{it})Y_{it}(0)$$

where D_{it} is an indicator that takes the value 1 if unit i is treated at time t, and 0 otherwise. In our application, $D_{it} = 1$ for NYC after implementing the congestion pricing policy, and $D_{it} = 0$ otherwise. The set of control cities remains untreated throughout the sample period, allowing them to serve as comparisons.

Following the GSC framework, we assume that the untreated potential outcome $Y_{it}(0)$ can be decomposed into a low-rank factor structure plus some idiosyncratic noise:

$$Y_{it}(0) = \alpha_i + \gamma_t + \boldsymbol{\lambda}_i^\top \mathbf{f}_t + \varepsilon_{it} \tag{2}$$

where α_i is a unit fixed effect, γ_t is a time fixed effect, $\boldsymbol{\lambda}_i$ is a vector of factor loadings specific to unit i, $\mathbf{f}_t$ is a vector of common factors, and ε_{it} is an idiosyncratic error term. The functional form can be extended to include time-varying unit characteristics, although we do not use such characteristics in our setting.

Estimation. We first estimate these factors and loadings from the panel data of outcomes observed in the pre-treatment period. The factors and loadings are estimated separately for each outcome variable. Intuitively, this step aims to identify the factors and loadings that best predict changes in the outcome variable over time in the soon-to-be-treated units. Using the estimated parameters, we can then predict *counterfactual* outcomes $\hat{Y}_{it}(0) = \hat{\alpha}_i + \hat{\gamma}_t + \hat{\boldsymbol{\lambda}}_i^\top \hat{\mathbf{f}}_t$ for each post-treatment period $t > T_0$, where T_0 is the time period that congestion pricing began.

The Average Treatment on the Treated (ATT) effect in period $t > T_0$ is given by:

$$\widehat{\text{ATT}}_t = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} Y_{it} - \hat{Y}_{it}(0) \tag{3}$$

where $\mathcal{I}$ is the set of treated units. We will often report the aggregate ATT, too, which is estimated as $\widehat{ATT} = \frac{1}{T - T_0} \sum_{t > T_0} \widehat{\text{ATT}}_t$. We implement this procedure using the *gsynth* package for R provided by Xu (2017). We use cross-validation to select the dimensionality of the factors and, for inference,

we use parametric standard errors.

Identification. The key assumption of the GSC approach is that the error term ε_{it} is independent of treatment assignment, time and unit fixed effects, and the unobserved cross-unit ($\boldsymbol{\lambda}_i$) and cross-time ($\mathbf{f}_t$) sources of heterogeneity. This assumption has three important implications. First, the factor loadings $\boldsymbol{\lambda}_i$ must be stable over time. If a structural break changes the fundamental relationship between cities, this stability may be violated, leading to poor counterfactual predictions during the treated period. Second, much like traditional differences-in-differences (DiD), the GSC requires that the exact timing of treatment is exogenous after conditioning on the factor structure. Finally, this assumption is violated if there are spillover effects from NYC to control cities (e.g., large-scale traffic diversions or copy-cat policy changes).

This assumption is generally weaker than the standard DiD assumption that, in the absence of treatment, treated and control groups would follow parallel trends. Instead, GSC allows for multiple time-varying confounders that affect units differently. This is particularly valuable in contexts like urban traffic, where many latent factors (e.g., seasonal weather patterns) may affect different cities in different ways.

4 Direct Effects in the CBD

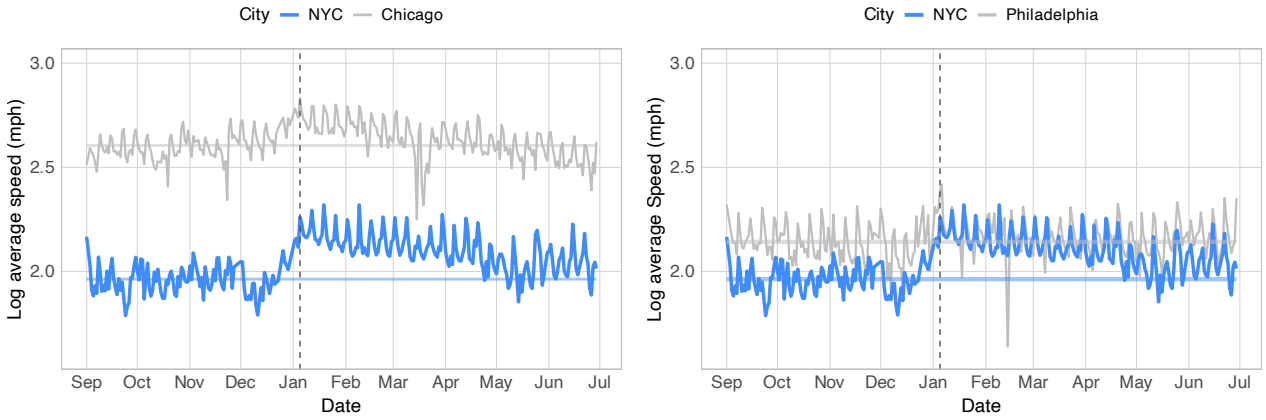
According to the MTA (MTA, 2025*b*), vehicle entries into the CBD decreased by approximately 10% following the introduction of the congestion pricing policy and have held steady (after accounting for seasonal fluctuations) through June, 2025. In this section, we examine the effects of congestion pricing on road conditions, air quality, consumer spending, and foot traffic in the NYC CBD. We find that the policy increased average speeds on CBD roads by 10% and improved estimated fuel efficiency by 3%. The biggest effects were on the slowest roads and at the most congested times of day. However, we find no evidence of impact on ambient air quality, visits to restaurants and retail establishments, and overall foot traffic.

4.1 Road Speeds within the CBD

The average speed on road segments in the NYC CBD increased from 7.1 mph in the four months preceding the policy's implementation to 8.2 mph in the six months after implementation (Figure 1). However, the speed increase before and after the policy's implementation is not unique to NYC. In fact, average speeds in the CBDs of Boston, Philadelphia, Atlanta, and Chicago were also higher after January 5th than in the preceding months (Table C.1), although the increase in NYC is the largest relative to pre-period speeds. Figure 1 shows head-to-head comparisons of CBD road speeds in NYC, Chicago, and Philadelphia.[14] In NYC, speeds increased sharply around the start of January 2025, and have remained above their pre-period average. Speeds in Chicago and Philadelphia also increased at the start of year, but have since fallen to their pre-period average in Philadelphia and below their pre-period average in Chicago.

[14] See Figure C.1 for average speeds over time in our other control cities.

Figure 1: Average speeds on CBD segments



Notes: This figure documents the volume-weighted average daily speed on road segments in the CBD of NYC, Philadelphia, and Chicago. The horizontal blue line represents the average speed in NYC between September 1 and December 15, 2024.

To what extent is the observed increase in speeds in NYC's CBD attributable to the effects of the congestion pricing policy? To evaluate this, we compare speeds in the NYC CBD to speeds in the CBDs of other cities by estimating Equation (3) using a panel of average speeds. Figure 2a plots the day-level ATT on log average speeds for CBD road segments during weekday priced hours.[15] Prior to the policy's launch, speeds in NYC were similar to the counterfactual speeds constructed from the set of comparison cities, suggesting that the GSC approach accurately captures trends across cities. After the onset of congestion pricing, average speeds in the CBD increased sharply by 10-15% relative to the synthetic control. The increase persisted through the first six months, with some reversion in May before rebounding in June. Averaging over this period, the implementation of congestion pricing increased speeds by 10%. This suggests that, under the assumptions of the GSC estimator, the policy's causal impact accounts for about three-fourths of the change in raw average speeds on segments within the CBD.

These average effects aggregate over substantial heterogeneity across times of day. Before the policy's implementation, average speeds in the CBD ranged from 14 mph in the early morning to half that in the late afternoon (Figure C.2). In Figure 2b we present ATTs on log CBD road speeds for each two-hour time window between 3am and 11pm. This breakdown of the policy's effects shows that the increases to average road speeds were modest in the early morning (3–5%), but rose to 15% in the afternoon peak and above 20% on weekend afternoons. That is, the relative impact of congestion pricing on speeds was largest during the times of day when congestion was most severe prior to the policy. In addition, we find that the increase in road speeds persisted after the end of peak hour pricing at 9pm, suggesting that the decreased entry volumes throughout peak hours had spillover effects to road conditions later in the evening.

These effects also aggregate over both extreme stop-and-go traffic and near-free-flow traffic. Most roads in the New York CBD are arterial roads (e.g., 5th Avenue) or substantially narrower local roads, both of which are generally slower (and with lower speed limits) than the handful of high-

[15] Figure C.3 includes similar plots for other measures of CBD traffic conditions, including median speeds.

volume highway segments included in the CBD area. The average speed on arterial roads prior to the policy was 7.1 mph, while on local roads it was only 5.5 mph (Table C.1). The 10th percentile of pre-policy speeds was under 3 mph for both local and arterial roads, while the 90th percentile was about 20 mph for local roads and 25 mph for arterials. The introduction of congestion pricing had a disproportionate impact on the slowest traffic conditions. The 10th percentile of road speeds experienced across the CBD increased by 13.8% on average, and the 90th percentile of road speeds increased by 2.8% on average. This suggests that one of the effects of the policy was to compress the tail of extreme delays—an outcome with potentially important welfare implications for reliability (Small, Winston and Yan, 2005). Measured directly, the dispersion of road speeds in the CBD as captured by the ratio of 90th to 10th percentile speeds decreased by 13.9%, and by the ratio of 80th to 20th percentile speeds, decreased by 9%

As a complementary exercise, we consider the impact of the policy at the level of trips that start and end inside the CBD (Figure 2d). This measure differs from the segment-level results above by capturing the end-to-end experience of travelers, aggregating over multiple links, intersections, and delays. However, it focuses only on trips that take place entirely within the CBD, while the segment-level results account for road conditions experienced by all drivers passing through the CBD regardless of origin or destination. In Section 5.3, we extend this exercise and discuss how different types of trips throughout the NYC metropolitan area were affected. For trips entirely within the CBD, we estimate that congestion pricing increased the average speed by about 7% since implementation—somewhat smaller than the segment-level estimates but persistent throughout the post-policy period.

Taken together, these results reinforce the canonical prediction that congestion pricing raises travel speeds on the priced roads (Vickrey, 1969), and provide new evidence on the magnitude, persistence, and distribution of effects across times and road types.

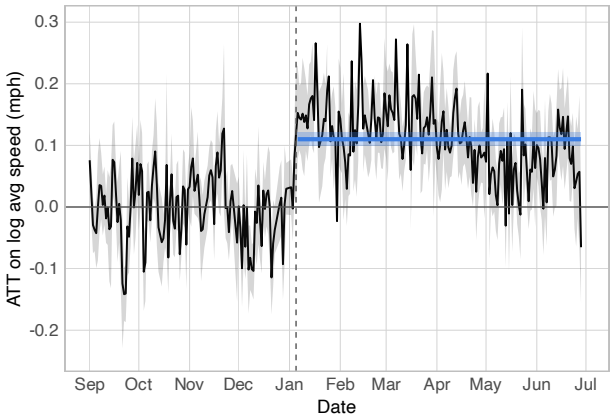
4.2 Fuel Consumption and Air Quality

Changes in traffic patterns also affect the quantity and location of vehicle emissions. CO_2 emissions impose externalities worldwide, while emissions of local air pollutants such as NOx, CO, and particulates impose an externality on nearby residents, which decays exponentially with distance from the road. Although it varies with factors such as terrain and weather, the decay rate for local air pollutants is steep, and air quality generally reaches 'background levels' beyond a few hundred meters (Brugge, Durant and Rioux, 2007; Liu, Chen and Xue, 2019). Exposure to vehicle emissions specifically has been found to have negative effects on infant (Currie and Walker, 2011; Knittel, Miller and Sanders, 2016) and adult (Buckeridge et al., 2002) health outcomes.

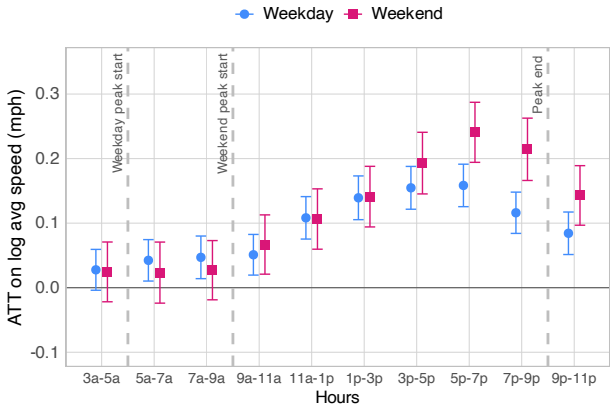
While we cannot observe vehicle emissions directly, we use predicted fuel consumption rates as a proxy for measuring how the change in traffic patterns affects potential emission rates. For a given vehicle type, the primary drivers of fuel consumption are vehicle speed profiles and the segment type. On a typical road segment, fuel consumption is most efficient at speeds between 60-80 mph, and is especially inefficient at speeds below 10 mph. Cars traveling at 5 mph—close to the average speed in the CBD during the late-afternoon—will consume about four times as much fuel per mile as those traveling at 60 mph (Figure A.1).

Figure 2: Effects on Speeds in the CBD

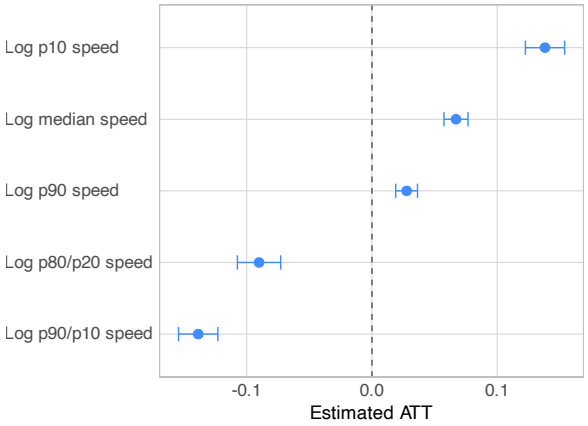
(a) ATT on log speeds: road segments



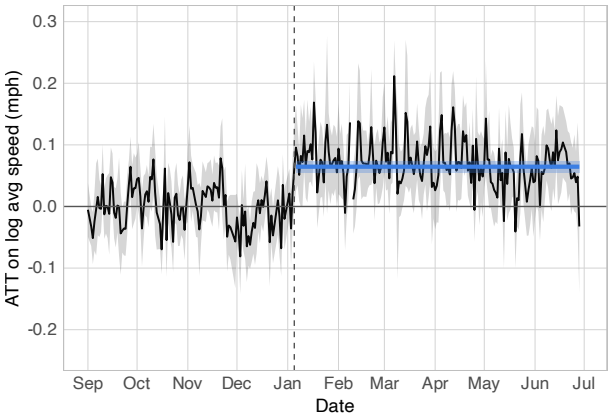
(b) ATT on log speeds: hour of day



(c) ATT on other traffic outcomes



(d) ATT on log speeds: all trips within CBD



Notes: Panel a) and d) documents day-level ATTs of congestion pricing over time on weekday peak hour CBD segment speeds and speeds on trips that start and end within the CBD. The data on average segment/trip speeds by two-hour bin in each CBD. The horizontal blue line is the aggregate ATT for all post-treatment periods. Panel b) estimates the treatment effect on log average speeds for each hour bin separately, using data from all other cities and hour bins as potential controls. Panel c) documents the ATTs on other traffic outcomes in the CBD, including median speeds and speeds on different types of roads. Shading and vertical bars denote 95% confidence intervals. Standard errors are clustered at the city-level.

Relative to the control cities, we find that the change in vehicle speeds and segments traveled translates into a 3.4% decrease in the estimated fuel consumption of vehicles traversing CBD segments (Figure 3). This estimate corresponds to the consumption *rate* for the average vehicle in the CBD, not the total fuel consumption across all vehicles. The effect on total fuel consumption—and the corresponding emissions—would be larger if accounting for any decreases in vehicle miles traveled in the CBD.

Despite the improved fuel efficiency, we find no evidence of an effect on ambient concentrations of fine particulates measured by PurpleAir sensors. Using sensor-day level data, we compare the $PM_{2.5}$ for PurpleAir sensors inside the NYC CBD to synthetic controls formed from sensors in other cities. The estimated treatment effect on particulate matter is statistically indistinguishable from zero (Figure 3). While the estimates are noisier than our fuel consumption estimates, we can reject changes of similar

magnitude to those found following the implementation in London (Green, Heywood and Paniagua, 2020) and Stockholm (Simeonova et al., 2021). The bottom end of the 95% confidence interval is -0.25μ/m^3, which would be a 3.3% decrease in the pre-period daily average $PM_{2.5}$ for sensors in NYC's CBD. One plausible explanation is that the London and Stockhom congestion charges—which began in 2003 and 2006, respectively—occurred when emissions standards were more relaxed and before the proliferation of electric vehicles, and so the marginal vehicle deterred by congestion pricing in NYC may be quite different than those deterred in London and Stockholm. For example, the 'Euro 3' emissions standards (2000-2005) allow for over seven times more NOx per kilometer than the 'Tier 3' US emissions standards (2017-) allow of NOx and NMOG combined (0.15 g/km vs. 0.018 g/km). Stockholm also exempted green vehicles from its congestion charge (Nilsson, Tarduno and Tebbe, 2024).

4.3 Commerce and Foot Traffic

The reported decrease in vehicle entries to the CBD raises the question of whether individual *visits* to the area decreased as well. Lower vehicle volumes may indicate fewer visits if drivers substituted toward driving to different locations or staying home, but not if drivers substituted toward carpooling or traveling by transit, walking, or other modes. To examine this, we apply our synthetic controls approach to estimate the treatment effects of the congestion pricing policy on visits to restaurants and retail establishments, as well as on overall foot traffic.

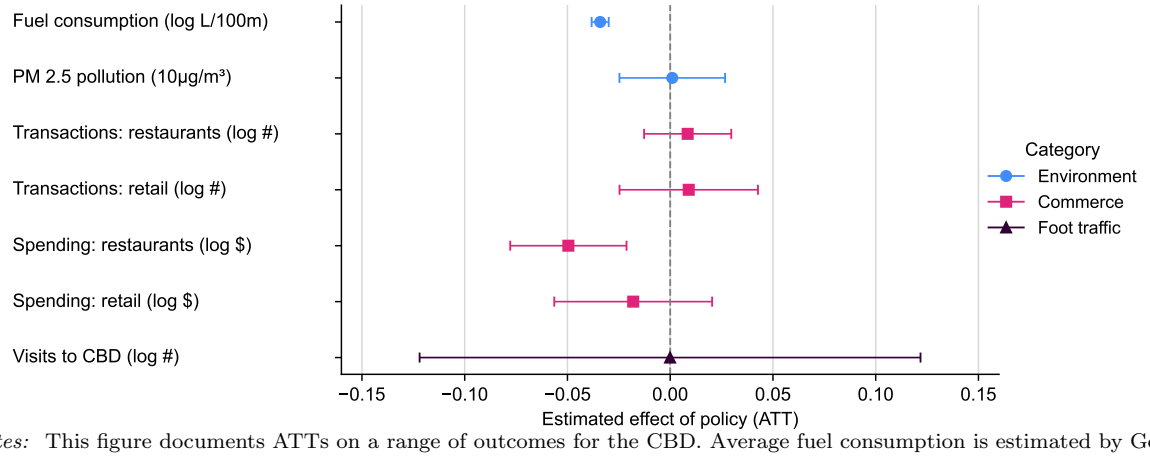
First, using the aggregated data from credit/debit card spending, we compare the number of transactions and the total amount spent at restaurant and retail establishments in the NYC CBD to those in our control cities. We estimate Equation (3) using zipcode-day level data and weight observations by the number of transactions in a given zipcode-category during 2024. The data capture all transactions within a day, without distinguishing peak from off-peak hours.

We find no evidence that congestion pricing reduced the number of transactions at either restaurants or retail establishments in the NYC CBD (Figure 3). Point estimates are close to zero, and the confidence intervals rule out decreases of over 1.2% at retail establishments and 2.5% at restaurants. However, we find that the total dollars spent decreased by 5% at restaurants and 2% at retail establishments, although the latter is not statistically different from zero. With transaction counts unchanged, this suggests that either customers are spending less per visit or establishments are lowering prices relative to control cities.

We also examine impacts on foot traffic in the CBD, again estimating Equation (3) with tract-day counts of CBD visitors during peak hours and weights corresponding to pre-policy total foot traffic. Here, too, we find no clear effect: the number of CBD visitors in the GPS data remains stable relative to our control cities, implying that any reduction in car trips has been offset by substitution toward other modes. This distinction is important. Speeds in the NYC CBD rise after the policy, consistent with fewer vehicles entering the cordon, yet restaurant/retail transactions and foot traffic remain steady. Together, the results suggest that travelers primarily adjust along the intensive margin of mode choice rather than the extensive margin of deciding whether to visit the CBD at all. However, while the point estimate for foot traffic is nearly zero, the confidence interval remains wide enough to

admit meaningful changes (Figure 3).

Figure 3: Effects on Other Outcomes in the CBD



Notes: This figure documents ATTs on a range of outcomes for the CBD. Average fuel consumption is estimated by Google's internal model, and we use the same specification as in Figure 2a. For pollution, we use day-sensor level measures of ambient air quality (PM2.5) from PurpleAir from sensors in the NYC CBD and in our control cities. For spending, we use day-merchant zipcode level data from MBHS3 on aggregate spending from a sample of credit and debit cards. We categorize retail and restaurant businesses based on their 3-digit NAICS code. For foot traffic, we use hour-tract level data from Advan on the number of visits by GPS-enabled devices, which we subset to peak hours and then aggregate to the day level. All data are subset to weekdays. Horizontal lines denote 95% confidence intervals.

5 Spillovers beyond the CBD

Cities are interconnected, and policies that affect one area invariably affect the entire city. A cordon toll directly prices only trips crossing into the cordon zone, but its impact on welfare depends as much on how traffic reallocates across the untolled parts of the network as on the priced zone itself (Verhoef, Nijkamp and Rietveld, 1996). If traffic diverts onto unpriced links, congestion may simply be displaced to other parts of the city. If, instead, overall volumes fall, traffic conditions may improve more broadly. Despite the importance of these spillovers, most prior evaluations of congestion pricing have primarily focused on traffic within the cordon zone only (e.g., Leape, 2006; Eliasson et al., 2009). In this section, we show that congestion pricing in NYC increased speeds on segments outside the CBD that had a high level of "co-occurrence" with trips to the CBD, resulting in lower observed travel times on trips throughout the NYC metropolitan area.

5.1 Road Speeds outside the CBD

To quantify the effects of spillovers to roads outside the CBD, we first develop a measure of the extent to which each road segment is exposed to the congestion pricing policy. The intuition behind our measure is simple: segments that are frequently traversed as part of trips to the CBD are more exposed to reductions in the number of CBD-bound trips than segments that are rarely part of such trips. We formalize this idea with the following measure of "co-occurrence" between each segment and the CBD of its metropolitan area:

Let S_{CBD} be the set of segments within a CBD, and let R be the set of observed trips in the

relevant metropolitan area within a given time span. Each trip $R_i = \{s_1, \ldots, s_N\}$ consists of a set of segments $s_j,\ j = 1, \ldots, N$ traversed between its origin and destination. For each segment s, we define its co-occurrence with the CBD C_s as:

$$C_s = \frac{|\{R_i \in R \mid s \in R_i \wedge S_{\text{CBD}} \cap R_i \neq \phi\}|}{|\{R_i \in R \mid s \in R_i\}|}. \tag{4}$$

In other words, C_s is the fraction of trips that passed through segment s, that also passed through at least one of the CBD segments during the period of observation. We compute the co-occurrence of each segment in NYC and each of our control cities with respect to its relevant CBD using data from 2024, and hold these values constant throughout our analysis. Figure C.4 plots the spatial distribution of road segments in NYC by their level of co-occurrence.

Our definition of co-occurrence provides a natural benchmark for evaluating spillover effects. If there were no diversion of driving to nearby areas such that all decreases in CBD entries corresponded to reductions in vehicles driving at all, and decreases in entries were proportional to pre-policy volumes, then co-occurrence would predict the policy's impact on volumes for roads outside the CBD. For instance, if a road segment had 50% co-occurrence with the CBD and the congestion pricing policy reduced CBD entries by 10%, then we would expect that road segment to have 5% fewer vehicle traversals after congestion pricing was implemented. If the relationship between road speeds and volumes were linear, then we would consequently expect to see a 5% increase in speeds. We do not directly observe volumes in our data, and, as we discuss in Section 7, the relationship between speeds and volumes is generally nonlinear. However, a comparison of the treatment effects on segments with different co-occurrence levels provides several insights about the network-wide effects of congestion pricing in NYC.

Figure 4 plots average ATTs over the first six months of congestion pricing, split by co-occurrence quintile and road type (i.e. arterial, highway, and local roads, or a volume-weighted average of all three). In each case, we estimate the treatment effects according to the approach outlined in Section 3.3 using segments of the same co-occurrence quintile and road type in other cities to form a synthetic control. Across all road types, the average non-CBD road segment with an 80-100% co-occurrence—that is, a segment for which 80-100% of traversals pertained to trips entering the CBD prior to the policy—experienced a 14% increase in speeds. The effects then decrease monotonically with co-occurrence, but even road segments with 0-20% co-occurrence experienced a 1.8% increase in speeds.

Notably, the estimated treatment effect on roads with 80-100% co-occurrence was larger than the estimated effect on speeds within the CBD in both relative and absolute terms. A breakdown of treatment effects by road type illustrates that this is primarily driven by the larger share of highway segments in the 80-100% co-occurrence bin than in the CBD. The average treatment effect on the highest co-occurrence highways, which include many of the major entrances to the CBD (e.g., the Lincoln and Holland Tunnels), was over 15%.[16] As the average pre-policy speed on highways with

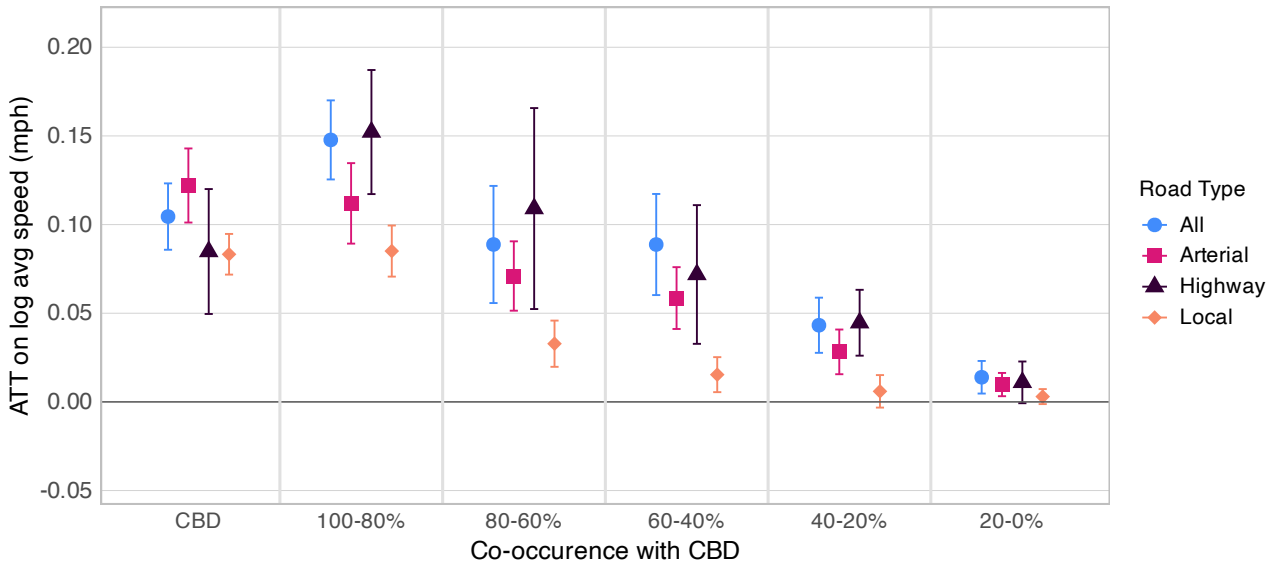
[16] This is consistent with early evidence from the Congestion Pricing Tracker, which showed larger changes in speeds on bridges leading into the CBD than within the CBD itself (Moshes and Moshes, 2025). There are only a few highway segments within the CBD itself, as the FDR Drive and West Side Highway connections to West Street are exempted from congestion pricing unless drivers exit into the CBD. The CBD highway segments include just the exit/entry segments

80-100% co-occurrence was 21 mph, this effect translated to a 3 mph increase in average speeds on the highway segments themselves and a 2 mph increase in average speeds on roads with 80-100% co-occurrence overall. By contrast, the average increase in road speeds within the CBD was about 0.7 mph.

More generally, congestion pricing substantially increased speeds on highways and arterial roads throughout the metro area. Even highways and arterial roads that had only 0-20% co-occurrence with the CBD experienced a 1% average increase in speeds. Highly exposed local roads, such as the roads with 80-100% co-occurrence just above the northern border of the CBD at 60th St., experienced a similar increase in speeds to the local roads inside the CBD. Farther away local roads with 60-80% and 40-60% co-occurrence experienced smaller increases in speeds, and local roads with less than 40% co-occurrence were largely unaffected.

Overall, the evidence suggests broad improvements in traffic conditions outside the CBD, with the largest gains on the main approaches into the CBD, but positive effects extending across much of the network. Importantly, we find no evidence of offsetting slowdowns on different road types or co-occurrence levels, suggesting that the policy reduced overall traffic volumes rather than simply displacing congestion to untolled roads. However, aggregating by co-occurrence may mask variation across neighborhoods, and concerns remain that some areas could face increased traffic. As we discuss in Section 5.3, we find little evidence of offsetting effects across neighborhoods of different income levels or in specific areas such as the Bronx and New Jersey.

Figure 4: Effect on Speeds by Co-Occurrence and Road Type



Notes: This figure documents treatment effects split by levels of co-occurrence and type of road segment. Each point is separately estimated using the average speeds in two-hour bins for segments with the corresponding level of co-occurrence and road segment type for both NYC and the comparison cities. Vertical bars represent 95% confidence intervals. Standard errors are clustered at the city-level.

connecting highways to the CBD.

5.2 Effects on Trips throughout the Metro Area

The spillover effects documented above at the segment level raise the question of how congestion pricing affected conditions for entire trips across the region. Travelers experience congestion not segment by segment but as the cumulative outcome of their journeys, and welfare ultimately depends on trip-level costs (Arnott, de Palma and Lindsey, 1993; Small, Winston and Yan, 2005). We classify trips into one of four categories depending on whether the origin and destination fall inside or outside the CBD: to CBD, from CBD, within CBD, and outside CBD. For passenger cars, only trips that start outside the CBD and enter the CBD ("to CBD") are subject to pricing, but all trips may be indirectly affected through the network spillovers documented above. For each trip type, we again use the GSC approach to estimate the treatment effect of congestion pricing on log average speeds over the entire trip.

We find that congestion pricing increased speeds not only on the priced trips to the CBD, but also on unpriced trips traveling within, from, and outside the CBD. Figure 5a plots the ATT on log average trip speeds within each origin-destination category. The policy increased speeds on trips to the CBD by about 9% and speeds on trips within and from the CBD by 6-7%. Even trips that took place entirely outside the CBD experienced modest gains of roughly 0.5%, although, at this level of aggregation, the effect is not statistically significant. While small in magnitude, these improvements apply to the majority of trips in the metro area, adding up to a substantial aggregate impact.

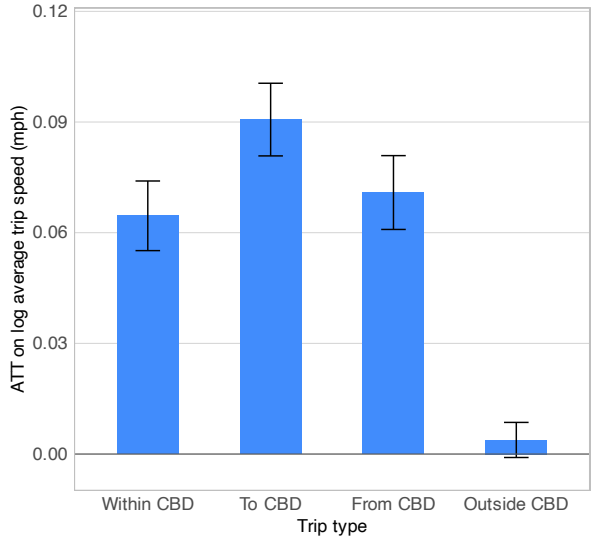
To examine the spatial pattern in more detail, we estimate effects across a 5-by-5 origin–destination matrix based on distance from the CBD (Figure 5b). We classify all origins and destinations based on the within-city percentile of their distance to the CBD (see Figure C.6 for a map). This disaggregation shows the largest gains for trips directly involving the CBD, but also nontrivial improvements for trips connecting outer areas. Trips traveling between the 0-20% distance bin and 60-100% bin experienced speed gains of about 2% in each direction. Disaggregating ODs to this level also lets us detect small changes even in trips on the outskirts of the metro area. For trips within the 60-100% distance bin, for example, we estimate a treatment effect of 0.7%. As we discuss in Section 6, these changes on unpriced trips—including the small gains on the many trips entirely outside of the CBD—have first-order implications for the policy's net welfare impact.

5.3 Distributional Effects

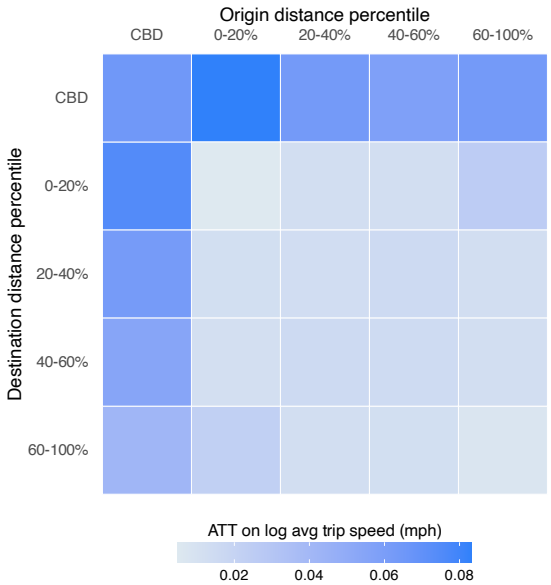
Proposals to implement congestion pricing often face concerns over potential adverse impacts on specific groups, such as residents of specific areas or lower-income drivers (Ecola and Light, 2009; Taylor, 2010). We estimate heterogeneous effects across income groups and geographic areas to assess whether the effects of congestion pricing in NYC were unevenly distributed. For each region, Figure 6 documents the treatment effects on three outcomes: average speeds for road segments within the region, average speeds on trips *to* the CBD that originate in the region, and average speed on trips *from* the CBD that end within the specified region. The estimates follow the same generalized synthetic control methodology as in previous sections. Because the spatial distribution of these regions may be related to CBD travel in different ways in different cities, we use the set of all corresponding regions in each of our control cities as potential controls (e.g., the controls for the ATT on the bottom income

Figure 5: Effect on Trip Speeds throughout NYC

(a) By trip type



(b) By OD distance to CBD



Notes: The first panel documents the ATTs for each of our four trip types. Vertical bars represent 95% confidence intervals. The second panel plots ATTs at a more spatially disaggregated level by binning trips based on their origin and destination's distance to the CBD. Each cell documents the ATT on log trip speeds for trips that originate in the bin indicated in the column labels and end in the bin indicated in the row labels. Grey hashes indicate ODs where the ATT is not different from zero at 5% level.

quintile are all income quintiles in other cities).

Household income. We begin by comparing effects across the neighborhood income distribution. We group Census tracts into within-city quintiles based on their median household income in the 2018–2023 American Community Survey (ACS), which range from roughly $44,500 in the bottom quintile to $182,000 in the top.

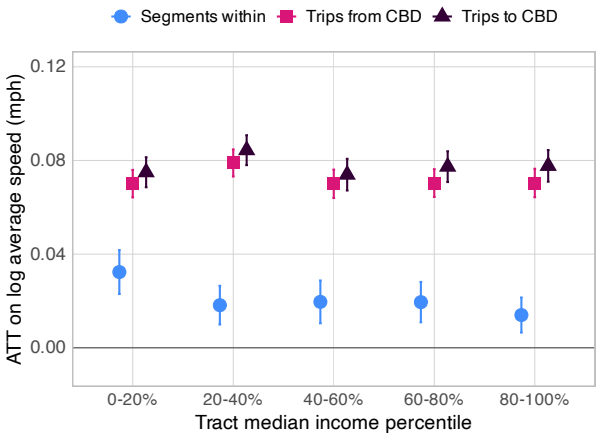
We find broadly similar treatment effects on tracts across the income distribution. If anything, the largest gains occurred in lower-income areas – road segments in the bottom quintile experienced speed increases of 3.2%, compared with slightly smaller gains for segments in higher-income tracts. Trips traveling to or from the CBD from any income quintile experienced a 7-9% increase in speeds, with no significant differences in effects across quintiles.

Specific regions of interest. We next examine effects across specific geographic areas. In New Jersey, average speeds increased by about 4.7% on road segments in Hudson and Bergen Counties, the two counties most directly connected to Manhattan. Trips from these counties to the CBD became 9.8% faster, while trips from the CBD to these counties became 9.6% faster. Within New York City, roads inside the Bronx became 2.8% faster, but trips to and from the CBD became 5.9% and 5.1% faster respectively. Finally, Long Island roads also became slightly faster on average, increasing in speeds by 2.3%, and travel between Long Island and the CBD became significantly faster; trip speeds to and from the CBD increased by 8% and 7.9%, respectively. Figure C.7 further breaks out the effects on speeds for trips to the CBD based on the entrance they cross to enter the CBD. The largest effects

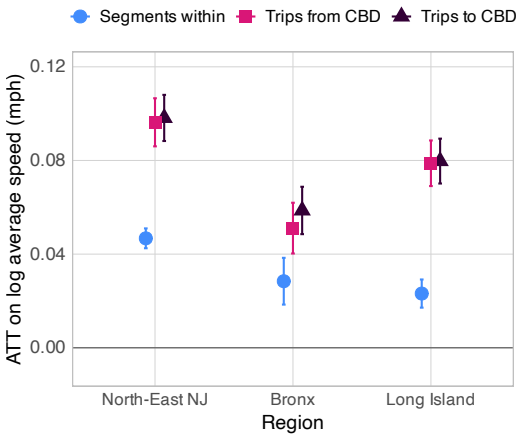
are for trips crossing the Holland Tunnel, Lincoln Tunnel, and Queensboro Bridge.

Figure 6: Effects on Speed by Geographies' Characteristics

(a) Tract median income



(b) Specific regions



Notes: These figures plot the ATTs on three measures of speeds: speeds on segments within a geography and speeds on trips to (from) the CBD that start (end) in the specified geography. Vertical lines represent 95% confidence intervals.

6 Driver Welfare

The implementation of congestion pricing in NYC increased speeds not only in the CBD, but across the New York metropolitan area. As a consequence, many trips benefited from speed increases without paying any congestion fee. In this section, we conduct a simple stylized exercise to compute a lower bound on the welfare effects generated by speed increases for trips between different origins and destinations. We combine these estimates to bound the net welfare impact of the congestion pricing policy on trips by private passenger vehicles and taxi/FHVs.

Pre-Policy Utility Consider an individual who would typically drive between an origin o and destination d during weekday peak hours prior to the implementation of congestion pricing, which we will denote as period 0. On an average day, the individual would face a travel time t^0_{od}, a cost of driving p^0_{od} that includes tolls, gas, depreciation, etc., and some baseline value for taking the trip ξ_{od}. We assume that these elements enter the driver's utility linearly:

$$u^0_{od} = \xi_{od} - p^0_{od} - \omega \times t^0_{od}. \tag{5}$$

Here, we normalize the marginal value of the cost of driving to 1 such that the coefficient on travel time ω can be interpreted as the value of travel time (VOTT) in dollars per unit of time.

There may be many individuals who would take trip (o, d) if it were sufficiently attractive relative to their next best option. We assume that each individual i considering a trip had an idiosyncratic outside option utility $\varepsilon^0_{i,od}$. By revealed preference, individuals who took trip (o, d) prior to the policy

preferred it to their outside option, $u^0_{od} \geq \varepsilon^0_{i,od}$, while individuals who did not take the trip preferred their outside option instead.

Note that for simplicity, we define a trip by its origin and destination such that the choice of whether to take the trip or not is binary and switching modes is not explicitly considered, and assume that the baseline value of taking the trip ξ_{od} is fixed over time. Because we do not place any restrictions on the outside option utilities, the outside options could include other modes of travel that may become more or less attractive relative to the baseline over time. As such, this assumption does not play an important role in our analysis.

Post-Policy Utility The congestion pricing policy had two major effects on the travel conditions faced by individuals: (1) if the trip required entering the CBD by vehicle, then the cost of driving increased by the congestion toll amount Δp_{od}; (2) the average travel time t decreased by an origin-destination-specific amount Δt_{od} (and auxiliary costs such as gas may have decreased as well as a result). Because all trips—defined as the OD pairs in Figure 5b—became weakly faster on average (i.e., $\Delta t_{od} \geq 0$), the utility that an individual would obtain for taking the same trip after the policy was implemented, which we will denote as period 1, is bounded from below as follows:

$$u^1_{od} \geq \xi_{od} - (p^0_{od} + \Delta p_{od}) - \omega \times (t^0_{od} - \Delta t_{od}). \tag{6}$$

For some individuals, these changes pushed them over the margin of deciding whether to take a trip or not. Individuals for whom the post-policy utility of taking their trip was higher than their next best option, $u^1_{od} \geq \varepsilon^1_{i,od}$, may have switched to taking the trip after the policy even if they hadn't taken the trip before. Conversely, individuals for whom $u^1_{od} < \varepsilon^1_{i,od}$ may have switched to taking their outside option even if they had been taking the trip before.

Bounding Welfare Contributions The net change in welfare from the policy is given by the difference in the sum of utilities for all affected individuals. For each trip (o, d), these individuals comprise three groups: those who took the trip before the policy, and continued to do so afterward; those who did not take the trip before the policy but started to afterward; and those who took the trip before the policy but stopped doing so afterward.[17] For each of the individuals who continued to make the same driving choice, the difference in utilities is just $u^1_{od} - u^0_{od}$. For the remaining individuals who switched either towards or away from their trip, the difference in their utilities relies on a comparison with their outside option.

Denote $\bar{\varepsilon}^0_{od,+}$ as the average outside option utility for individuals who only started taking the trip after the policy and $\bar{\varepsilon}^1_{od,-}$ as the average outside option utility for individuals who only took the trip prior to the policy but stopped afterwards. The average difference in utility obtained by the new trip takers is given by $u^1_{od} - \bar{\varepsilon}^0_{od,+}$. Similarly, the average difference in utility obtained by the former trip takers is given by $\bar{\varepsilon}^1_{od,-} - u^0_{od}$. We cannot credibly estimate the change in trip volumes due to the policy using our data, and so we cannot infer the values of $\bar{\varepsilon}^0_{od,+}$ and $\bar{\varepsilon}^1_{od,-}$ in order to compute an

[17] We assume that any trips that are unaffected by either the congestion toll price or the change in road speeds experienced no change in utility as a consequence of the policy.

exact welfare impact. However, we can construct a simple bound on the welfare contribution of each trip based on *pre-policy* volumes, using the assumption that an individual would only start (stop) to take a given trip if that trip provided strictly higher (lower) utility after the policy than it did before. We describe the bounds formally here; Appendix B.2 offers a simple graphical illustration of the argument.

Formal Bound Derivation Let N^0_{od} be the number of individuals who took trip (o, d) prior to the policy, and let $N^{1,s}_{od}$, $N^{1,+}_{od}$, and $N^{1,-}_{od}$ be the numbers of individuals who stuck with their driving choice, newly started taking the trip or stopped taking the trip after the policy, respectively. The net contribution of trip (o, d) to the welfare impact of the policy is then bounded from below as follows:

$$\begin{aligned}
\Delta W_{od} &= N^{1,s}_{od} \times (u^1_{od} - u^0_{od}) + (N^{1,+}_{od}) \times (u^1_{od} - \bar{\varepsilon}^0_{od,+}) + (N^{1,-}_{od}) \times (\bar{\varepsilon}^1_{od,-} - u^0_{od}) && (7)\\
&\geq N^{1,s}_{od} \times (u^1_{od} - u^0_{od}) + (N^{1,-}_{od}) \times (\bar{\varepsilon}^1_{od,-} - u^0_{od}) && (8)\\
&\geq N^{1,s}_{od} \times (u^1_{od} - u^0_{od}) + (N^{1,-}_{od}) \times (u^1_{od} - u^0_{od}) && (9)\\
&\geq N^0_{od} \times (u^1_{od} - u^0_{od}) && (10)\\
&\geq N^0_{od} \times (\omega \times \Delta t_{od} - \Delta p_{od}). && (11)
\end{aligned}$$

Here, line (7) breaks up the net welfare contribution between individuals who took trip (o, d) prior to the policy and continued to do so afterward, individuals who only started taking the trip after the policy, and individuals who only stopped taking the trip after the policy. For individuals who newly started taking the trip, $u^1_{od} - \bar{\varepsilon}^0_{od,+}$ must be greater than zero by revealed preference, yielding the inequality in line 8. For individuals who newly stopped taking the trip, $\bar{\varepsilon}^1_{od,-}$ must be greater than u^1_{od}, yielding the inequality in line (9). Line (10) is given by the accounting identity $N^0_{od} = N^{1,s}_{od} + N^{1,-}_{od} - N^{1,+}_{od}$. Line (11) results from subtracting equation (5) from inequality (6).

Bounding Net Welfare The net welfare impact of the policy across the metropolitan area is given by the sum of welfare contributions across all OD pairs:

$$\Delta W = \sum_o \sum_d \Delta W_{od}. \tag{12}$$

Applying the inequality from line (11) we obtain the following lower bound on the net welfare impact:

$$\Delta W \geq \sum_o \sum_d \left[N^0_{od} \times (\omega \times \Delta t_{od} - \Delta p_{od}) \right]. \tag{13}$$

Connecting to Data. To estimate the bounds empirically, we limit our attention to welfare on trips by passenger vehicles and taxi/FHVs, setting aside effects on other road users (e.g., buses and commercial trucks) and how the city uses the toll revenue. In practice, the toll revenue has been committed to improvements to public transit in NYC, and speeds for other road users have also increased. Both of these suggest that our lower bound on welfare would likely increase by accounting

for revenue recycling and other road users, although for the latter it depends on other road users' VOTT and the share traveling on priced trips.

As in Figure 5b, we group origins and destinations by five bins based on their distance to the CBD. For each OD pair, the effect of congestion pricing on travel time is given by

$$\widehat{\Delta t_{od}} = \bar{t}_{od} \times \left(\frac{-\widehat{\text{ATT}}_{od}}{1 - \widehat{\text{ATT}}_{od}} \right), \tag{14}$$

where $\widehat{\text{ATT}}_{od}$ are the estimated ATTs on trip speed from Figure 5b and $\bar{t}_{od}$ is the pre-period average travel time. We assume that passenger vehicles and taxis are affected by road speed improvements in the same way, so that $\widehat{\Delta t_{od}}$ does not depend on the vehicle type. To estimate the change in tolls ($\widehat{\Delta p_{od}}$), we use data from the MTA on CBD entries by vehicle class and entry point to compute the average congestion fee paid by passenger vehicles and taxi/FHVs. For passenger vehicles, we assume they enter a single time per day (i.e., we do not amortize their per-day cost across entries) and compute the average price net of any rebates for entering certain already-tolled entries.[18]

To estimate the pre-period trip volumes N^0_{od} for each vehicle type, we use volumes reported by Replica, a company that combines a variety of data sources including GPS devices, sensors embedded in roads, taxi/FHV companies, and administrative sources to simulate mobility patterns throughout entire metro areas. Replica's primary customers are in the public sector, and include the MTA. We use their estimated tract-to-tract travel flows by passenger vehicles and taxi/FHV for the NYC metro area for a 'typical weekday' in 2024 Q4, and sum across tracts in each of our origin-destination groups. For trips that enter the CBD, we supplement the data from Replica with data on CBD entries from the MTA (MTA, 2025*b*). The estimates by Replica are below the reported daily CBD entries by the MTA for 2024, so we scale the counts of trips to the CBD from Replica up to exactly match the aggregate entries reported by the MTA. We describe this process further in Appendix B.1.

The final component needed to estimate the welfare bounds is a value of travel time (VOTT) coefficient ω. We use \$40/hour as a benchmark—which is approximately the average hourly wage in the NYC metropolitan area (BLS, 2025)—and consider several threshold values that might change the qualitative interpretation of our analysis.

For passenger vehicles, we find that aggregate driver welfare increased by at least \$0.8 million, with the positive welfare gains coming entirely from *unpriced* trips. Figure 7a plots the per-trip welfare gain for each of the origin-destination pairs. The average passenger vehicle trip to the CBD paid \$7.9 to save just 2.7 minutes, which, with a VOTT of \$40/hour, adds up to a minimum total welfare effect of -\$1.7 million across the 274,000 trips to the CBD each day. For these trips, the average driver would only be guaranteed to be better off following the implementation of congestion pricing if her VOTT were above \$176/hour. Meanwhile, while unpriced trips save just 8 seconds per trip, there are far more unpriced trips than priced trips. Across the 27.3 million daily unpriced passenger vehicle trips,

[18] Drivers entering via the Lincoln Tunnel, Holland Tunnel, Queens Midtown Tunnel, Hugh L. Carey Tunnel receive a \$3 'cross credit' so pay an effective price of \$6 to enter the CBD, while drivers entering via the other entrances pay the full amount. We compute the weighted average price paid using the total passenger vehicle crossings for each entry (MTA, 2025*b*). We do not account for the additional 50% discount that low-income drivers can receive for trips after their first 10 trips each month.

these time savings add up to a welfare gain of at least $2.5 million. Taking into account both priced and unpriced trips, the 'break-even' VOTT of drivers is $29/hour.

Passengers of Taxi/FHVs enjoyed similar time savings but at a lower price than passenger vehicles, assuming that the only effect of the policy on taxi/FHV pricing is the congestion fee itself. For these trips, the toll on any travel to, from, or within the CBD was set at $0.75 for taxi and $1.50 FHV after the policy was implemented. We assume an average toll price of $1, because about two-thirds of such trips are by taxis in the entry counts published by the MTA (MTA, 2025*b*). Thanks to the cheaper congestion fees, both priced and unpriced taxi/FHV trips experienced an aggregate net welfare gain. The minimum per-trip welfare gains are largest for trips going to the CBD, which offered the most time savings. Taxi/FHV trips within the CBD were worse-off, because their ex-ante shorter trip durations translated to smaller time savings. The break-even VOTTs are $27/hour for the average priced trip and $25/hour for unpriced trips. There are fewer taxi/FHVs trips overall, so these gains add up to only $191,600.

In total, our estimates suggest that the congestion pricing policy had a positive net welfare impact of at least $1.0 million per day for passenger vehicle drivers and taxi/FHV passengers, driven almost entirely by the large welfare gains on *unpriced* passenger vehicle trips. These welfare gains are before even accounting for the revenue raised. The implied peak hour revenue from the CBD entry counts and average prices is $/2.5/ million per weekday, although this is only an approximation as it does not account for other vehicle classes, the low-income discount plan, payment methods (it costs more to pay by mail), or repeated entries in a single day.[19] In Appendix B.3, we document all of the input values for our welfare calculations so that skeptical readers can explore the results under alternative inputs.

7 Mechanisms

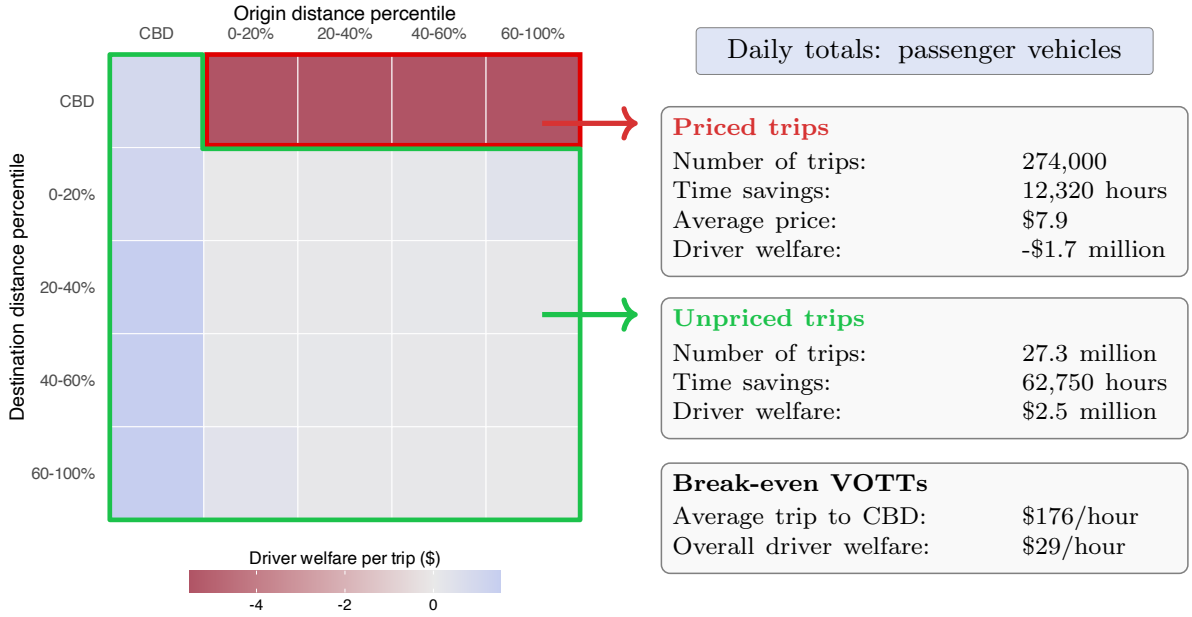
The net effects of introducing congestion pricing depend on three underlying ingredients: 1) how pricing affects volumes; 2) the exposure of other drivers to changes in volumes (e.g., via the co-occurrence of segments they traverse); and 3) the relationship between volumes and speeds on different road segments. In this section, we evaluate how these ingredients contributed to the effectiveness of the policy in NYC, and whether similar effects may be obtained in other cities. To simplify the problem, we will consider a scenario in which all cities set prices such that there are 10% fewer trips to the CBD from all origins, similar to the 11% decrease in CBD entries estimated by MTA (2025*b*) for NYC. We then quantify how this change in volumes propagates throughout the road network in each city, and whether the magnitude of the time savings on tolled and untolled trips is large enough to outweigh the costs imposed on now-tolled trips to the CBD.

Exposure to changes in volumes. Travel patterns vary across cities, such that the average trip in different cities may be more or less exposed to any changes in volumes stemming from the introduction

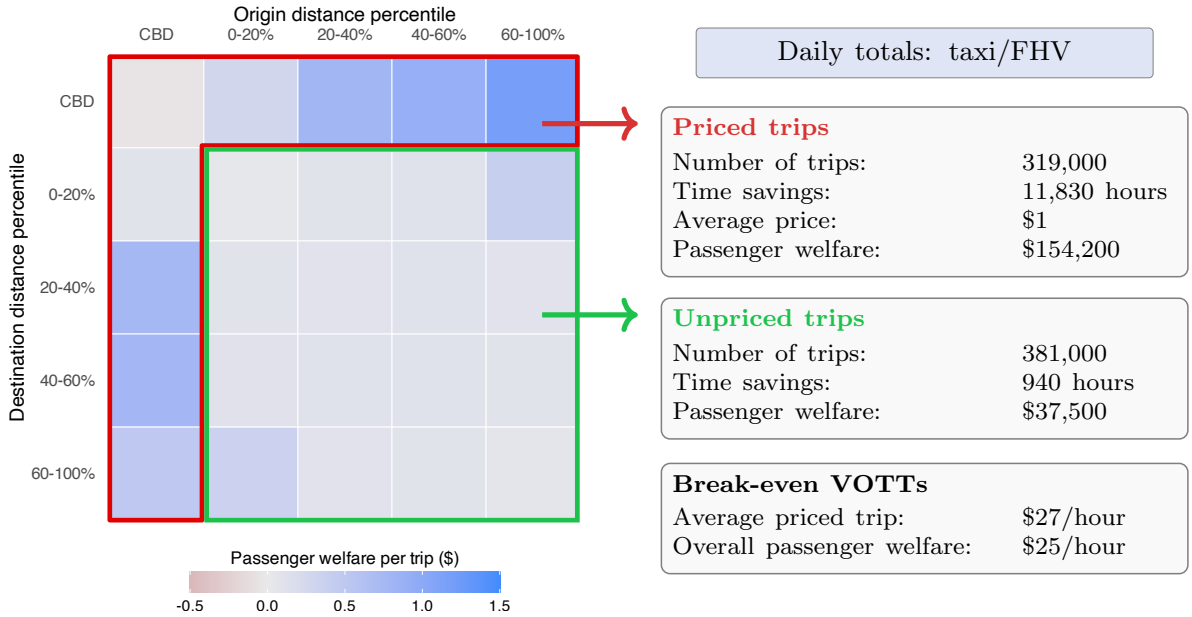
[19] The MTA reported CBD toll revenues of $159 million for January to March 2025 and operating expenses of $35.9 million (MTA, 2025*c*).

Figure 7: Driver welfare estimates

(a) Passenger vehicles



(b) Taxi/FHV trips



Notes: Grey hashes indicate ODs where the ATT is not different from zero at 5% level; these are excluded for aggregate time savings and welfare. The number of trips is the average tract-to-tract passenger car trips or taxi/FHV trips estimated by Replica for a typical weekday in 2024Q4. For trips to the CBD, we further scale the estimates from Replica to match aggregate counts of pre-period entries from the MTA (MTA, 2025*b*). Average prices are computed based on entries data. The passenger vehicle price is below $9 because of crossing credits for some entries, and the taxi/FHV price is set between the per-trip prices for taxis ($0.75) and FHV ($1.50). The welfare estimates are computed for each cell separately and assume a Value of Travel Time (VOTT) of $40/hour and no revenue recycling.

of cordon-based congestion pricing. Consider two extreme travel patterns: one where all trips to the CBD come from the farthest reaches of the metro area, and a second where all trips to the CBD begin just outside of its boundaries. In the former case, reducing the number of trips to the CBD will have spillovers throughout the metro area, as many non-CBD trips will travel along segments with high co-occurrence. In the latter case, few segments will have high co-occurrence, and, while congestion in the CBD itself may improve, we would expect minimal spillovers beyond the CBD.

We combine our segment-level measures of co-occurrence with data from trips between origins and destination pairs to quantify how exposed different trips are to changes in CBD volumes. For each trip type t, we measure the average time spent on each type of road segment, where we divide road segments based on their co-occurrence bin $c \in \{\text{CBD}, 100\text{-}80\%, 80\text{-}60\%, 60\text{-}40\%, 40\text{-}20\%, 20\text{-}0\%\}$ and road type $k \in \{\text{local}, \text{arterial}, \text{highway}\}$. Using the same matrix of ODs as in Figure 5b to define trip types, we compute the average time on each segment type as

$$\bar{t}_{od}^{cr} = \frac{1}{|R_{od}|} \sum_{i \in R_{od}} \sum_{s \in R_i} t_{is} \times \mathbb{1}_{c_s = c} \times \mathbb{1}_{r_s = s} \tag{15}$$

where R_{od} is the set of observed trips for a given OD pair, each trip $R_i \in R_{od}$ is composed of segments $R_i = \{s_1, s_2, \ldots, s_J\}$, $\mathbb{1}_{c_s=c}$ is an indicator for whether segment s is in co-occurrence bin c, and $\mathbb{1}_{r_s=r}$ indicates whether the segment is of road type r.

We also measure the overall exposure for the average trip to the CBD, from the CBD, and outside of the CBD based on the co-occurrence of segments traversed, weighted by the time on each segment. For each trip type $l \in \{\text{To CBD}, \text{From CBD}, \text{Outside CBD}\}$, we aggregate across different ODs to compute its average exposure as:

$$\bar{e}_l = \frac{1}{|R_l|} \sum_{i \in R_l} \sum_{s \in R_i} \frac{t_{is} \times c_s}{t_i} \tag{16}$$

where R_l is the set of trips of type l, t_i is the total travel time for trip i, and we overload notation slightly to let c_s be the actual co-occurrence of segment s (not the co-occurrence bin).

The first three columns of Table 1 document this measure of exposure for NYC and each of our control cities. Naturally, trips to and from the CBD are more exposed to the policy than trips entirely outside of the city. Including the part of the trip that is within the CBD, the average segment-level co-occurrence on trips to/from the CBD ranges from about 55-75% (weighted by the duration on each segment). Across cities, NYC drivers to/from the CBD have some of the largest exposure to other CBD drivers, second only to Chicago. Cities with high average exposure, such as Chicago, have travel patterns that suggest more potential for positive spillovers onto the speeds of unpriced trips.

Congestion functions. The second key ingredient is the local relationship between volumes and road speeds. If all roads were already operating under free-flow conditions, then further decreases in CBD volumes would have no effect on travel times, regardless of a given trip's exposure to the policy. The empirical relationship between speeds and the volume (or density) of cars is often used by traffic engineers to evaluate the effects of past and potential changes to traffic flows (Greenshields et al., 1935; Greenberg, 1959; Seo et al., 2017), and is called the 'congestion function' of a road.

Table 1: Potential Effects in Other Cities

	Average co-occurrence ($\bar{e}_l$)			Effects of ↓ 10% CBD trips	
City	To CBD	From CBD	Outside CBD	Time savings (hr/day)	Break-even price ($)
NYC	71.3%	71.4%	4.8%		
Philadelphia	63.0%	63.7%	3.0%		
Chicago	73.9%	74.2%	9.7%		
Boston	59.1%	59.1%	9.3%		
Atlanta	54.8%	55.0%	8.1%		
Baltimore	55.8%	55.5%	5.2%		

Notes: TBD

These congestion functions may vary across roads based on their physical characteristics (e.g., width, curvature, lane structure, etc.) and typical driving behavior (e.g., spacing between cars). A curvy road full of potholes will operate at slower speeds for a given density of cars than a well-maintained highway (Currier, Glaeser and Kreindler, 2023).

We estimate congestion functions for groups of road segments g using samples of simultaneous speeds v_g (in miles per hour) and densities ρ_g (in vehicles per mile). We cannot directly observe total volumes or road density, as our data cover only a sample of cars on the road, so we infer density by assuming a constant scaling factor on observed traversals and normalizing relative to the maximum density observed. We then fit the following functional form, taken from the Bureau of Public Roads (BPR):

$$\frac{1}{v_g} = \frac{1}{v_{\text{FF}}}\left[1 + \alpha\left(\frac{\rho_g}{\kappa}\right)^{\beta}\right] \tag{17}$$

In this specification, v_{FF} denotes the free-flow speed, ρ_s is the density of vehicles on the segment, κ is the capacity (the density at which congestion effects become significant), α is a scaling parameter, and β governs the curvature of the congestion response.

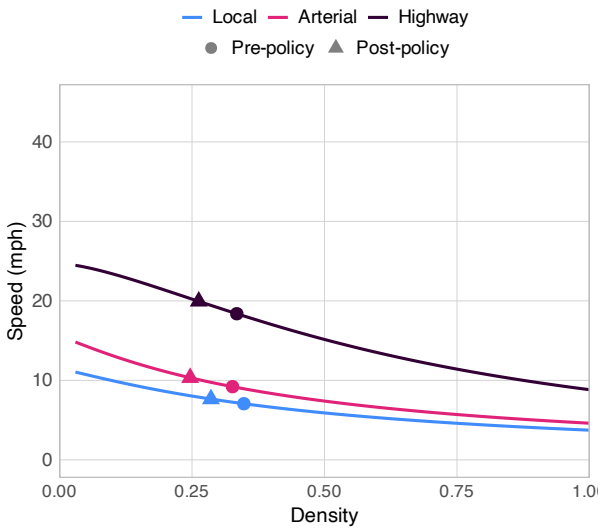
We estimate separate congestion functions for each combination of road type and co-occurrence bin in each of our cities.[20] Figures 8a and 8b plot the congestion functions for the CBD and 80-60% co-occurrence bins by road type in NYC; the other NYC congestion functions are shown in Figure C.8. The circles on each represent average speeds on the average day prior to the policy's launch, and the triangles represent the post-policy speeds computed as the pre-period speeds times one plus the estimated ATT for this road type and co-occurrence bin. This sample of congestion functions illustrates a few key points that are generally true across most cities and co-occurrence bins. First, congestion functions for local and arterial roads are lower in both levels and slopes than those for highways, so reductions in density have more limited effects on speeds. Second, congestion functions for some roads—especially highways—are non-linear, and the effects of any changes in density depend on the slope at the relevant levels of density. A highway that is already operating near the speed limit

[20] These congestion functions are subtly different than those studied in traffic engineering, because we use data from a collection of heterogeneous road segments within a group rather than a single segment. As a result, different levels of density in our data may correspond to different patterns of crowding across roads – some segments may be near free-flow while others are heavily congested. The estimated function should therefore be interpreted as a reduced-form mapping between average speeds and average densities, rather than a physical law of traffic flow on a particular road.

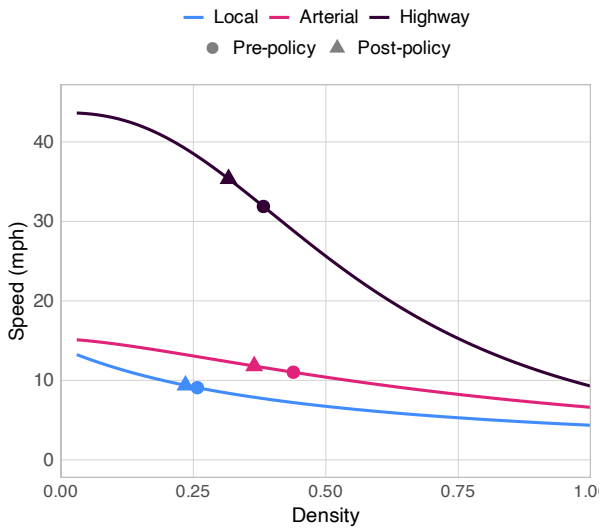
will see little gain from further reductions in density.

Figure 8: Example Congestion Functions

(a) CBD roads (NYC)



(b) 80-60% co-occurrence roads (NYC)



Notes: This figure plots estimated congestion functions for roads in NYC. The congestion functions are estimated using pre-period observed speeds and densities, based on Equation (17). Circles are added to each line based on the raw average speed before the policy's implementation. Triangles are based on the post-period speed, which is computed as the pre-period speed plus the estimated ATT on speeds.

Simulating potential effects in other cities. We now combine our different ingredients to evaluate the potential effects of introducing congestion pricing in other cities. To do so, we simulate a scenario in which each city has introduced a price that reduced trips to the CBD by 10% from all origins (with no change in other volumes), and compute the estimated total time savings and the 'break-even' price for which this policy would be welfare improving.

For the average trip in a given OD pair, we estimate the change in average travel time by summing over co-occurrence bins c and road types r and using the congestion functions (v^{cr}) to infer the time savings:

$$\Delta\hat{t}_{od} = \sum_{c}\sum_{r} \bar{t}_{od}^{cr} \times \left(\frac{v^{cr}(\rho^{cr} - 0.1 \times c \times \rho^{cr})}{v^{cr}(\rho^{cr})} - 1 \right) \tag{18}$$

where ρ^{cr} is the pre-period average density for this co-occurrence and road type combination. To compute the expected density under the counterfactual world with congestion pricing, we subtract $10\% \times c$ from pre-period density, where c is the midpoint of the corresponding co-occurrence bin. Note that this uses the assumption that the simulated decrease in CBD trips is uniform across origins, such that the decrease in cars on any given segment is proportional to its co-occurrence with the CBD (i.e. for a 10% decrease in CBD trips, a segment with 50% co-occurrence would have 5% fewer cars in the post-period).

We aggregate over ODs to compute the total time savings as

$$\Delta T = \sum_{o} \sum_{d} N_{od}^{\text{PV}} \times \Delta \hat{t}_{od} \tag{19}$$

where N_{od}^{PV} are the pre-period passenger vehicle trip volumes from Replica. We focus just on passenger vehicle trips for this exercise, as taxis are a small share of the non-NYC markets.

The price that achieves a 10% decrease in volumes may vary by city depending on the demand for CBD travel. While we cannot estimate the exact price required without a full model of driver decisions, we compute the 'break-even' price, i.e. the price for which the estimated time savings would exactly offset the costs paid by drivers to the CBD. We assume a VOTT of $40/hour and use the average number of pre-period CBD passenger vehicle trips from Replica. The final two columns of Table 1 document the total time savings and the break-even prices in each city. [Results pending].

8 Conclusion

The introduction of congestion pricing in New York City offers new evidence on how cordon-based pricing can reshape urban mobility in ways that extend well beyond the tolled area. While the policy had the expected direct effect of improving speeds in the CBD, the larger finding is that the greatest welfare gains arise from spillovers onto unpriced trips *outside* of the CBD. Because the road networks are interconnected, reductions in CBD-bound traffic have increased speeds across the NYC metro area. This highlights that evaluations limited to the cordon itself understate the effects of congestion pricing, as unpriced travelers may be the primary beneficiaries.

The welfare implications underscore a tension central to the political economy of congestion pricing. Priced drivers face salient, concentrated losses unless their value of time is very high, while unpriced travelers enjoy diffuse, less visible benefits. That dynamic may explain both the opposition that delayed New York's program for decades and the persistence of concerns over its distributional impact. We find broadly similar speed improvements across neighborhoods with different household income levels or in specific areas of interest, such as New Jersey, suggesting that equity concerns tied to travel times may be overstated in the short run.

Whether similar effects can be expected in other cities depends, in part, on each city's travel patterns and road network. [Results pending].

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." *Journal of the American Statistical Association*, 105(490): 493–505.

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel G. Hickok, and Tobias Salz. 2024. "Optimal Urban Transportation Policy: Evidence from Chicago." NBER Working Paper No. 32185.

Arnott, Richard, André de Palma, and Robin Lindsey. 1993. "A Structural Model of Peak-Period Congestion: A Traffic Bottleneck with Elastic Demand." *The American Economic Review*, 83(1): 161–179. Publisher: American Economic Association.

Ater, Itai, Adi Shany, Brad Ross, Eray Turkel, and Shoshana Vasserman. 2025. "Can Usage-Based Pricing Reduce Congestion?" *Working paper*.

Barkjohn, Karoline K., Amara L. Holder, Samuel G. Frederick, , and Andrea L. Clements. 2022. "Correction and Accuracy of PurpleAir PM2.5 Measurements for Extreme Wildfire Smoke." *Sensors*, 22(24): 9669.

Barwick, Panle Jia, Shanjun Li, Andrew Waxman, Jing Wu, and Tianli Xia. 2024. "Efficiency and Equity Impacts of Urban Transportation Policies with Equilibrium Sorting." *American Economic Review*, 114(10): 3161–3205.

Bento, Antonio, Kevin Roth, and Andrew R. Waxman. 2020*a*. "Avoiding Traffic Congestion Externalities? The Value of Urgency." https://www.nber.org/papers/w26956.

Bento, Antonio, Kevin Roth, and Andrew R. Waxman. 2020*b*. "Avoiding Traffic Congestion Externalities? The Value of Urgency." NBER Working Paper No. 26956.

BLS. 2025. "Occupational Employment and wages in New York-Newark-jersey city - May 2024." https://www.bls.gov/regions/northeast/news-release/occupationalemploymentandwages_newyork.htm.

Brooker, Aaron, Jeffrey Gonder, Lijuan Wang, Eric Wood, Sean Lopp, and Laurie Ramroth. 2015. "FASTSim: A model to estimate vehicle efficiency, cost and performance." *SAE Technical Paper*.

Brugge, Doug, John L Durant, and Christine Rioux. 2007. "Near-highway pollutants in motor vehicle exhaust: a review of epidemiologic evidence of cardiac and pulmonary health risks." *Environmental health*, 6: 1–12.

Buckeridge, David L, Richard Glazier, Bart J Harvey, Michael Escobar, Carl Amrhein, and John Frank. 2002. "Effect of motor vehicle emissions on respiratory health in an urban area." *Environmental health perspectives*, 110(3): 293–300.

Clark Center Forum. 2024. "Congestion Pricing in New York." University of Chicago.

Cook, Cody, and Pearl Z. Li. 2024. "Value Pricing or Lexus Lanes? The Distributional Effects of Dynamic Tolling." *Working paper*.

Currie, Janet, and Reed Walker. 2011. "Traffic congestion and infant health: Evidence from E-ZPass." *American Economic Journal: Applied Economics*, 3(1): 65–90.

Currier, Lindsey, Edward L Glaeser, and Gabriel E Kreindler. 2023. "Infrastructure Inequality: Who Pays the Cost of Road Roughness?" National Bureau of Economic Research.

Davis, Lucas W. 2008. "The Effect of Driving Restrictions on Air Quality in Mexico City." *Journal of Political Economy*, 116(1).

De Palma, Andre, and Robin Lindsey. 2000. "Private toll roads: Competition under various ownership regimes." *The Annals of Regional Science*, 34(1): 13–35.

Department for Energy Security and Net Zero. 2023. "Greenhouse gas articleing: Conversion factors 2023." https://www.gov.uk/government/publications/greenhouse-gas-articleing-conversion-factors-2023.

Department of Energy (DOE). 2004. "Success Story: PSAT Accurately Simulates Advanced Vehicles." FreedomCAR and Vehicle Technologies Program.

Department of Energy (DOE). 2020. "Vehicle energy consumption and performance analysis." https://www.autonomie.net/expertise/VehicleEnergyConsumptionAndPerformanceAnalysis.html.

Durrmeyer, Isis, and Nicolás Martínez. 2024. "The Welfare Consequences of Urban Traffic Regulations." *Working paper.*

Ecola, Liisa, and Thomas Light. 2009. "Equity and congestion pricing." *Rand Corporation*, 1–45.

Eliasson, Jonas, Lars Hultkrantz, Lena Nerhagen, and Lena Smidfelt Rosqvist. 2009. "The Stockholm congestion–charging trial 2006: Overview of effects." *Transportation Research Part A: Policy and Practice*, 43(3): 240–250.

Ersal, Tulga, Jeffrey Stein, Zoran Filipi, Huei Peng, Dennis Assanis, Hosam Fathy, and Lucas Louca. 2012. "VESIM - Automotive Research Center (ARC) VEhicle SIMulation." https://deepblue.lib.umich.edu/handle/2027.42/98080.

Faris, W., R. Kafafy, M. Idres, and S. Elmoselhy. 2011. "Vehicle fuel consumption and emission modelling: An in-depth literature review." *International Journal of Vehicle Systems Modelling and Testing*, 6: 318–395.

Galdon-Sanchez, Jose Enrique, Ricard Gil, Felix Holub, and Guillermo Uriz-Uharte. 2022. "Social Benefits and Private Costs of Driving Restriction Policies: The Impact of Madrid Central on Congestion, Pollution, and Consumer Spending." *Journal of the European Economic Association*, 21(3): 1227–1267.

Gallego, Francisco, Juan-Pablo Montero, and Christian Salas. 2013. "The Effect of Transport Policies on Car Use: Evidence from Latin American Cities." *Journal of Public Economics*, 107: 47–62.

Gibson, Matthew, and Maria Carnovale. 2015. "The effects of road pricing on driver behavior and air pollution." *Journal of Urban Economics*, 89: 62–73.

Gordon, Aaron, Priyanjana Bengani, Jeff Kao, Marie Patino, and Jason Kao. 2025. "Is NYC's Congestion Pricing Working? Fewer Private Cars Are On the Road (For Now)." Bloomberg.

Greenberg, Harold. 1959. "An analysis of traffic flow." *Operations research*, 7(1): 79–85.

Green, Colin P., John S. Heywood, and Maria Navarro Paniagua. 2020. "Did the London congestion charge reduce pollution?" *Regional Science and Urban Economics*, 84.

Greenshields, Bruce D, J Rowland Bibbins, WS Channing, and Harvey H Miller. 1935. "A study of traffic capacity." In *Highway research board proceedings*. Vol. 14, 448–477, Washington, DC.

Hall, Jonathan D. 2018. "Pareto Improvements from Lexus Lanes: The effects of Pricing a Portion of the Lanes on Congested Highways." *Journal of Public Economics*, 158: 113–125.

Herzog, Ian. 2024. "The City-Wide Effects of Tolling Downtown Drivers: Evidence from London's Congestion Charge." *Journal of Urban Economics*, 144.

Hierons, Thomas. 2024. "Spreading the Jam: Optimal Congestion Pricing in General Equilibrium." *Working paper.*

Holden, Jacob, Nicholas Reinicke, and Jeffrey Cappellucci. 2020. "RouteE: A vehicle energy consumption prediction engine." *Society of Automotive Engineers Technical Paper Series*, 2(NREL/JA-5400-78089).

Hu, Winnie, Ana Ley, and Nate Schweber. 2025. "Congestion Pricing Results Are Mixed but Some Commutes Improve." The New York Times.

Hu, Winnie, and Ana Ley. 2024. "How the Resurrected Congestion Pricing Plan Could Die in the Courts." The New York Times.

Knittel, Christopher R, Douglas L Miller, and Nicholas J Sanders. 2016. "Caution, drivers! Children present: Traffic, pollution, and infant health." *Review of Economics and Statistics*, 98(2): 350–366.

Leape, Jonathan. 2006. "The London Congestion Charge." *Journal of Economic Perspectives*, 20(4): 157–176.

Ley, Ana. 2022. "A Plan to Push Cars Out of Manhattan Could Make the Bronx's Air Dirtier." *The New York Times*.

Ley, Ana, Winnie Hu, and Keith Collins. 2025. "Less Traffic, Faster Buses: Congestion Pricing's First Week." *The New York Times*.

Liu, Shi V, Fu-lin Chen, and Jianping Xue. 2019. "A meta-analysis of selected near-road air pollutants based on concentration decay rates." *Heliyon*, 5(8).

Li, Zhenlong, Huan Ning, Fengrui Jing, and M. Naser Lessani. 2024. "Understanding the bias of mobile location data across spatial scales and over time: A comprehensive analysis of SafeGraph data in the United States." *PLOS ONE*, 19(1): e0294430. Publisher: Public Library of Science.

Montero, Juan-Pablo, Felipe Sepúlveda, and Leonardo Basso. 2025. "Pricing Congestion to Increase Traffic: The Case of Bogota." *Journal of the European Economics Association*.

Moshes, Joshua, and Benjamin Moshes. 2025. "Congestion Pricing Tracker." https://www.congestion-pricing-tracker.com/.

MTA. 2022. "Central Business District (CBD) Tolling Program Environmental Assessment Executive Summary." https://www.mta.info/document/162396.

MTA. 2025*a*. "Congestion Relief Zone Tolling Week One Update." https://www.mta.info/document/162396.

MTA. 2025*b*. "MTA Finance Committee: Financial Performance Report." https://metrics.mta.info/?cbdtp/vehiclereductions.

MTA. 2025*c*. "Reduction in Vehicle Entries to the CBD." https://www.mta.info/document/170771.

National Institute of Environmental Health Sciences. 2024. "Air Pollution and Your Health." *Environmental Health Topics*.

Nilsson, Peter, Matthew Tarduno, and Sebastian Tebbe. 2024. "Road Pricing with Green Vehicle Exemptions: Theory and Evidence." *CESifo Working Paper*.

NYC Department of City Planning. 2011. "Public Parking Study." Manhattan Core.

Ostrovsky, Michael, and Frank Yang. 2024. "Effective and Equitable Congestion Pricing: New York City and Beyond." *Working paper*.

Pigou, Arthur Cecil. 1920. *The Economics of Welfare.* Macmillan. Google-Books-ID: 10kPAQAAIAAJ.

Seo, Toru, Alexandre M Bayen, Takahiko Kusakabe, and Yasuo Asakura. 2017. "Traffic state estimation on highway: A comprehensive survey." *Annual reviews in control*, 43: 128–151.

Simeonova, Emilia, Janet Currie, Peter Nilsson, and Reed Walker. 2021. "Congestion Pricing, Air Pollution, and Children's Health." *Journal of Human Resources*, 56(4): 971–996. Publisher: University of Wisconsin Press Section: Articles.

Sleiman, Léa Bou. 2024. "Displacing Congestion: Evidence from Paris." *Working Papers*.

Small, Kenneth A., Clifford Winston, and Jia Yan. 2005. "Uncovering the Distribution of Motorists' Preferences for Travel Time and Reliability." *Econometrica*, 73(4): 1367–1382. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2005.00619.x.

Tang, Cheng Keat, and Jos van Ommeren. 2022. "Accident externality of driving: evidence from the London Congestion Charge." *Journal of Economic Geography*, 22(3): 547–580.

Taylor, Brian D. 2010. "How fair is road pricing? Evaluating equity in transportation pricing and finance." National Transportation Policy Project.

Tschofen, Peter, Inês L. Azevedo, and Nicholas Z. Muller. 2019. "Fine Particulate Matter Damages and Value Added in the US Economy." *Proceedings of the National Academy of Sciences*, 116(40): 19857–19862.

Verhoef, Erik, Peter Nijkamp, and Piet Rietveld. 1996. "Second-best congestion pricing: the case of an untolled alternative." *Journal of Urban Economics*, 40(3): 279–302.

Verhoef, Erik T. 2002. "Second-best congestion pricing in general static transportation networks with elastic demands." *Regional Science and Urban Economics*, 32(3): 281–310.

Vickrey, William. 1969. "Congestion Theory and Transport Investment." *American Economic Review*, 59(2): 251–60. Publisher: American Economic Association.

Vickrey, William S. 1963. "Pricing in Urban and Suburban Transport." *The American Economic Review*, 53(2): 452–465. Publisher: American Economic Association.

Xu, Yiqing. 2017. "Generalized synthetic control method: Causal inference with interactive fixed effects models." *Political Analysis*, 25(1): 57–76.

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." *American Economic Journal: Economic Policy*, 12(1): 418–453.

Appendices

A Data Appendix **34**

A.1 CBD Definitions 34

A.2 Estimating Fuel Efficiency 34

A.3 PurpleAir Pollution Measures 36

A.4 Consumer Spending Data from MBHS3 37

A.5 Foot Traffic Data from Advan 37

B Welfare Calculations Appendix **39**

B.1 Origin-Destination Volumes 39

B.2 Illustrative Example 40

B.3 Welfare Calculation Inputs 41

C Supplementary Tables and Figures **43**

A Data Appendix

A.1 CBD Definitions

Table A.1 contains references to the CBD definitions we use for the control cities and New York City. Each links to a snapshot of the corresponding web page on the Internet Archive. This avoids links breaking due to website changes of the city-affiliated organizations from which we derive the CBD definitions.

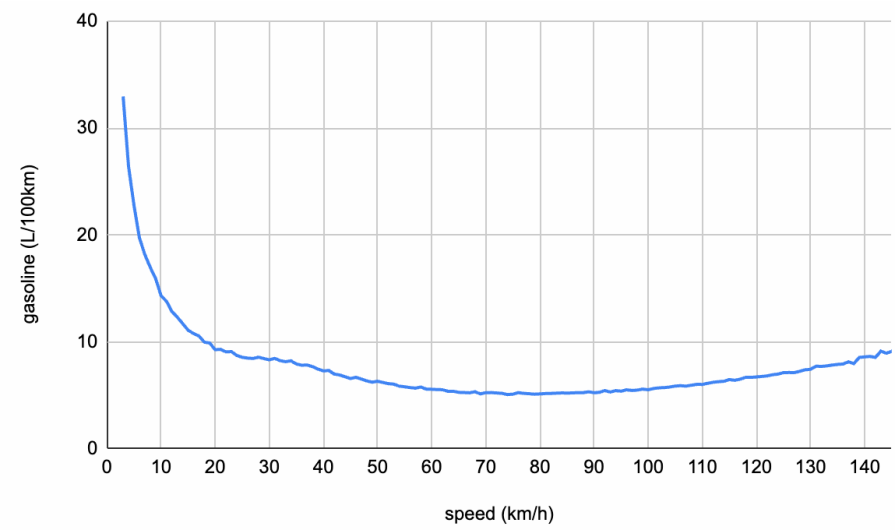
Table A.1: CBD Definitions Overview

Link	City name	Defined by	CBD shape
🔗	Atlanta	Atlanta Downtown	The "Downtown Boundary".
🔗	Baltimore	Downtown Partnership of Baltimore	The "Downtown Management Authority" shape on page 5.
🔗	Boston	City of Boston	The outline of the "Downtown" neighborhood.
🔗	Chicago	City of Chicago	The "Downtown Zone Area".
🔗	New York City	Metropolitan Transportation Authority	The outline of the "Congestion Relief Zone".
🔗	Philadelphia	Center City District Philadelphia	The boundaries of the district on page 8.

A.2 Estimating Fuel Efficiency

We evaluate the environmental impact of this experiment by measuring fuel consumption rates on road segments. CO_2 emission rates then naturally arise from these estimates of fuel consumption, as carbon emissions are typically modeled as proportional to the fuel consumption in transportation settings (Department for Energy Security and Net Zero, 2023). Since direct measurements of fuel or CO_2 emissions are impractical, we use scalable methods for estimating fuel consumption rates instead.

Figure A.1: Example Speed-Fuel Consumption Rate Curve



Notes: This figure documents an example of the empirical relationship between speeds and fuel consumption for one class of roads. Similar data are used to model the overall relationship on various roads and road conditions.

Fuel consumption modeling has been widely studied in the literature (Faris et al., 2011). Models of this form roughly fall into two categories: principled models (Faris et al., 2011), which aim to model the physics underlying energy usage, and empirical models (Department of Energy , DOE; Ersal et al., 2012; Department of Energy , DOE), which fit often non-parametric models to ground-truth fuel consumption data. For this study, we collaborate with *National Renewable Energy Laboratory* (NREL), integrating their models which fall roughly under both categories. These models at their core rely on FASTSim (Brooker et al., 2015), a physics-based simulator (hence a principled model) which calculates the power required to meet a given drive cycle speeds provided other inputs such as road grade and vehicle specifications such as drag, transmission, and rolling resistance. Its methodology and data are validated from dynamometer testing data via collaboration with other labs (e.g., Argonne National Laboratory), so this is a high-fidelity model with many parameters to calibrate towards specific vehicle models. However, FASTSim requires significant computation power and high frequency GPS location data, which makes it challenging to run for all segments or trips. To address this issue, we use an empirical machine learning model on top of FASTSim, similar to NREL's RouteE model (Holden, Reinicke and Cappellucci, 2020). This family of models significantly reduces the computational burden and works well with segment-level speeds, eliminating the need for high-fidelity GPS location data. The ML-based model takes as features the properties from segments and estimates the fuel consumption for each segment. Features commonly used by these models include the segment-level speeds, road grade, and length.

Figure A.1 depicts the empirical relationship between speed and fuel consumption on one class of roads, which is used to model the overall relationship. Notably, the convex shape of the model is a commonly known feature by energy modeling practitioners, and denotes that vehicles operating at intermediate driving speeds generally experience the highest levels of fuel efficiency.

A.3 PurpleAir Pollution Measures

For each city, we collect hourly $PM_{2.5}$ CF1 and relative humidity readings for all PurpelAir sensors in the corresponding CBSA. For redundancy, each sensor includes two independent measurement channels, labeled as "A" and "B" channels. We fill in missing relative humidity readings of any one channel with the readings from the other channel and censor any outlier $PM_{2.5}$ values larger than 300 $\mu g/m^3$.

We follow EPA's PurpleAir data calibration methodology to improve the quality of our air pollution data (Barkjohn et al., 2022), taking the following steps to build a dataset of sensor-day level $PM_{2.5}$ level.

1. We drop records for which the difference between the A and B channel $PM_{2.5}$ values are greater than 5 $\mu g/m^3$ or the relative percentage difference is greater than 70% (computed with the lowest value in the denominator).

2. Still following Barkjohn et al. (2022), we compute the quantities

$$\begin{aligned} x_1 &= 0.524\ PM_{2.5,\ raw} - 0.0862\ RH + 5.75 \\ x_2 &= 4.21 \cdot 10^{-4}\ PM_{2.5,\ raw}^2 + 0.392\ PM_{2.5,\ raw} + 3.44 \\ x_3 &= 0.0244\ PM_{2.5,\ raw} - 13.9, \end{aligned} \tag{A.1}$$

   where $PM_{2.5,\ raw}$ is the raw hourly CF1 $PM_{2.5}$ channel reading and RH is the relative humidity channel reading;

   We then use these values to compute the estimated hourly $PM_{2.5}$ for a specific channel in a given sensor as

$$PM_{2.5} = \begin{cases} x_1 & \text{if } PM_{2.5,\ raw} < 570 \\ x_1\,(1 - x_3) + x_2 x_3 & \text{if } 570 \leq PM_{2.5,\ raw} < 611 \\ x_2 & \text{if } 611 \leq PM_{2.5,\ raw} \end{cases} \tag{A.2}$$

3. We aggregate to the sensor-hour level by taking the average of the channel-specific adjusted levels of $PM_{2.5}$, or taking the value of just one of the two channels' hourly adjusted $PM_{2.5}$ levels when it is missing for the other channel.

4. Finally, we compute the daily adjusted $PM_{2.5}$ averages, dropping any comprising of fewer than 18 hours of observations.

The outcome variable we use in our analysis of the air pollution effects in the NYC CBD is that of sensor-date level $PM_{2.5}$. Some sensors are missing information for one or more days. We impute missing sensor-date $PM_{2.5}$ values using a regression with sensor $\times$ month-of-sample and sensor $\times$ day-of-week fixed effects.

A.4 Consumer Spending Data from MBHS3

The data from MBHS3 are aggregated to the zipcode, date, and 3-digit NAICS code level. The zipcode is based on the merchant's physical location. The data exclude online transactions. Table A.2 documents the total number of transactions and aggregate transaction amount by category. The final three columns subset to zipcodes within the CBD. The average transaction size in both restaurant and retail categories tends to be higher for merchants in the CBD than in the overall city.

We define zipcodes as lying within a city's CBD if its centroid is contained by the CBD boundaries. For NYC, we manually include zipcode 10004. This zipcode includes Governor's Island, which pulls it's centroid out into the river, but most of its merchants are in the southern portion of the Financial District. Figure A.2 Panel a) maps the 'treated' zipcodes in the NYC CBD. The Baltimore CBD is small and there are no zipcodes whose centroids lie within it, so we do not report any spending for Baltimore's CBD.

Table A.2: Transaction Counts and Amounts: January 2024 - April 2025

		All			**CBD-only**		
City	Category	Count (millions)	Amount (millions $)	Avg. size ($/trans.)	Count (millions)	Amount (millions $)	Avg. size ($/trans.)
NYC	Restaurant	164.7	4613.2	28.0	26.1	941.7	36.0
	Retail	286.1	13247.0	46.3	39.6	2196.0	55.4
PHL	Restaurant	88.5	2291.9	25.9	3.9	125.3	32.1
	Retail	168.6	7549.0	44.8	3.2	115.0	35.5
CHI	Restaurant	114.2	2773.8	24.3	5.8	172.5	29.7
	Retail	148.2	7425.8	50.1	3.7	137.5	37.3
BOS	Restaurant	59.2	1433.1	24.2	0.8	25.6	32.4
	Retail	68.4	3390.8	49.6	0.2	15.9	79.2
BAL	Restaurant	61.7	1066.6	17.3	–	–	–
	Retail	58.4	2644.8	45.3	–	–	–
ATL	Restaurant	141.7	3368.8	23.8	3.6	90.3	25.2
	Retail	190.0	8995.0	47.3	1.2	43.2	37.2
Total	Restaurant	629.9	15547.4	24.7	40.2	1355.6	33.7
	Retail	919.7	43252.4	47.0	47.9	2507.5	52.3

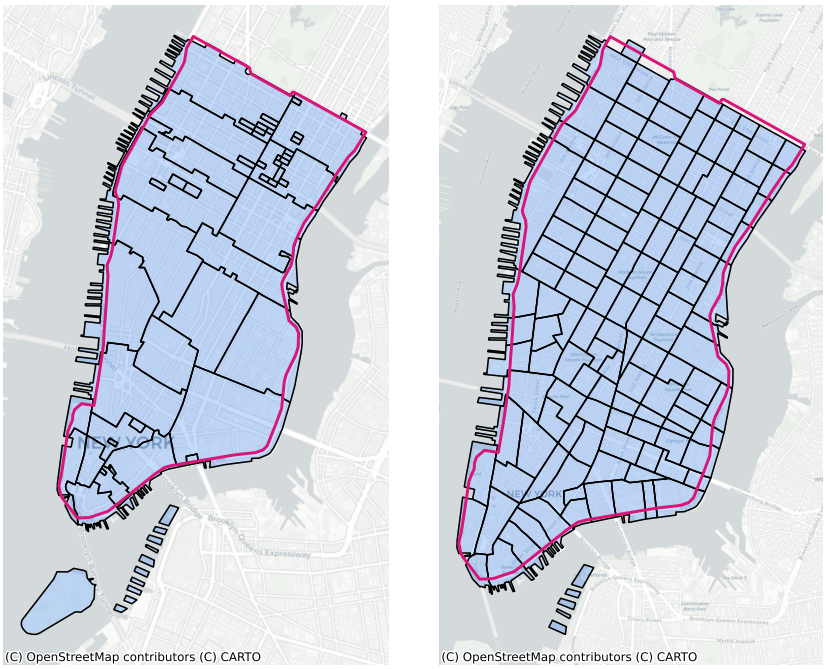
Notes: This table documents the aggregate number of transactions and total spending in the MBHS3 data for zipcodes within each of the sample cities between January 2024 and April 2025. The Baltimore CBD is small and there are no zipcodes whose centroids lie within it.

A.5 Foot Traffic Data from Advan

The raw data from Advan Neighborhood Patterns include the number of stops by GPS devices at the hour-block group level. We restrict to peak hours and combine across all block groups within the same Census tract. Figure A.3 plots the average daily foot traffic in the CBDs of our sample cities. There is a clear drop in all cities at the start of 2025. The drop is substantial in magnitude—about 30%—and is likely due to a change in the sources from which Advan purchases location data. The decline in measured foot traffic is similar in magnitude and timing across all of our sample cities.

Figure A.2: Intersection of CBD and Zipcodes/Tracts

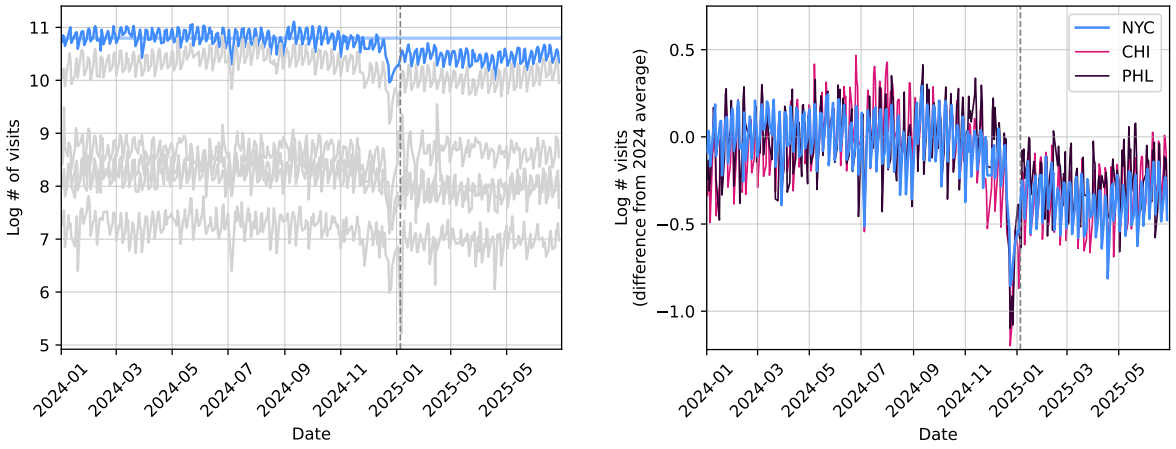
(a) Treated zipcodes (b) Treated tracts



Notes: This figure documents the Census tracts and zipcodes that are defined as 'treated' by the policy because their centroids lie within the CBD. Treated units are shaded blue and the boundary of the CBD is in pink. We manually include zip 10004, which includes Governor's Island and parts of the Financial District as its centroid lies outside the CBD but the vast majority of its merchants are within the CBD.

Figure A.3: CBD Foot Traffic

(a) Log average stops per day (b) Log stops relative to 2024 average



Notes: These figures plot daily CBD foot traffic, restricted to weekday peak hours (5am to 9pm). The first panel plots the log total number of visits. The blue horizontal line is the pre-period average for the NYC CBD. The second panel restricts attention to the three largest cities and plots the difference in log trips between a given date and the 2024 average.

B Welfare Calculations Appendix

B.1 Origin-Destination Volumes

One limitation of our data is that we cannot reliably observe volumes, much less changes in volumes over time. For computing the aggregate welfare effects in Section 6, we need the number of personal vehicle and taxi/FHV trips between each set of origins and destinations. We estimate these values using a combination of data from Replica and the MTA.

Our starting point are estimated tract-to-tract flows for a 'typical' weekday in the fourth quarter of 2024 from Replica, an data platform company that spun out of Google. Replica ingests data from dozens of sources, including location data from GPS devices, traffic patterns from in-road sensors, administrative data on local demographics, parcel land use, public transit usage, taxi/FHV trips, travel surveys, and more. Replica then builds population-scale simulations that aim to accurately capture the population and travel patterns of metro areas. Underlying these simulations are activity-based models trained to match the 'ground truth' data (e.g., the number of cars driving over a specific road segment with an embedded sensor). Key for our use case, Replica estimates the mode choice for each trip, including separately estimating use of private autos and taxi/FHVs.

We supplement data from Replica with data from the MTA on CBD entries. CBD entries are an especially important source of welfare in our setting, and the MTA provides realized flows. First, we use data on realized entries since congestion pricing launched, split by entry point, 10-minute interval, and vehicle class. We use these data to compute the number of taxi/FHVs ("TLC Taxi/FHV") and passenger car ("Cars, Pickups and Vans") entries during peak hours on the average weekday. Together, these two categories account for 85% of all entries (Table B.1). We scale the estimated entries by Replica to match these the passenger vehicle and taxi/FHV counts. Second, as these entries are for 2025 but our desired volumes are those in the absence of congestion pricing, we further scale each CBD entry estimate up by 13% to match the baseline CBD entries reported by in MTA (2025*b*). This approach essentially combines Replica data on relative distribution of origin tracts for CBD trips with MTA data on the overall levels.

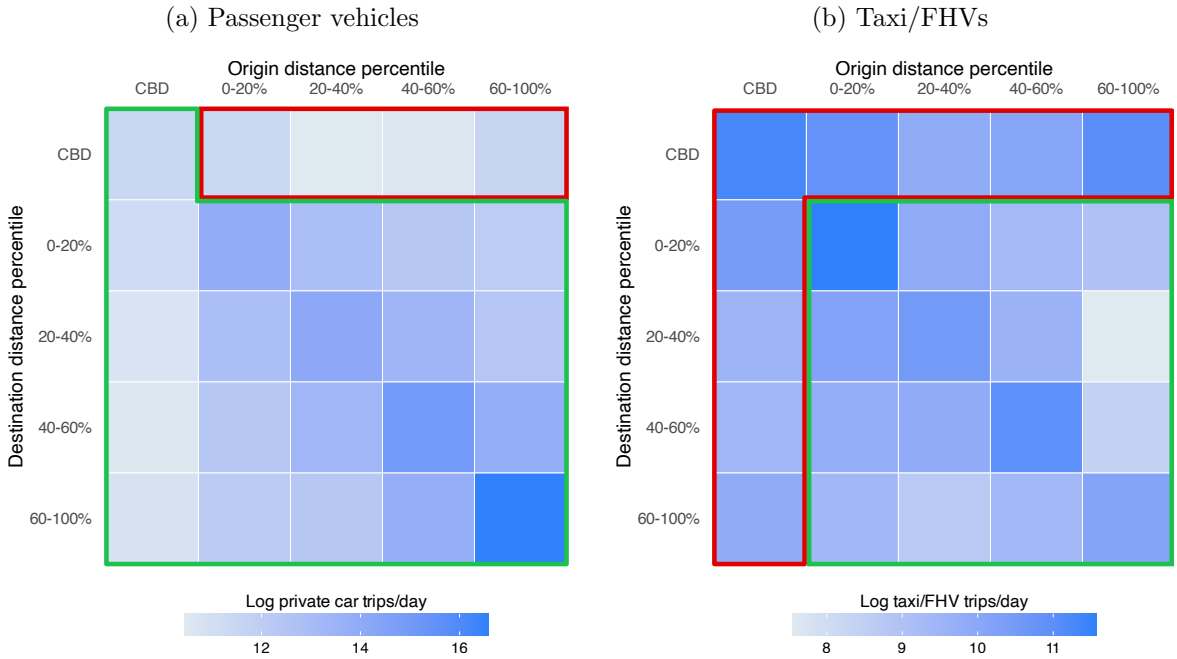
Finally, while we do not use within-day variation in entries, Figure B.2 shows the number of entries to the CBD by 10-minute interval on the average day, including all vehicle classes. There is some clear bunching just before and after the peak pricing time.

Table B.1: CBD entries by vehicle class

Vehicle class	Avg. daily entries	Frac. of all entries
Cars, Pickups and Vans	247934	0.594
Taxi/FHV	142775	0.342
Single-Unit Trucks	16365	0.039
Buses	7882	0.019
Motorcycles	1636	0.004
Multi-Unit Trucks	1126	0.003

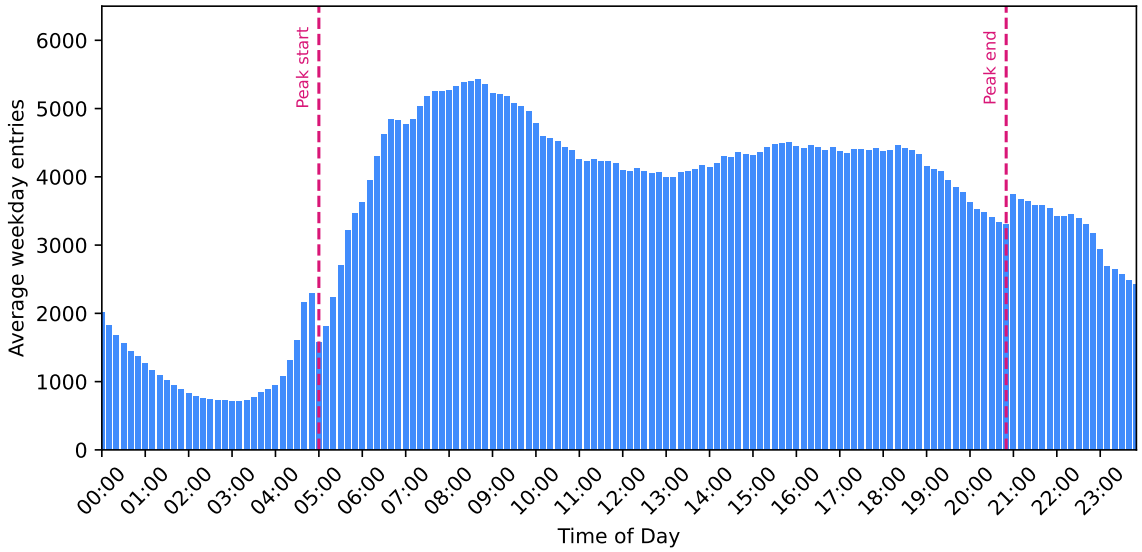
Notes: This table documents the average entries into the CBD for weekday peak hours between Jan 5, 2025 to June 30, 2025.

Figure B.1: Estimated Daily Trip Counts



Notes: These figures plot the log daily estimated trip counts between bins of origins and destinations. The trip counts are estimated using a combination of data from Replica and the MTA and reflect pre-congestion pricing flows.

Figure B.2: CBD Entries by Time of Day (MTA)



Notes: This figure plots the average number of weekday CBD entries by 10-minute intervals. The underlying data are from the MTA and include all vehicle classes.

B.2 Illustrative Example

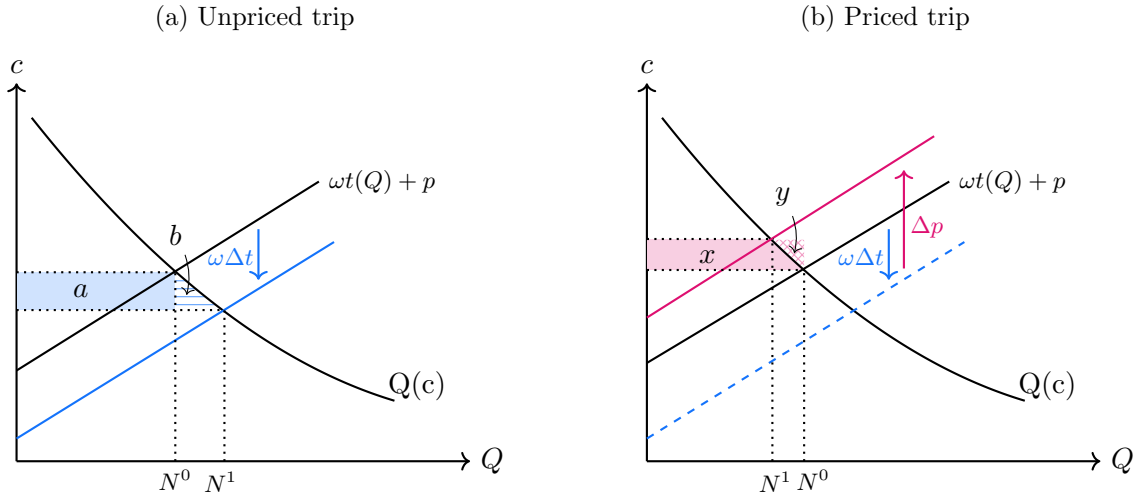
We provide an illustrative example of the argument underlying the welfare bounds derived in Section 6. Figure B.3 illustrates the effects of congestion pricing for a representative priced and unpriced trip. The y-axes are the cost of the trip, inclusive of the value of the travel time $\omega t(Q)$ and the congestion

fee p, setting aside all other costs. Here, the travel time depends linearly on the quantity Q—which is often not the case for real-world congestion functions—so the marginal cost line is described by $c = \omega t(Q) + p$. The demand curve $Q(c)$ can be thought of as the CDF of outside option values (ε_i) across the population of prospective drivers for a trip. Although we allow the outside option values to change before and after the policy in the formal bounds, here we assume that the outside option value for each driver is fixed such that demand is unchanged pre- and post-policy.

For unpriced trips (Figure B.3a), introducing congestion pricing reduces travel time costs by $\omega\Delta t$. With elastic demand, this would encourage additional driving and move the equilibrium volume from N^0 to N^1. The true change in welfare stemming from the policy change is the area $a + b$. By using fixed volumes in our welfare calculation, we report welfare of just $\widehat{\Delta W}_{\text{unpriced}} = a$ for unpriced trips, which understates the true welfare gains by the area of b.

For priced trips (Figure B.3b), the countervailing price increase (Δp) offsets the improved travel times ($\omega\Delta t$) and, in this illustration, reduces total volumes. The true change in driver welfare is the area x. Here, by using fixed volumes, our welfare estimate of $\widehat{\Delta W}_{\text{priced}} = -x - y$ overstates the true welfare loss by the area y. If the gains due to time savings are greater than the losses due to paying the toll—as is the case for many taxi trips—then the figure would look more similar to Figure B.3a and we would understate the net gains of priced trips.

Figure B.3: Illustration of Driver Welfare Calculations



B.3 Welfare Calculation Inputs

Table B.2 documents the primary inputs for the welfare estimates presented in Section 6. The columns correspond to estimated treatment effect on log speeds ($\widehat{\text{ATT}}_{od}$), the average pre-period duration ($\bar{t}_{od}$), and the volumes (N^j_{od}) and average price (p^j_{od}). Note that we assume that prices are equal across all priced trips within a vehicle class and that each vehicle class faces the same travel times and ATTs conditional on an origin-destination pair.

Table B.2: Welfare estimate inputs

				Passenger Vehicles		Taxi/FHVs	
Origin	Destination	$\widehat{\text{ATT}}_{od}$	$\bar{t}_{od}$	N_{od}^{PV}	p_{od}^{PV}	N_{od}^{FHV}	p_{od}^{FHV}
CBD	CBD	0.065	23.6	96114	\$0	81616	\$1.00
CBD	0-20%	0.075	23.6	74760	\$0	35357	\$1.00
CBD	20-40%	0.062	45.0	46499	\$0	14228	\$1.00
CBD	40-60%	0.055	51.5	39417	\$0	12652	\$1.00
CBD	60-100%	0.042	56.0	51783	\$0	19933	\$1.00
0-20%	CBD	0.083	25.3	89467	\$7.90	50924	\$1.00
0-20%	0-20%	0.005	10.8	1126305	\$0	105315	\$0
0-20%	20-40%	0.012	11.5	391971	\$0	27024	\$0
0-20%	40-60%	0.012	14.0	237659	\$0	18111	\$0
0-20%	60-100%	0.023	24.0	181999	\$0	12245	\$0
20-40%	CBD	0.063	45.5	34772	\$7.90	18861	\$1.00
20-40%	0-20%	0.013	11.5	386737	\$0	20140	\$0
20-40%	20-40%	0.014	11.1	1400166	\$0	37744	\$0
20-40%	40-60%	0.016	11.2	625676	\$0	19312	\$0
20-40%	60-100%	0.013	11.8	236182	\$0	6067	\$0
40-60%	CBD	0.058	52.1	41253	\$7.90	24424	\$1.00
40-60%	0-20%	0.013	14.0	238968	\$0	11050	\$0
40-60%	20-40%	0.016	11.2	656745	\$0	14498	\$0
40-60%	40-60%	0.015	11.2	3366824	\$0	57407	\$0
40-60%	60-100%	0.012	11.4	972889	\$0	11656	\$0
60-100%	CBD	0.063	55.0	108412	\$7.90	60927	\$1.00
60-100%	0-20%	0.025	24.3	179963	\$0	8418	\$0
60-100%	20-40%	0.013	11.9	253298	\$0	1907	\$0
60-100%	40-60%	0.013	11.4	1040177	\$0	4600	\$0
60-100%	60-100%	0.008	11.3	15740553	\$0	25781	\$0

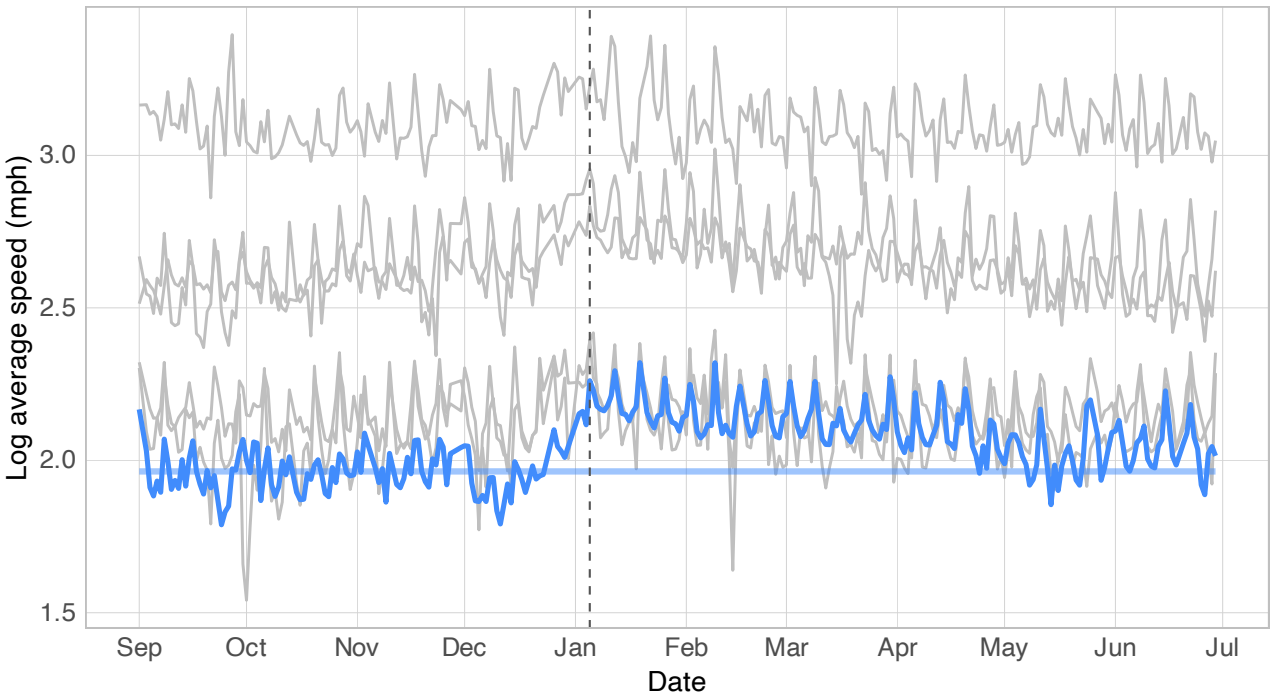
C Supplementary Tables and Figures

Table C.1: Average CBD speeds before and after congestion pricing

		Average speeds (mph)	
City	Road type	Before Jan 5, 2025	After Jan 5, 2025
New York	All	7.14	8.17
	Highway	16.95	18.45
	Arterial	7.19	8.31
	Local	5.54	6.24
Atlanta	All	22.37	22.26
	Highway	33.86	33.57
	Arterial	11.51	11.36
	Local	10.22	10.2
Baltimore	All	8.19	8.29
	Highway	8.54	8.83
	Arterial	8.3	8.28
	Local	7.6	7.84
Boston	All	13.99	14.82
	Highway	32.16	33.11
	Arterial	7.09	7.37
	Local	6.61	6.79
Chicago	All	13.57	13.88
	Highway	25.21	26.68
	Arterial	10.32	10.43
	Local	8.69	8.88
Philadelphia	All	8.56	8.77
	Highway	35.77	35.01
	Arterial	7.93	8.15
	Local	6.59	6.71

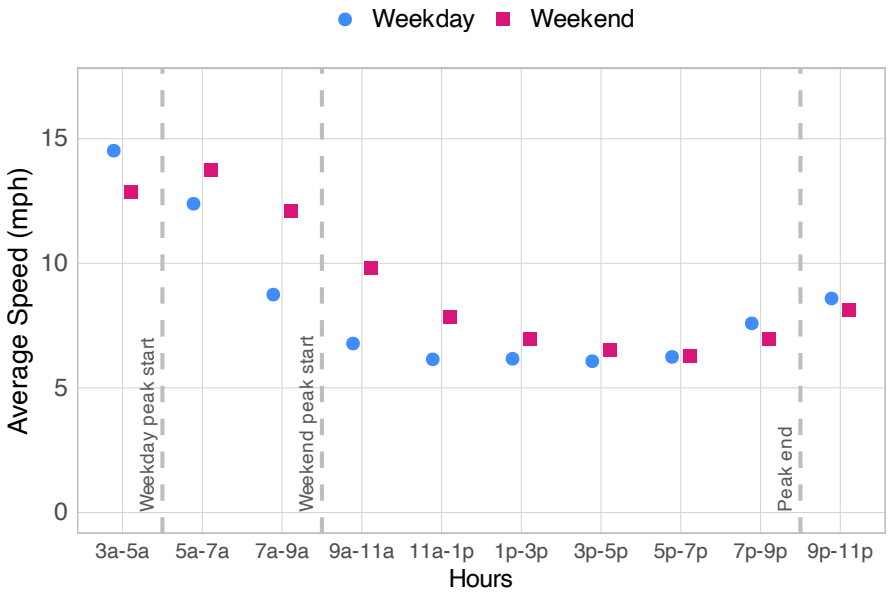
Notes: This table reports the average volume-weighted speeds on segments in each of the cities used throughout our analyses for weekday peak hours. The "before" period data starts September 1st, 2024 and the "after" period data ends on June 30th.

Figure C.1: Average CBD speeds



Notes: This figure documents the volume-weighted average daily speed on road segments in the CBD of each city. The horizontal lines represent the average speed between September 1 and December 15, 2024.

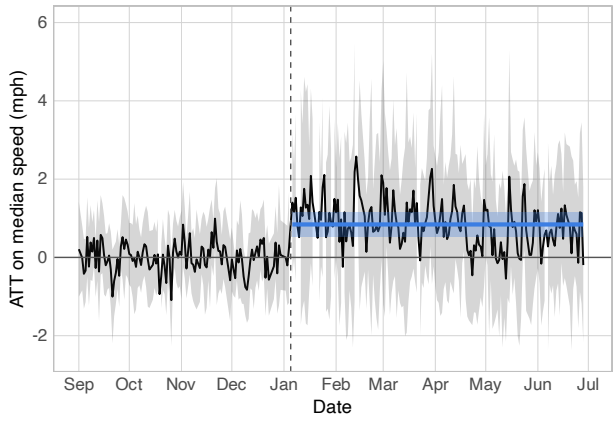
Figure C.2: Average segment speeds by hour of day



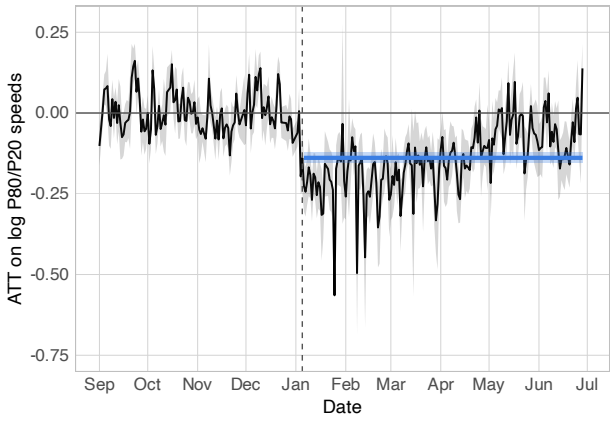
Notes: This figure reports the average pre-policy speeds in the NYC CBD by hour of the day and whether or not it is a weekend.

Figure C.3: Effects on CBD Median Speed and P90/P10 Speed

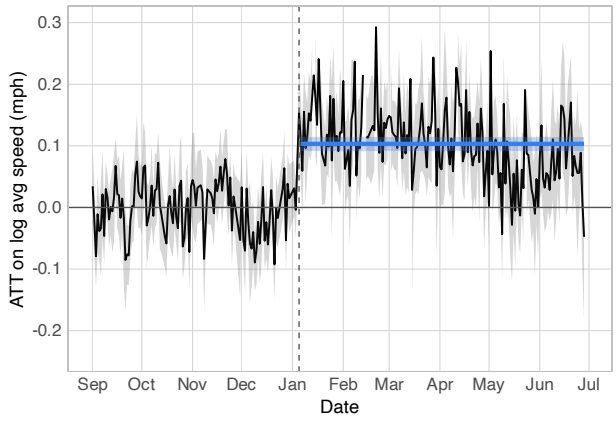
(a) ATT on median speed



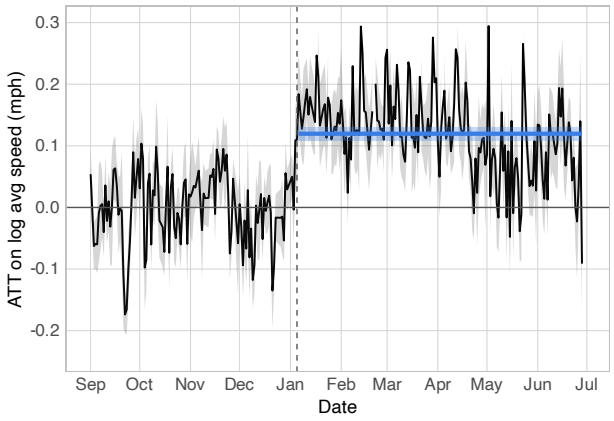
(b) ATT on p80/p20 speeds



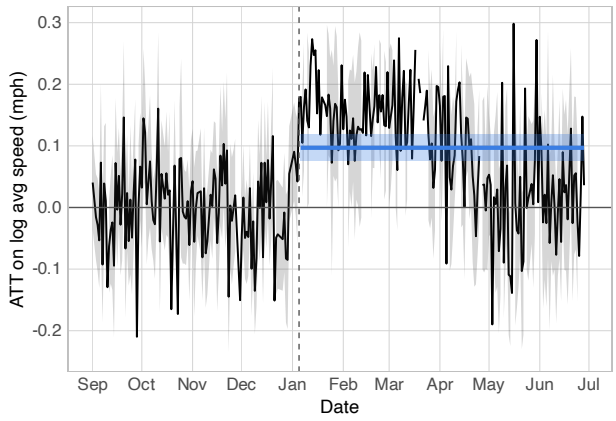
(c) Local roads: ATT on avg speeds



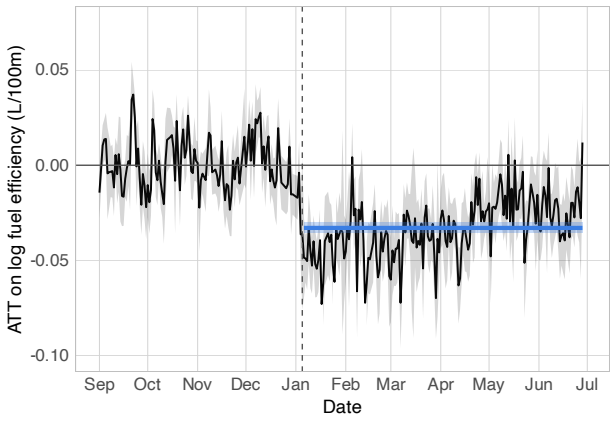
(d) Arterials: ATT on avg speed



(e) Highways: ATT on avg speed



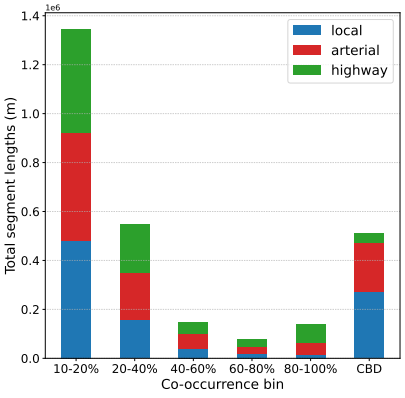
(f) ATT on fuel economy (L/100m)



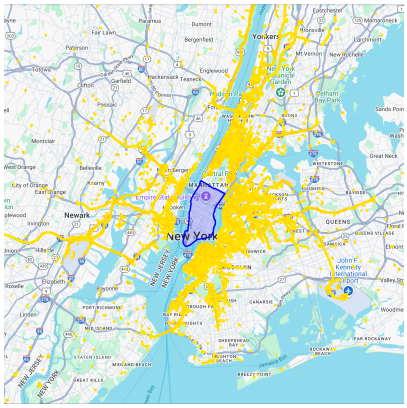
Notes: This figure replicates Figure 2a for additional CBD traffic measures. Each figure plots day-level ATTs of congestion pricing over time. The underlying data are two-hour bin outcomes in each CBD. The horizontal blue line is the aggregate ATT for all post-treatment periods. Shading denotes 95% confidence intervals. Standard errors are clustered at the city-level.

Figure C.4: Segments by bins of co-occurrence

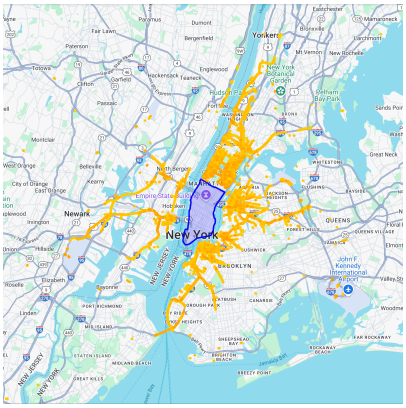
(a) Bin compositions



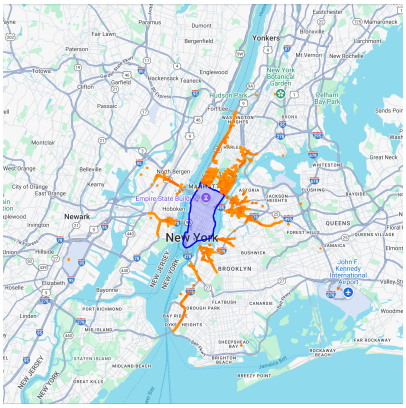
(b) 10-20%



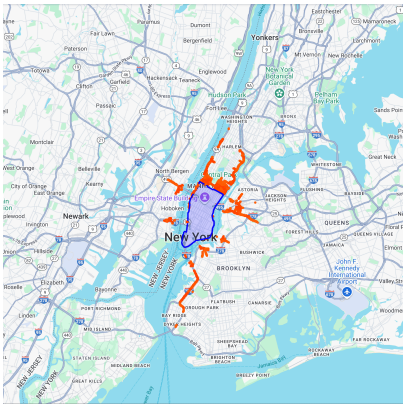
(c) 20-40%



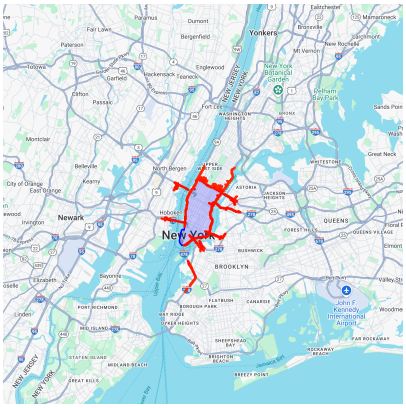
(d) 40-60%



(e) 60-80%

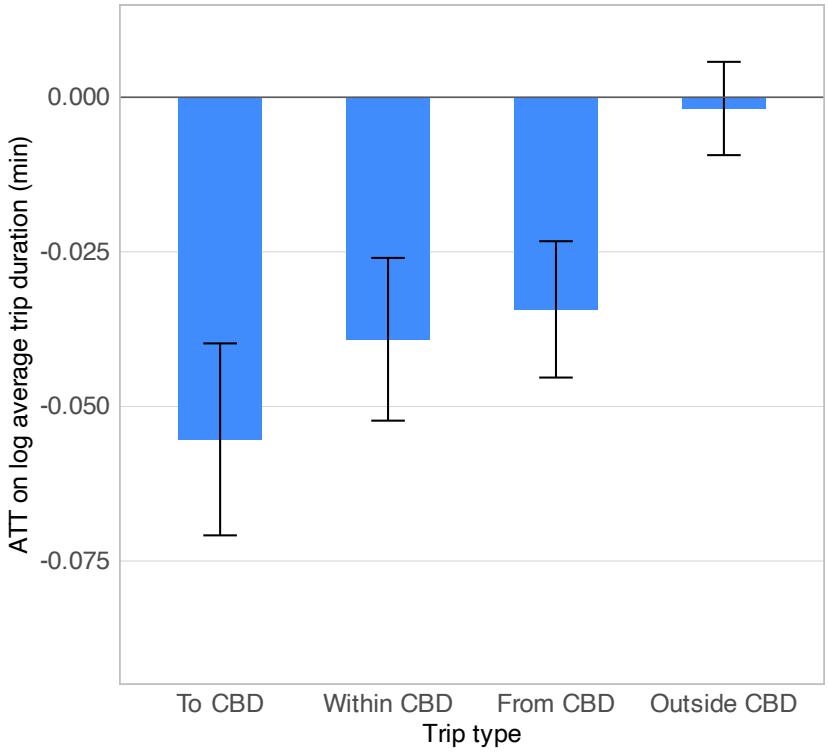


(f) 80-100%



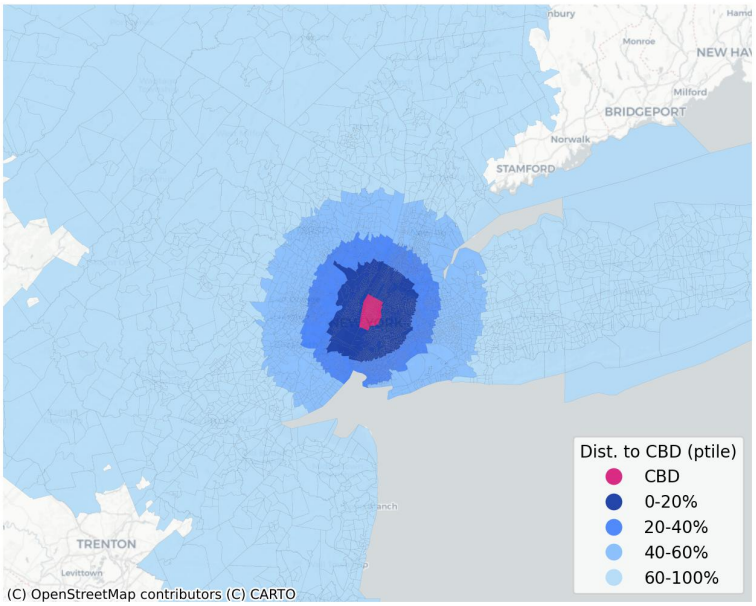
Notes: These figures map the road segments by their level of co-occurrence with the CBD, which is defined as the share of trips crossing the road segment that ultimately enter the CBD. The first panel plots the distribution of road types by co-occurrence.

Figure C.5: Treatment effect on trip durations



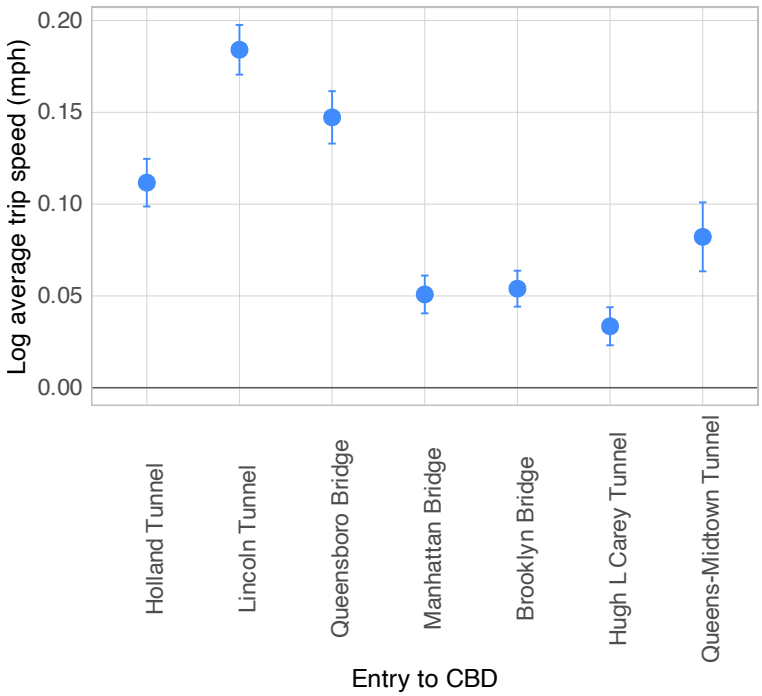
Notes: This figure documents day-level ATT on average realized trip durations split by whether the trip starts and/or ends in the CBD. The underlying data are the aggregate speed for each city, OD type, and two-hour bin. The horizontal blue lines are the aggregate ATTs for all post-treatment periods. Shaded areas represent 95% confidence intervals. Standard errors are clustered at the city-level.

Figure C.6: Bins of distance to CBD



Notes: The boundaries are based on the Core Based Statistical Area boundaries. The map is zoomed in slightly to make the CBD more visible; the CBSA extends further in each direction, especially east down Long Island. Distance to CBD is based on the distance between centroid and the closest point of the CBD.

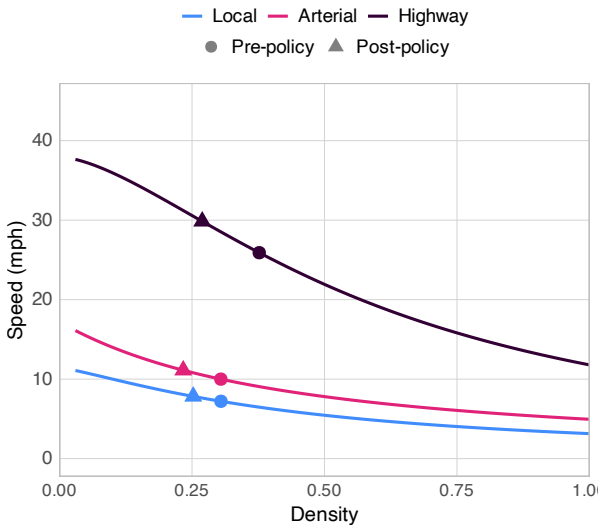
Figure C.7: ATTs on trip speeds by CBD entrance



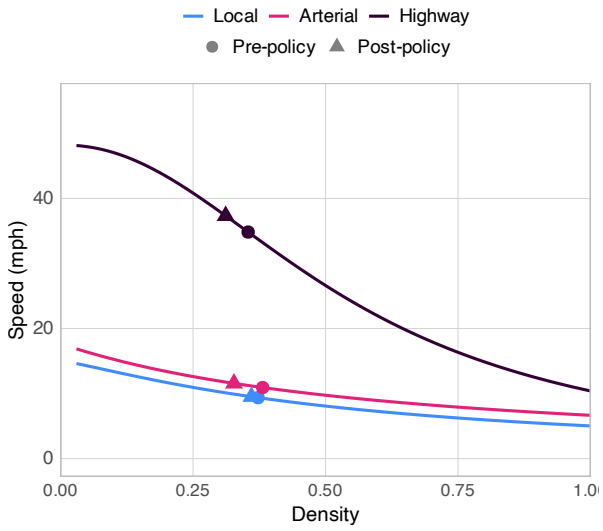
Notes: This figure plots the estimated ATT on log trip speeds for trips entering the CBD, split by which entrance they cross to enter the CBD. The set of potential controls include two-hour aggregates of trip speeds on trips to the control city CBDs from any origin. Vertical lines represent 95% confidence intervals.

Figure C.8: Other congestion functions: NYC

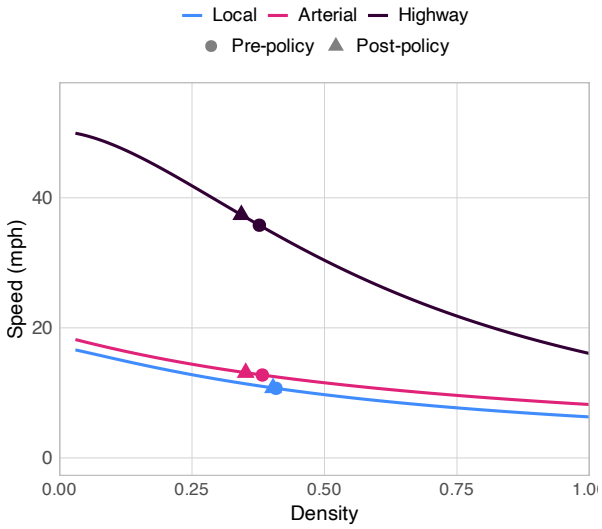
(a) 100-80% co-occurrence roads



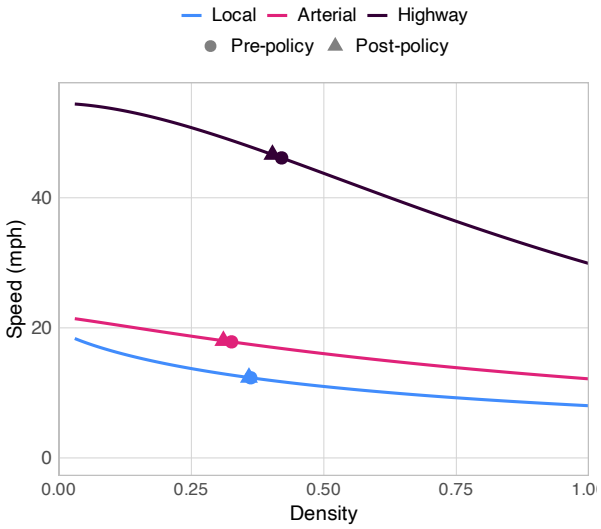
(b) 60-40% co-occurrence roads



(c) 40-20% co-occurrence roads



(d) 20-0% co-occurrence roads



Notes: This figure plots the remaining estimated congestion functions for roads in NYC; see Figure 8 for the CBD and 80-60% co-occurrence congestion functions. The congestion functions are estimated using pre-period observed speeds and densities, based on Equation (17). Circles are added to each line based on the raw average speed before the policy's implementation. Triangles are based on the post-period speed, which is computed as the pre-period speed plus the estimated ATT on speeds.