The Network-Wide Effects of Congestion Pricing: Evidence from New York City*

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

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

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

August 3, 2026

Abstract

Congestion pricing programs toll only a subset of roads and times, with theoretically ambiguous spillovers. We study New York City's cordon-based congestion pricing using granular data from Google Maps and a generalized synthetic control design. The toll raised cordon-area speeds by 11%, with little effect on air quality, consumer spending, or foot traffic. Speeds also rose outside the cordon, especially on roads frequently traversed by cordon-bound drivers, making many *unpriced* trips faster. The magnitudes reflect steep congestion functions, where small volume reductions yield large speed gains. Because unpriced trips vastly outnumber priced ones, their small per-trip gains dominate aggregate driver welfare.

JEL codes: **R41, R48, R5, H2, H71**

Keywords: congestion pricing, traffic, externalities, Pigouvian taxation

*Cook, Kreidieh and Vasserman are joint first authors. Shoshana Vasserman is the corresponding author (svass@stanford.edu). We are grateful to Milena Almagro, Costas Arkolakis, Felipe Barbieri, Lanier Benkard, Steve Berry, Victor Couture, Rebecca Diamond, Matthew Freedman, Ed Glaeser, Stephan Heblich, Drew Hill, Emir Kamenica, Zi Yang Kang, Gabriel Kreindler, Ivan Kuznetsov, Pearl Li, John List, Preston McAfee, Willa Ng, Carolina Osorio, Michael Ostrovsky, Jim Poterba, Prem Ramaswami, Stephen Redding, Katja Seim, Andrew Stober, and Matt Yarri for their many helpful comments. This paper subsumes an earlier draft titled "The Short-Run Effects of Congestion Pricing in New York City."

1 Introduction

Congestion pricing is a textbook application of Pigouvian taxation (Pigou, 1920), and, in theory, can achieve first-best traffic allocations by internalizing congestion externalities (Vickrey, 1963, 1969). In practice, existing congestion pricing programs are far from the theoretical first-best: they price only a subset of roads and times, often at a flat rate. Because pricing covers only part of the network, the welfare effects depend on how traffic reallocates across both priced and unpriced roads. The direction of these spillovers is theoretically ambiguous. Pricing a subset of roads may divert traffic to unpriced routes and worsen congestion there, or reduce total vehicle volume and ease congestion city-wide.[1] Because unpriced roads and trips vastly outnumber priced ones, effects that are individually small can still be large in sum. Empirical evidence is nonetheless scarce, in part because detecting these spillovers requires isolating small, diffuse changes across the whole road network.

In this paper, we study the network-wide effects of congestion pricing in the context of New York City's (NYC) cordon-based program, the first of its kind in the United States. The policy launched in January 2025 following decades of debate over its likely effects and over which residents would bear the costs.[2] Under the policy, passenger vehicles must pay $9 to enter NYC's central business district (CBD) between 5am–9pm on weekdays and 9am–9pm on weekends.[3] Our primary data include road-segment- and origin-destination-level traffic statistics derived from anonymized trips taken using Google Maps. The granularity and scale of the data allow us to detect even small effects—often less than one percent—on speeds throughout the metro area. We complement our traffic data with information on air quality from PurpleAir sensors, spending at shops and restaurants from a sample of credit/debit cards, and foot traffic from a sample of GPS devices. To estimate the policy's effects, we use a generalized synthetic control (GSC) design and other large US cities as controls (Xu, 2017). We then build a model to map the estimated changes in prices and travel times into bounds on driver welfare, decomposing the policy's total effect on drivers into the direct effects on priced trips into the CBD and the spillover effects on unpriced trips throughout the NYC metropolitan area.

We start with the direct effects inside the cordon, where we show that the introduction of

[1] In a survey of economists before the policy's implementation, 90% 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 priced zone (Clark Center Forum, 2024).

[2] Opponents argued that a Manhattan cordon would shift traffic to outer-borough roads rather than reduce driving overall (Schaller, 2010; Ley, 2022). William Vickrey first proposed pricing for Manhattan in the 1950s. Subsequent plans were blocked by Congress in 1977, withdrawn in 1987, and failed in the state Assembly in 2008. The current program was authorized in 2019, paused indefinitely in June 2024, and revived in November 2024 at $9 rather than the $15 approved by the MTA board.

[3] The priced zone covers lower Manhattan south of 60th Street, excluding the FDR Drive and West Side Highway. Prices are per-day, and are lower for motorcycles, higher for trucks and buses, and 75% lower during off-peak hours. For-hire vehicles pay a per-trip fee of $0.75 for taxis and $1.50 for ridesharing companies.

congestion pricing increased road speeds in the CBD and had little-to-no effect on air quality, consumer spending, or foot traffic. Relative to a synthetic control formed from other cities, average speeds on CBD road segments in NYC increase by 11% after the launch of congestion pricing. The speed gains attenuate from 15% in the first four months to 10% thereafter, consistent with marginal drivers returning to the road as they learn that driving has become faster.[4] Unlike for road capacity expansions, where "induced demand" can fully erode speed gains (Downs, 1962; Duranton and Turner, 2011), congestion pricing puts a permanent wedge in the cost of priced trips, and speeds settle above their pre-policy level. The effects on CBD speeds are largest during the afternoon—historically the most congested time—and persist even after peak-hour pricing ends. Introducing congestion pricing also reduced the dispersion in speeds by compressing the tail of extreme delays and improved estimated fuel efficiency.[5] Beyond traffic conditions, however, we find no detectable treatment effect on ambient concentrations of fine particulates ($PM_{2.5}$), spending at shops and restaurants, or overall foot traffic in the CBD.

To evaluate the spillover effects outside the CBD, we first develop a road segment-level measure of exposure to the policy based on each segment's *co-occurrence* with the CBD. We define co-occurrence as the share of drivers crossing the focal segment who also cross CBD segments as part of their trip. Many drivers not bound for the CBD still traverse segments exposed to the policy, and, through this exposure to the now-priced CBD trips, can experience speed changes after the policy's implementation. For example, if the policy's main effect is to reduce CBD trips by 10% uniformly across space, a segment with 50% co-occurrence—i.e. a segment where half of the cars are headed to/from the CBD—would see a proportional 5% decrease in cars. For each segment, we compute its co-occurrence using pre-period traversals so that the measure is unaffected by any post-policy rerouting.

We find that the introduction of congestion pricing increased speeds on roads throughout the metro area, with little evidence of any offsetting slowdowns. The increases in speeds scale with each segment's co-occurrence with the CBD. For road segments with 80-100% co-occurrence—which include the Lincoln Tunnel, Holland Tunnel, and other main roads leading into the CBD—speeds increased by an average of 6.4%. Because these high-exposure segments start from higher baseline speeds, the effects on absolute speed *levels* are larger than in the CBD itself. Across road types, highways generally experienced larger increases in speeds than local and arterial roads. For roads with lower levels of co-occurrence, we find correspondingly smaller effects on speeds. However, even for those with just 0-20% co-occurrence we find modest increases in average speeds of 1–1.5% on highway and arterial segments.

[4]Consistent with such learning, Dreyfuss (2025) and Yang (2025) find that voter sentiment toward congestion pricing improves after launch, which Dreyfuss (2025) attributes to people's difficulty anticipating the equilibrium effects of the policy.

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

The changes in road segment speeds add up to aggregate travel-time savings for trips throughout the city, including for many *unpriced* trips. For trips to the CBD, average trip speeds increase by 3.7% following the launch of congestion pricing. While drivers on trips to the CBD pay for these time savings, passenger vehicles on other trips within, from, or outside the CBD benefit from increased speeds without paying the congestion fee. The policy increased speeds on trips from and within the CBD by 6.2% and 4%, respectively, and even trips entirely outside the CBD experienced modest gains of 0.4–1.9% depending on the distance of their origin and destination to the CBD. The speed gains are also broadly shared, both across the neighborhood income distribution and in many outlying areas. A common concern with congestion pricing is whether the costs fall most heavily on particular places or on lower-income drivers (Ecola and Light, 2009; Taylor, 2010; Hu and Ley, 2024). At the segment-level, average speeds rose in Census tracts across the household income distribution, with larger increases in lower-income tracts. Similarly, segments in the Bronx, Hudson and Bergen counties in New Jersey, and Long Island experienced net increases in speed. The patterns are similar at the trip-level, with speeds on trips to and from the CBD rising by 4–8% for each income quintile and for the Bronx, New Jersey counties, and Long Island.

To evaluate the implications for driver welfare, we build a model in which drivers weigh time savings against price changes. Because our data come from trips routed through Google Maps rather than a census of all cars, they are ill-suited for measuring changes in volumes.[6] However, we show that the model can provide a lower bound for the welfare change using only pre-period trip volumes and the observed changes in prices and travel times, similar to the "social savings" approach in Fogel (1964). By revealed preference, drivers who stop taking a trip lose no more than those who continue, so the bound accounts for forgone trips without observing them.

For trips traveling to the CBD, the average passenger vehicle driver pays $7.9 to save 2.3 minutes. At this price, a driver headed to the CBD is only made better off by the policy if their Value of Travel Time (VOTT) exceeds $207/hour.[7] The average trip from, within, and outside the CBD saves just 9 seconds, but drivers do not pay for the speed gains. Moreover, these unpriced passenger vehicle trips outnumber priced trips by more than a hundredfold. Accounting for both priced and unpriced passenger vehicle trips, the break-even VOTT is at most $25/hour. This pattern of concentrated losses and diffuse gains is commonly associated with policies that are politically fragile even when efficient (Olson, 1965; Wilson, 1980), which may explain why congestion pricing remains rare. Taxi and for-hire vehicle trips enjoy similar time savings at a lower per-trip fee, and even their priced trips come close to breaking even for reasonable

[6]The observed counts reflect Google Maps usage, which may itself respond to changing traffic conditions. Speeds and travel times are less affected by changes in the sample, as long as marginal drivers travel at speeds similar to those of inframarginal drivers.

[7]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.

VOTTs. Assuming all drivers and taxi/FHV passengers have a VOTT of $40/hour—about 80% of the average hourly wage in the NYC metro area[8]—the policy generates weekly driver welfare gains of at least $7.5 million, before accounting for any revenue recycling, environmental benefits, or the benefits for other road users (e.g., buses and emergency vehicles). The weekly revenue net of operating expenses is roughly $11.4 million and is earmarked for investments in public transit (MTA, 2026*c*). The "break-even" VOTT for the policy is $26/hour before accounting for revenue and just $5/hour if the net revenue is fully redistributed.

Finally, we investigate the mechanisms behind the effects and assess whether similar results are plausible in other cities. The network-wide effects of cordon pricing depend, in part, on two factors measurable in our data: the exposure of other drivers to changes in CBD traffic and the congestion functions that relate traffic density to speed. We measure *exposure* as the duration-weighted co-occurrence of segments traversed on a trip. For a driver sampled at a random moment on their trip, this measure captures the expected share of surrounding drivers who are on CBD trips. We estimate *congestion functions* by road type and co-occurrence bin using contemporaneous speeds and vehicle densities, following the methodology in Choudhury et al. (2024) to impute densities from observed partial flows. Consistent with the traffic engineering literature (see, e.g., Seo et al., 2017), the estimated functions are non-linear, so the effect of removing traffic depends on the local elasticity of each congestion function.

Of these two factors, the steepness of the congestion functions in NYC most plausibly accounts for the positive spillovers on trip speeds throughout the city. NYC's roads operated on unusually elastic regions of their congestion functions, where even modest volume reductions translate into large speed gains.[9] Average local elasticities for NYC trips are over 80% higher than for comparable trips in Philadelphia or Baltimore and 20% higher than for trips in Boston or Atlanta. Only Chicago is higher, with average local elasticities about 25% above those in NYC. While NYC's exposure on trips to and from the CBD is also among the highest (second only to Chicago), its exposure on trips entirely outside the CBD is the lowest in our sample. Across cities, cordon pricing may be least effective in Philadelphia and Baltimore (both lower exposure and lower elasticities than NYC) and most effective in Chicago (higher exposure and higher elasticities). However, the realized effects in any city will also depend on factors we do not measure, such as the price elasticity of driving demand. This demand elasticity may vary with the reach of public transit in each city, the ease of substituting away from car trips, and the distribution of drivers' price sensitivity.

[8] Goldszmidt et al. (2020) estimate a value of time of 75% of the average hourly wage using experimental variation in wait times for ride-sharing. For drivers in Prague, Buchholz et al. (2025) estimate a value of time that is greater than the average hourly wage.

[9] Kreindler (2024) finds that congestion functions in Bangalore are relatively flat and linear, which mutes potential speed gains from "peak spreading" there. The differences across cities may stem from variations in infrastructure investments, which Akbar et al. (2023) show are a stronger determinant of cross-city road speeds than congestion levels.

Related literature. The spillovers of congestion pricing have often been recognized in theory, but rarely measured empirically. Early theoretical work emphasized the case in which diversion worsens congestion on unpriced substitutes (Verhoef, Nijkamp and Rietveld, 1996; De Palma and Lindsey, 2000).[10] However, pricing can also reduce the number of cars on the road rather than just relocate them. Recent quantitative spatial frameworks make the network-wide propagation explicit, showing that a change on one link alters travel costs and traffic throughout the network, but the sign and magnitude of these effects remain an empirical question (Fajgelbaum and Schaal, 2020; Allen and Arkolakis, 2022).[11] Using data for a subset of major road segments, Herzog (2024) finds evidence of reduced traffic on highways leading towards London's downtown cordon zone, while Gibson and Carnovale (2015) find that Milan's policy increased traffic on radial roads skirting the congestion zone. We offer three main contributions to this literature. First, using granular data on *all* road segments, we quantify the net effect of spillovers and show that the policy improved average speeds on trips throughout the metro area. Second, we trace the magnitude of the spillovers to the slopes of the relevant congestion functions. NYC roads were operating at steep points on their congestion functions, where each additional driver imposes large travel time externalities on other drivers. Finally, we show that accurately measuring the spillovers is first-order for welfare accounting. Imposing a corrective tax (Diamond, 1973; Sandmo, 1975) on CBD-bound trips yields positive welfare gains even before using the revenue, driven largely by time savings for the many unpriced trips outside the cordon zone.

There is more existing evidence on the *direct* effects of cordon-based programs on outcomes within the priced zone. The 2003 London congestion charge reduced traffic, pollution, and accidents within the zone (Leape, 2006; Tonne et al., 2008; Green, Heywood and Paniagua, 2020; Tang and van Ommeren, 2022). Similar effects have been documented for Stockholm (Eliasson et al., 2009; Börjesson et al., 2012; Simeonova et al., 2021; Nilsson, Tarduno and Tebbe, 2024) and Milan (Gibson and Carnovale, 2015), and journalistic accounts of the NYC policy reached similar conclusions about within-cordon speeds (Gordon et al., 2025; Ley, Hu and Collins, 2025; Hu, Ley and Schweber, 2025). Most prior programs predate Google Maps and other large-scale data sources, so researchers typically relied on a handful of sensors from the focal city and studied outcomes on roads within the cordon, using "untreated" roads elsewhere in the same city as controls. However, when roads elsewhere in the city are themselves affected,

[10] 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, 2024; Cook and Li, 2024).

[11] Other place-based policies also generate spillovers beyond the targeted area. Because the spillovers are difficult to detect empirically, evaluations typically focus on the targeted zone or impose structural assumptions to recover their magnitude (Busso, Gregory and Kline, 2013; Kline and Moretti, 2014; Criscuolo et al., 2019; Adao, Arkolakis and Esposito, 2025).

within-city comparisons are biased in a direction that depends on the sign of the spillovers (Butts, 2021). In NYC, where the spillovers are positive, such a design would understate the within-cordon gains. In contrast, our data span the near-universe of segments across many cities, allowing us to use untreated cities as controls. Moreover, these richer data allow us to estimate treatment effect heterogeneity by time of day, road type, and for specific geographies (e.g., higher- versus lower-income tracts), which has not been documented for prior programs.

Finally, a broader literature studies the overall costs of urban congestion (see, e.g., Couture, Duranton and Turner, 2018; Akbar et al., 2023) and the policies intended to reduce it. Aside from pricing, common policy instruments for reducing congestion include road capacity expansions, which induced demand tends to undo (Downs, 1962; Duranton and Turner, 2011), driving restrictions and car-free zones, which are typically poorly targeted and raise no revenue (Davis, 2008; Gallego, Montero and Salas, 2013; Sleiman, 2024), and transit improvements (Anderson, 2014; Gu et al., 2021). Within pricing, experimental and quasi-experimental evidence shows that drivers respond to time-varying charges, though the resulting travel-time gains depend on the underlying congestion functions (Kreindler, 2024; Cook and Li, 2024; Ater et al., 2026). A complementary structural literature simulates optimal pricing schemes for cities without existing programs (Barwick et al., 2024; Durrmeyer and Martínez, 2024; Hierons, 2024; Almagro et al., 2025). We contribute causal evidence from the first cordon-based congestion policy in the United States, and show that even a coarse congestion pricing policy can increase driver welfare and improve traffic conditions across the metropolitan area.

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 by time of day, vehicle type, and whether the vehicle is equipped with an E-ZPass transponder. The Metropolitan Transportation Authority (MTA) plans to use the toll revenue 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 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.[12] There are a few exempted vehicles (e.g., emergency

[12]The charges are shown as separate charges on receipts, although it is unclear whether base fares for FHV

vehicles), and vehicles entering via certain bridges or tunnels that are already tolled receive a partial credit. Low-income residents can also apply for 50% off after their first 10 trips in a month. The passenger car rate is set to increase to $12 in 2028 and $15 in 2031.

Whether the program would take effect, and at what price, remained unresolved until weeks before launch. The state legislature authorized the program in the April 2019 budget, initially targeting a 2021 launch, but federal environmental review and litigation delayed it for years. Critics argued that congestion pricing in NYC would impose an additional financial burden on residents and shift traffic and pollution to other parts of the city (Ley, 2022). Nonetheless, the MTA board voted in March 2024 to adopt a $15 peak toll and set a June 30, 2024 start date. However, weeks before the intended launch date, the state announced an indefinite pause, citing costs to drivers. The program was revived in early November 2024 with a $9 toll—40% below the approved rate—and the final agreement between the MTA and the federal government was signed on November 21st, 2024. Tolling began on January 5th, 2025, following unsuccessful last-minute attempts to block it in court. Because both the existence and the level of the toll were in doubt until shortly before implementation, drivers had little basis for adjusting their behavior in advance, limiting anticipation effects.

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. In Appendix B.1, we show that our main results are robust to this choice of cities. For each city, we define its boundaries according to the corresponding Core Based Statistical Area (CBSA) and define its CBD using versions of a CBD or "downtown" drawn by a city government-affiliated organization.[13] Appendix A contains additional details on each data source.

3.1 Google Maps

Our primary data on traffic conditions are anonymized, aggregated statistics from trips taken using Google Maps during the study period. The data include two primary sets of statistics: 1) hourly road segment-level outcomes, which we then further aggregate in time and space across sets of road segments; and 2) hourly origin-destination (OD) level outcomes, which are

operators adjust to absorb some of the charge. Ostrovsky and Yang (2025) 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.

[13] In NYC, the priced 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.

aggregated across trips based on OD characteristics. The sample covers traffic conditions from September 2024 through June 2026, giving us approximately 4 months of data prior to the policy's launch and 18 months since launch.[14] Except where otherwise noted, all analyses focus on data during priced hours (5am–9pm on weekdays and 9am–9pm on weekends).

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. 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 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 estimate the liters of gasoline per 100 kilometers consumed by a typical sedan on a road segment, based on the segment characteristics and speed profile, following the approach in Brooker et al. (2015). Appendix A.2 provides additional details on how we model fuel consumption.

Origin-Destination (OD) Level Outcomes. Our second set of data covers 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. 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 pair, we measure realized travel times and trip speeds (e.g. trip length divided by travel time), aggregated across all trips within a given OD-hour bin. 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 OD-hour bin between September and December 2024.

[14] These data were not retained before the project began, so earlier months of data are unavailable.

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. 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). 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 use data from January 2024 through June 2026. 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. The data cover transactions between January 2024 and July 2025, and do not extend through early 2026 (unlike our other data sources). We focus on transactions at restaurants or retail establishments, which we identify using their NAICS codes.[15] In NYC, the data include 130 million restaurant transactions ($3.6 billion total spend) and 302 million retail transactions ($14.1 billion total spend) during the sample period (Table A.2). The boundaries of NYC zipcodes almost exactly align with the CBD boundaries, so we define a zipcode as treated by the policy if its centroid lies within the NYC CBD.

Foot Traffic. We use Advan Research's Neighborhood Patterns to measure the amount of foot traffic in the CBD over time. 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 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. We use data from January 2024 through March 2026,[16] and define a tract as treated by the policy if its centroid lies within the NYC CBD.

[15] We use the 3-digit NAICS code 722 to identify restaurants and the 2-digit NAICS codes 44-45 to identify retail establishments.

[16] The data are released with a lag and were not yet available for more recent months at the time of writing.

3.3 Empirical Strategy: Generalized Synthetic Controls (GSC)

A direct comparison of average outcomes before and after the policy combines the policy's effects with any other time-varying factors influencing traffic conditions. As such, we adopt the generalized synthetic control (GSC) methodology introduced by Xu (2017), which compares changes in traffic conditions in NYC to contemporaneous changes in comparison cities.

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 i is treated. 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.

For each outcome, we estimate these factors and loadings using pre-treatment data. We 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 $t > T_0$. The Average Treatment on the Treated (ATT) in period t is defined as:

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

where $\mathcal{I}$ is the set of treated units. We often report the aggregate ATT, given by $\widehat{ATT} = \frac{1}{T - T_0} \sum_{t > T_0} \widehat{\mathrm{ATT}_t}$, where T_0 is the last day before pricing begins. We implement this procedure using the *gsynth* package for R, provided by Xu (2017). We use cross-validation to select the factor dimension and report parametric standard errors for inference.

A key advantage of this design is that it uses other cities as controls rather than other roads within NYC. Most prior evaluations of cordon pricing compare treated roads inside the zone to untreated roads elsewhere in the same city. However, the spillovers we document extend throughout the NYC metro area, so roads outside the cordon are themselves affected by the policy. Using them as controls would violate the stable unit treatment value assumption (SUTVA), which requires that a unit's outcome be unaffected by others' treatments (Rubin, 1980). Such interference biases within-city comparisons toward zero since the "control" roads also speed up, mirroring identification challenges found in other interventions with significant positive externalities (Miguel and Kremer, 2004). By drawing controls from separate metro areas, our design sidesteps this interference.

4 Direct Effects in the CBD

According to the MTA, vehicle entries into the CBD decreased by about 10% following the introduction of the congestion pricing policy and remained below the expected volumes throughout 2025 (MTA, 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 11% and improved estimated fuel efficiency by 2.5%. Beyond traffic conditions, however, we find no evidence of effects on ambient air quality, visits to restaurants and retail establishments, or foot traffic.

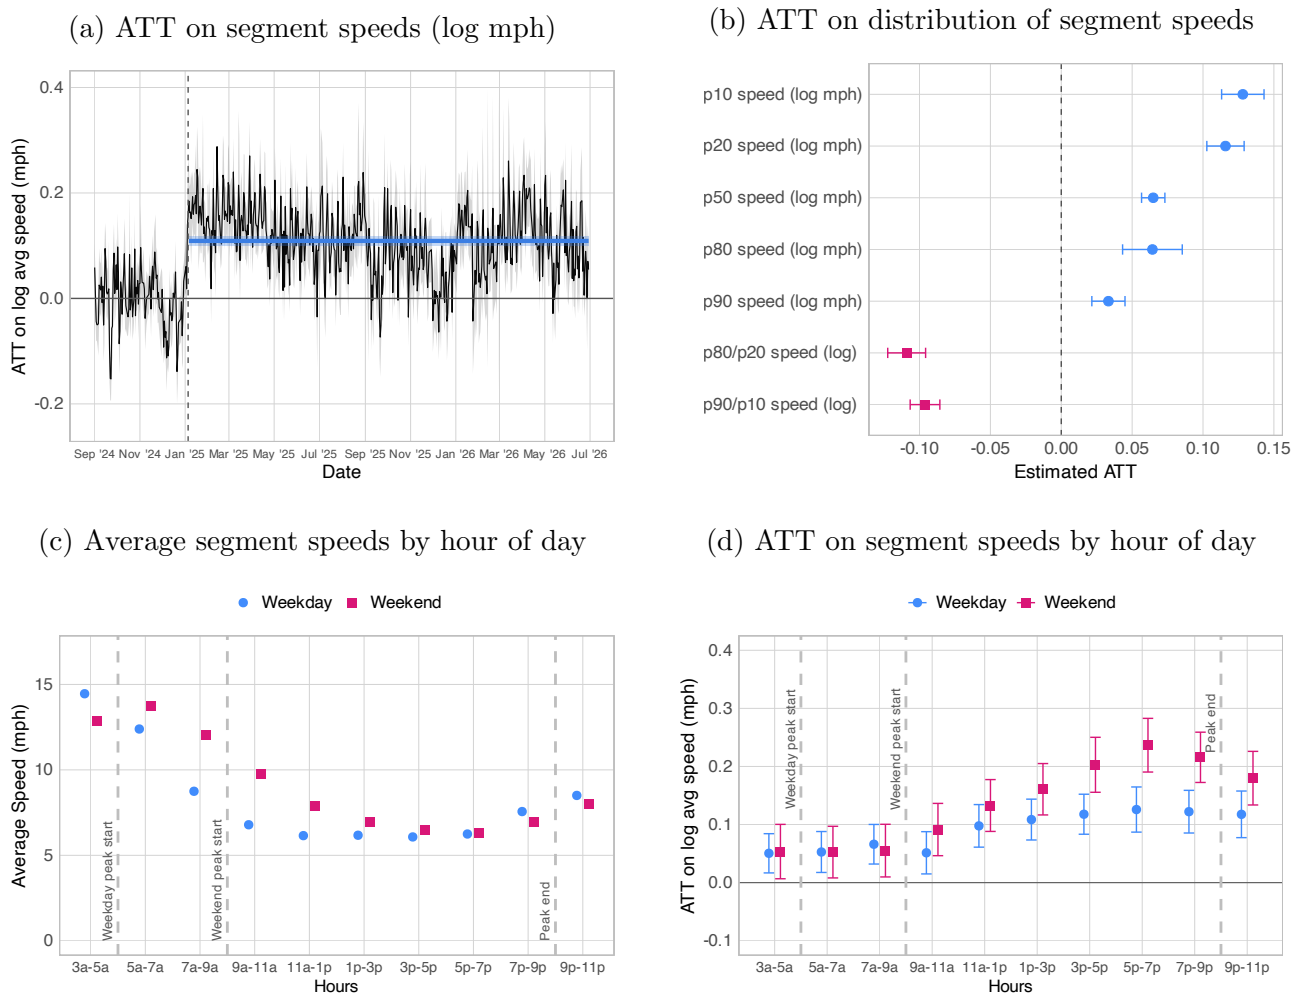
4.1 Road Speeds within the CBD

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 7.7 mph in the eighteen months after implementation (Table C.1). To evaluate the extent to which the observed increase in speeds is attributable to the policy, we estimate Equation (3) using a panel of log average speeds in the CBDs of NYC and our control cities. Figure 1a plots the day-level ATT for CBD road segments during priced hours.[17] Before the policy's launch, speeds in NYC were similar to the counterfactual speeds constructed from the set of comparison cities. Immediately after the onset of congestion pricing, average speeds in the CBD increased sharply relative to the synthetic control. Averaging over the entire post-period, the implementation of congestion pricing increased speeds by 11%. This suggests that, under the assumptions of the GSC estimator, the policy's causal impact can fully account for the change in raw average speeds on CBD segments.

These average effects aggregate over both extreme stop-and-go traffic and near-free-flow traffic. Most roads in the CBD are arterial roads (e.g., 5th Avenue) or substantially narrower local roads, both of which are slower (and have lower speed limits) than the handful of high-volume highway segments included in the CBD area. The average speed on arterial roads before the policy was 7.2 mph, while on local roads it was only 5.5 mph (Table C.1). The 10th percentile of pre-policy speeds was below 3 mph on both local and arterial roads, and the 90th percentile was about 20 mph on local roads and 25 mph on arterials. The introduction of congestion pricing disproportionately impacted the slowest traffic conditions. The 20th percentile of road speeds experienced across the CBD increased by 12% on average, while the 80th percentile of road speeds increased by 6.4% on average. The average dispersion of road speeds in the CBD, as captured by the ratio of 80th to 20th percentile speeds, decreased by 11% (Figure 1b). Therefore, one effect 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).

[17] Figure C.2 includes similar plots for other measures of CBD traffic conditions.

Figure 1: Effects on Speeds in the CBD



Notes: Panel a) documents day-level ATTs of congestion pricing over time on peak-hour CBD segment speeds. The data are average segment speeds by two-hour bin in each CBD. The horizontal blue line is the aggregate ATT for all post-treatment periods. Panel b) documents the ATTs on the distribution of speeds in the CBD. Panel c) documents the average pre-period speeds by hour bin and weekday vs weekend. Panel d) 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. Shading and vertical bars denote 95% confidence intervals. Standard errors are clustered at the city-level.

The average effects also 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 less than half that in the late afternoon (Figure 1c). In Figure 1d, we present ATTs on log CBD road speeds for each two-hour time window between 3am and 11pm. The increases in average road speeds are modest in the early morning (about 5%), but rise to about 15% at the weekday afternoon peak and over 25% in the early evening on weekends. That is, the relative impact of congestion pricing on speeds was largest during the times of day when congestion was most severe before the policy. Moreover, increases in road speeds persist after peak-hour pricing ends at 9pm, suggesting that reduced entry volumes during peak hours had spillover effects on road conditions later in the evening.

The effects of the policy on CBD speeds were largest during the first four months. The average treatment effect from launch until the start of May 2025 was 15%, then fell to 10% and remained largely stable through the end of our sample. This partial attenuation is consistent with demand re-equilibrating to the faster speeds the policy itself created. Individuals have difficulty forecasting the equilibrium effects of policy changes (Dreyfuss, 2025), so beliefs formed before the policy's launch may have underappreciated the potential for improved travel times. Consistent with this possibility, Yang (2025) shows that public opinion of congestion pricing in both London and NYC was initially negative but improved after each policy launched. As NYC travelers observed that trips had become faster, marginal drivers may have returned to the road, dissipating part of the initial speed gains. The dynamics echo the "induced-demand" response to highway capacity expansions, in which latent demand erodes initial speed improvements (Downs, 1962; Duranton and Turner, 2011). A congestion charge differs in a critical way, however. Added capacity lowers the cost of driving until demand absorbs the gains, whereas a toll introduces a permanent wedge in the cost of priced trips. So long as priced demand responds to the toll, equilibrium speeds must settle above their pre-policy level rather than fully rebounding, which is the pattern we observe.

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

4.2 Other CBD Outcomes

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 local air pollutants such as NOx, CO, and particulates impose externalities on nearby residents (Buckeridge et al., 2002; Currie and Walker, 2011; Knittel, Miller and Sanders, 2016). While we cannot directly observe vehicle emissions, we use predicted fuel consumption rates as a proxy for emissions. For a given vehicle type, the primary determinants of fuel consumption are vehicle speed profiles and the segment type. Cars traveling at 5 mph—close to the average CBD speed in late afternoon—will consume about four times as much fuel per mile as those traveling at 60 mph (Figure A.1). Compared with the control cities, estimated fuel consumption rates on NYC CBD segments decreased by 2.5% after the launch of congestion pricing (Figure 2).

Despite the more efficient speed profiles and reduced vehicle traffic in the CBD, we find no evidence of effects on air quality. The estimated treatment effect on $PM_{2.5}$ inside the NYC CBD is statistically indistinguishable from zero (Figure 2), and 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 -0.36$\mu g/m^3$, which would correspond to a 4.8% decrease relative to the pre-period daily average. One plausible explanation is that the London and Stockholm programs launched when emissions standards were less stringent, so the marginal vehicles deterred by congestion pricing in NYC may differ from those deterred in London and Stockholm. For example, the Euro 3 emissions standards (2000–2005) allow for over eight times more NOx per kilometer than the Tier 3 US emissions standards (2017–) allow for NOx and NMOG combined (0.15 g/km vs. 0.018 g/km). The results for London and Stockholm also consider a longer time horizon and may capture changes such as a shift towards green vehicles, which are exempt from the Stockholm fee (Nilsson, Tarduno and Tebbe, 2024) but face the same price as gas-powered vehicles in NYC.[18]

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 the same synthetic controls approach to estimate treatment effects on transactions at restaurants and retail establishments, as well as on overall foot traffic.

First, using aggregated credit/debit card spending data, we compare the number of transactions and the total amount spent at restaurant and retail establishments in the NYC CBD with 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 between peak and off-peak hours.

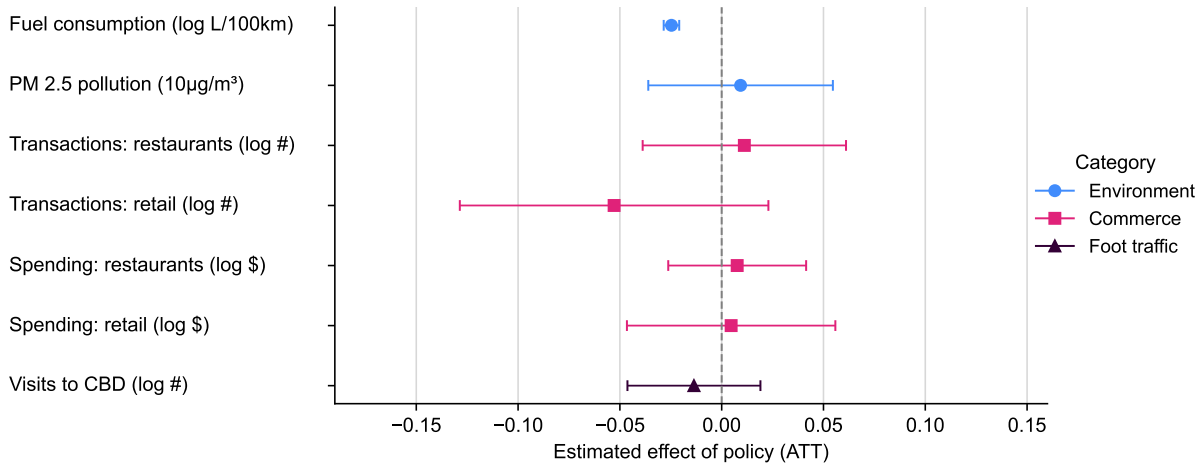
We find no evidence that congestion pricing reduced the number of transactions or total spending at either restaurants or retail establishments in the NYC CBD (Figure 2). For total spending, the point estimates are close to zero, and the 95% confidence intervals rule out decreases of over 4.7% at retail establishments and 2.6% at restaurants. Similarly, the estimates for effects on the total number of transactions are not statistically different from zero, although the results are noisier and we cannot rule out large effects, especially for the number of transactions at retail establishments.

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

[18]One additional concern is that the low-cost optical laser counters used by PurpleAir sensors are less reliable for detecting ultrafine particles (see, e.g., He, Kuerbanjiang and Dhaniyala, 2019). However, while fresh tailpipe emissions are concentrated in the ultrafine range (<0.1 μm), a substantial portion of vehicle-induced particulate mass is generated above this threshold via non-tailpipe pollution sources, including brake pad abrasion and tire-wear friction.

been offset by substitution toward other modes. 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 adjusted along the intensive margin of mode choice rather than the extensive margin of deciding whether to visit the CBD at all. However, we caution that the confidence intervals for consumer spending and foot traffic are wide enough to admit meaningful changes (Figure 2).

Figure 2: 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 1a to compute the ATT. 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 by their NAICS codes. 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. Horizontal lines denote 95% confidence intervals. Standard errors are clustered by city.

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 (e.g., Leape, 2006; Eliasson et al., 2009), where the effect sizes are often large enough to detect even with more limited data. In this section, we show that congestion pricing in NYC also increased speeds on segments outside 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 policy. The intuition behind our measure is simple: segments that are frequently traversed as part of CBD trips are more exposed to reductions in CBD 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:

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

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 September to November 2024, and hold these values constant throughout our analysis. Figure C.3 plots the spatial distribution of road segments in NYC by their level of co-occurrence. Roads closer to the CBD and highways leading toward the CBD tend to exhibit higher levels 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 and decreases in entries were proportional to pre-policy volumes, then co-occurrence would exactly predict the policy's impact on volumes. 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.

Figure 3 plots average ATTs of congestion pricing, split by co-occurrence bin and road type (i.e., arterial, highway, and local roads, or a traversals-weighted average of all three). In each case, we estimate the treatment effects using the approach outlined in Section 3.3, in which we construct a synthetic control by combining segments from other cities that share the same co-occurrence bin and road type. 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 going to or from the CBD prior to the policy—experienced a 6.4% increase in speeds. The effects then decrease with co-occurrence. In the lowest co-occurrence bin, which contains the majority of road segments outside the CBD, speeds increased by 1% on highways and 1.5% on arterial roads.

The effects on high co-occurrence roads are large and, when measured in levels, often exceed the direct effects on the CBD. For example, the ATT for the highways in the highest co-occurrence bin, which include many of the major entrances to the CBD (e.g., the Lincoln and Holland Tunnels), is 7.2%. As the average pre-policy speed on highways with 80-100% co-occurrence was 23 mph, this effect translates to about 1.8 mph faster average speeds on the highway segments themselves. By contrast, the implied average increase in road speeds within the CBD, with a much lower baseline (7.1 mph), is about 0.8 mph.

At all levels of co-occurrence, the speed gains tend to be larger for highways and arterial roads than for local roads. Even highways and arterial roads with 0-20% co-occurrence with the CBD experienced an average increase of 1% to 1.5% in speeds. Among local roads, those with 80-100% co-occurrence, such as the streets just above 60th St. at the CBD's northern edge, also show some speed gains. Further out, local roads are generally less affected, and we see no speed increases for roads beyond the 60-80% co-occurrence bin. Some local roads between 20-60% co-occurrence may even be slightly slower now, although the effects are small and there are relatively few local roads in these bins. As we discuss in Section 7, highways in NYC operated at densities where removing a few cars raises speeds far more than on local or arterial roads, which helps explain why the gains are concentrated there.

Overall, we find that traffic conditions improved broadly outside the CBD, with the largest gains on the main approaches into and out of the zone and smaller but measurable gains across much of the network. Averaging by co-occurrence could still mask variation across neighborhoods, leaving open the concern that some areas absorbed diverted traffic. As we discuss in Section 5.3, we find no such offsetting slowdowns across neighborhoods of different income levels or in areas of particular concern, such as the Bronx and New Jersey.

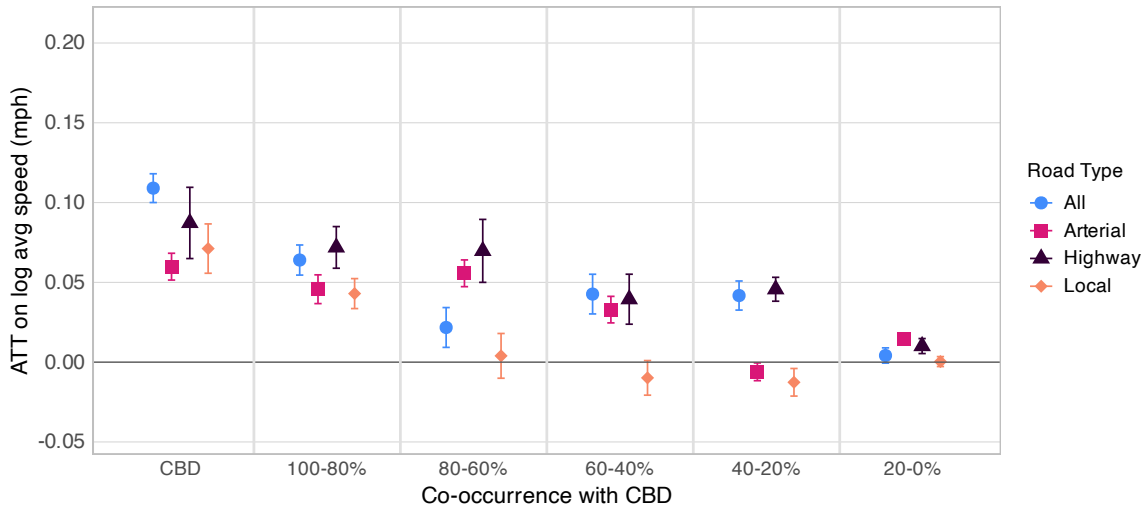
5.2 Effects on Trips throughout the Metro Area

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). For passenger cars, only trips that start outside the CBD and enter the CBD are subject to pricing, but all trips may be indirectly affected through the network spillovers documented above.[19] For each trip type, we again use the GSC approach to estimate the treatment effect of congestion pricing on log average speeds over entire trips.

We find that congestion pricing increased speeds not only on the priced trips to the CBD, but also on unpriced trips traveling from, within, and outside the CBD. Figure 4a plots the ATT on log average trip speeds for each trip type. The policy increased speeds on trips to the

[19]Note that trips that start outside the CBD and pass through it are also priced, even if their ultimate destination is outside of the CBD. We include such trips in the definition of 'to CBD.'

Figure 3: 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 estimated separately using 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.

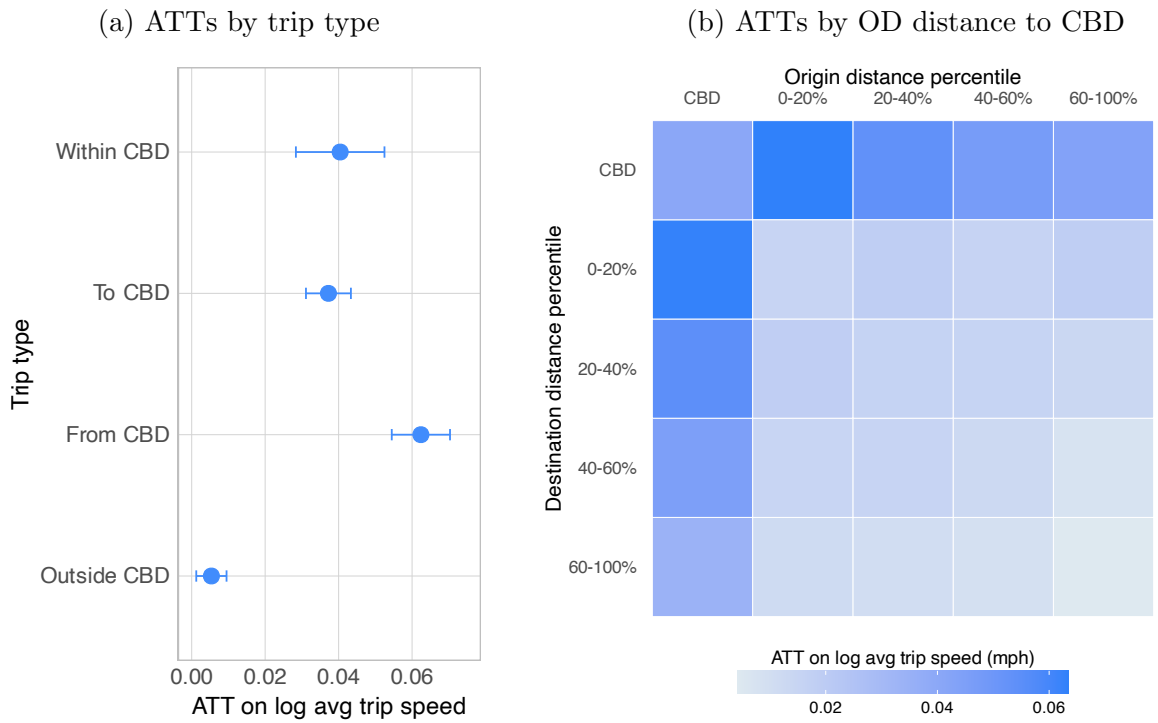
CBD by 3.7% and speeds on trips from the CBD by 6.2%. Outside the CBD, the average trip experienced modest gains of 0.54%. Figures C.2c–C.2f plot the ATTs for each trip type over time, which show little-to-no attenuation over time.

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 4b). We classify all origins and destinations by their within-city percentile distance to the CBD. This disaggregation shows the largest gains for trips directly involving the CBD, as well as nontrivial improvements for trips connecting outer areas. Trips traveling between the 0-20% distance bin and 60-100% bin, for example, experienced speed gains of 1–2% in each direction. Even for trips within the 60-100% distance bin, we estimate a statistically significant treatment effect of just 0.4%. As we discuss in Section 6, these changes on unpriced trips—including the small gains on the many trips outside of the CBD—have first-order implications for the policy's aggregate welfare impact.

5.3 Distributional Effects

Proposals to implement congestion pricing often face concerns over potential adverse impacts on residents of specific areas or on 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 5 documents the treatment effects on three outcomes: average speeds for road segments within the region (split by road type), average speeds on trips *to* the CBD that originate in the region, and

Figure 4: Effect on Trip Speeds throughout NYC



Notes: Panel a) documents the ATT on trip speeds for each class of trips. Panel b) 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. Each ATT is significantly different from zero at the 95% level, with standard errors clustered at the city-level.

average speeds on trips *from* the CBD that end within the specified region. Because a region's spatial relationship to the CBD may differ across cities, we use the set of all corresponding regions in our control cities as potential controls (e.g., the controls for the bottom income quintile are all income quintiles in other cities).

Household income. We find that speeds improved in tracts across the neighborhood income distribution, with larger increases in lower-income tracts. We group Census tracts into within-city quintiles based on their median household income in the 2019–2023 American Community Survey (ACS), which range from roughly $44,500 in the bottom quintile to $182,000 in the top. Speeds on roads in bottom quintile tracts increased by 3.8% compared to 1.9% for roads in tracts in the top income quintile. Similarly, for trips traveling to or from the CBD, speeds increased by 5-7% for each income quintile.

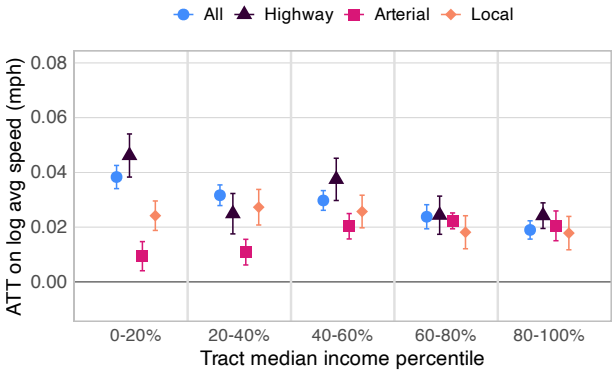
Specific regions of interest. We next examine effects across specific geographic areas. In New Jersey, average speeds increased by 12.9% on road segments in Hudson and Bergen Counties, the two counties most directly connected to Manhattan. This large effect is driven primarily by highways, many of which have high levels of co-occurrence with the CBD. Trips from these counties to the CBD became 6.7% faster, while trips from the CBD to these counties are now

6.9% faster. In the Bronx, which was often highlighted as an area that may experience negative spillovers because the Cross Bronx Expressway offers an alternative to going through the CBD (see, e.g., Rozner, 2024), average speeds on road segments increased by 9.8%, again driven mostly by improvements on highway segments. Bronx trips to and from the CBD became 5.4% and 4.5% faster, respectively. Finally, Long Island roads became 3.9% faster on average, and trips to and from the CBD increased by about 7%.

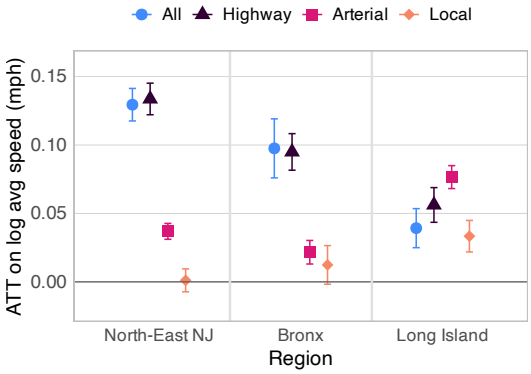
Figure C.4 further breaks out the effects on speeds for trips to the CBD based on the entrance they cross to enter the CBD. The largest effects are for trips crossing the Lincoln Tunnel, Holland Tunnel, Queens-Midtown Tunnel, and Queensboro Bridge. Effects on the main entrances from Brooklyn, including the Brooklyn Bridge and Manhattan Bridge, are generally smaller.

Figure 5: Effects for Specific Geographies

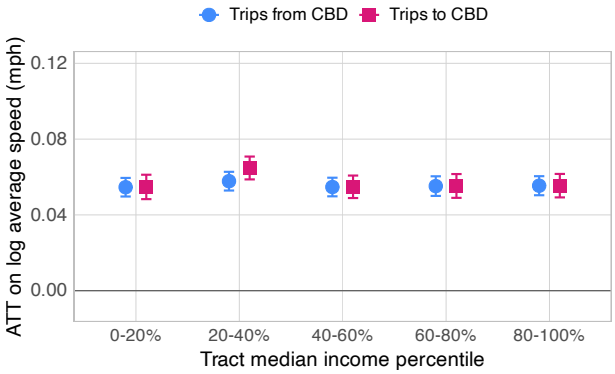
(a) Segment speeds by tract income



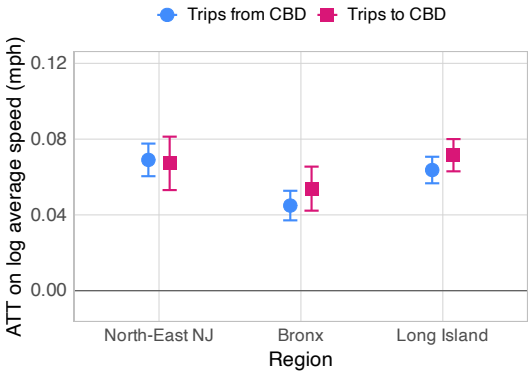
(b) Segment speeds for specific regions



(c) Trip speeds by tract income



(d) Trip speeds for specific regions



Notes: These figures plot the ATTs on 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 with standard errors clustered by city.

6 Driver Welfare

Congestion pricing raised speeds for priced and unpriced trips throughout the NYC metro area. In this section, we build a model to translate the changes in prices and travel times into an estimate of the welfare effects for both private passenger vehicle drivers and taxi/FHVs passengers. We show that the model can provide bounds on the welfare effects given data on pre-period volumes, without observing the demand response.

Pre-Policy Utility. Consider an individual who would typically drive between an origin o and destination d during peak hours prior to the implementation of congestion pricing, which we will denote as period 0. On an average day, the individual faces 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 hour.

There may be many individuals who would take the trip (o, d) if it were sufficiently attractive relative to their next-best option. We assume that each individual i considering a trip has 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 to take the trip is binary and switching modes is not explicitly considered. We assume that the baseline value of taking the trip ξ_{od} is fixed over time, and that the outside option—which could include other modes of travel that become more attractive with infrastructure investments—is non-decreasing: $\varepsilon^1_{i,od} \geq \varepsilon^0_{i,od}$.

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}; and 2) the average travel time t decreased by an origin-destination-specific amount Δt_{od}.[20] The utility that an individual obtains for taking the same trip after the policy was implemented, denoted as period 1, is given by:

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

[20] Auxiliary costs, such as gas, may have also decreased as a result of the improved fuel economy; however, we set aside any non-toll price changes for our main welfare analyses.

For some individuals, these changes pushed them over the margin of deciding whether to take a trip. 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 on priced trips for whom $u^1_{od} < \varepsilon^1_{i,od}$ may have switched to their outside option even if they had been taking the trip before.

Bounding Welfare Contributions. The net change in driver 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.[21] For each of the individuals who continued to make the same driving choice, the difference in utilities is $u^1_{od} - u^0_{od}$. 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 exact welfare impact. However, we can bound the welfare contribution of each trip based on *pre-policy* volumes, assuming that an individual would only start taking a given trip if it now provided higher utility (and vice versa for individuals who chose to stop taking a trip). This approach is similar in spirit to the "social savings" argument of Fogel (1964).

We describe the bounds formally here; Appendix B.3 offers a graphical illustration of the argument. 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, respectively, the numbers of individuals who stay with their driving choice ("stayers"), newly started taking the trip ("starters"), and stopped taking the trip after the policy ("stoppers"). The net contribution of trip (o, d) to the welfare impact of the policy

[21] 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.

is then bounded from below as follows:

$$\begin{aligned}
\Delta W_{od} &= \overbrace{N_{od}^{1,s} \times (u_{od}^{1} - u_{od}^{0})}^{\text{Stayers}} + \overbrace{N_{od}^{1,+} \times (u_{od}^{1} - \bar{\varepsilon}_{od,+}^{0})}^{\text{Starters}} + \overbrace{N_{od}^{1,-} \times (\bar{\varepsilon}_{od,-}^{1} - u_{od}^{0})}^{\text{Stoppers}} && (7)\\
&\geq N_{od}^{1,s} \times (u_{od}^{1} - u_{od}^{0}) + (N_{od}^{1,-}) \times (\bar{\varepsilon}_{od,-}^{1} - u_{od}^{0}) && (8)\\
&\geq N_{od}^{1,s} \times (u_{od}^{1} - u_{od}^{0}) + (N_{od}^{1,-}) \times (u_{od}^{1} - u_{od}^{0}) && (9)\\
&= N_{od}^{0} \times (u_{od}^{1} - u_{od}^{0}) && (10)\\
&= N_{od}^{0} \times (\omega \times \Delta t_{od} - \Delta p_{od}) && (11)
\end{aligned}$$

Here, line (7) breaks up the net welfare contribution by stayers, stoppers, and starters. For individuals who newly started taking the trip, $u_{od}^{1} - \bar{\varepsilon}_{od,+}^{1}$ must be greater than zero by revealed preference, and, by assumption, $\bar{\varepsilon}_{od,+}^{1} \geq \bar{\varepsilon}_{od,+}^{0}$, yielding the inequality in line (8). Similarly, for individuals who newly stopped taking the trip, $\bar{\varepsilon}_{od,-}^{1}$ must be greater than u_{od}^{1} by revealed preference, yielding the inequality in line (9). Line (10) is given by the accounting identity $N_{od}^{0} = N_{od}^{1,s} + N_{od}^{1,-}$. Finally, Line (11) results from subtracting equation (5) from equation (6).

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:

$$\Delta W \geq \sum_{o} \sum_{d} \left[N_{od}^{0} \times (\omega \times \Delta t_{od} - \Delta p_{od}) \right]. \tag{13}$$

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, for now, 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.

As in Figure 4b, 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

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

where $\widehat{\text{ATT}}_{od}$ are the estimated ATTs on trip speed from Figure 4b 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.[22]

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 (e.g., GPS devices, sensors embedded in roads, taxi/FHV companies, and administrative data) to simulate mobility patterns throughout entire metro areas.[23] Replica's primary customers are in the public sector, and include the MTA and other transit organizations. We use their estimated tract-to-tract travel flows by passenger vehicles and taxi/FHV for the NYC metro area for a typical weekday and weekend day in 2024 Q4, which we scale up to the full week and then sum across tracts in each of our origin-destination groups. For trips that enter the CBD, we supplement the data from Replica with MTA CBD entry data (MTA, 2026*b*), and scale the counts of trips to the CBD from Replica to exactly match the aggregate entries reported by the MTA for 2024. We describe this process further in Appendix B.2.

The final component needed to estimate the welfare bounds is a VOTT (ω). We use \$40/hour as a benchmark, which is just over 80% of the average hourly wage in the NYC metropolitan area.[24] We also identify several threshold VOTT values that would change the qualitative interpretation of our analysis.

Estimated Welfare Bounds. For private passenger vehicles, we find that aggregate driver welfare increased by at least \$7.4 million per week, with the positive welfare gains coming entirely from *unpriced* trips. Figure 6a plots the per-trip welfare gain for each of the OD pairs. The average passenger vehicle trip to the CBD paid \$7.9 to save just 2.3 minutes. With a VOTT of \$40/hour, this adds up to a total welfare change of no less than -\$10.0 million across the 1.57 million trips to the CBD each week. For these trips, the average driver would only be guaranteed to be better off on this trip following the implementation of congestion pricing if her VOTT were above \$207/hour, although these drivers likely also take other unpriced trips (e.g.,

[22]Drivers entering via the Lincoln Tunnel, Holland Tunnel, Queens-Midtown Tunnel, and the Hugh L. Carey Tunnel receive a \$3 'crossing 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, 2026*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.

[23]Replica volumes are estimated using both contemporaneous and historical data, and are not designed to capture sharp year-to-year changes in mobility patterns.

[24]According to the American Community Survey, the average hourly wage for workers aged 25-64 was \$48 per hour in 2024.

when they leave the CBD). If we assume each driver to the CBD is on a round-trip journey that ends back at their origin, the break-even VOTT falls to $109/hour.

Unpriced trips from, within, and outside the CBD save an average of just 9 seconds per trip. However, there are far more unpriced trips than priced trips. Across the 180.4 million weekly unpriced passenger vehicle trips, these time savings add up to a welfare gain of at least $17.3 million per week. Taking into account both priced and unpriced trips, the 'break-even' VOTT—i.e., the VOTT where the time savings of all drivers exactly offset the prices paid by paying drivers—for passenger vehicle trips is at most $25/hour.

Passengers of taxi/FHVs enjoy similar time savings but at a lower price than passenger vehicles, assuming that the only effect of the policy on their cost 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 for FHV after the policy was implemented. Nearly 70% of these trips are by FHVs rather than taxis, according to MTA (2026*a*), so we assume an average price paid across all taxi/FHV trips of $1.27. Thanks to the cheaper prices, priced taxi/FHV trips come closer to breaking even. The estimated lower bounds for per-trip welfare gains are highest (and often positive) for priced trips between the CBD and the outskirts of the metro area, which benefited from large time savings at relatively low prices. Welfare gains for taxi/FHV trips within the CBD are negative because the shorter trip durations yield smaller time savings for the same price. The break-even VOTTs are $46/hour for the average priced trip and $37/hour when including unpriced trips. While there are fewer taxi/FHV trips than private passenger vehicle trips, the gains for taxi/FHV passengers still add up to $0.2 million per week.

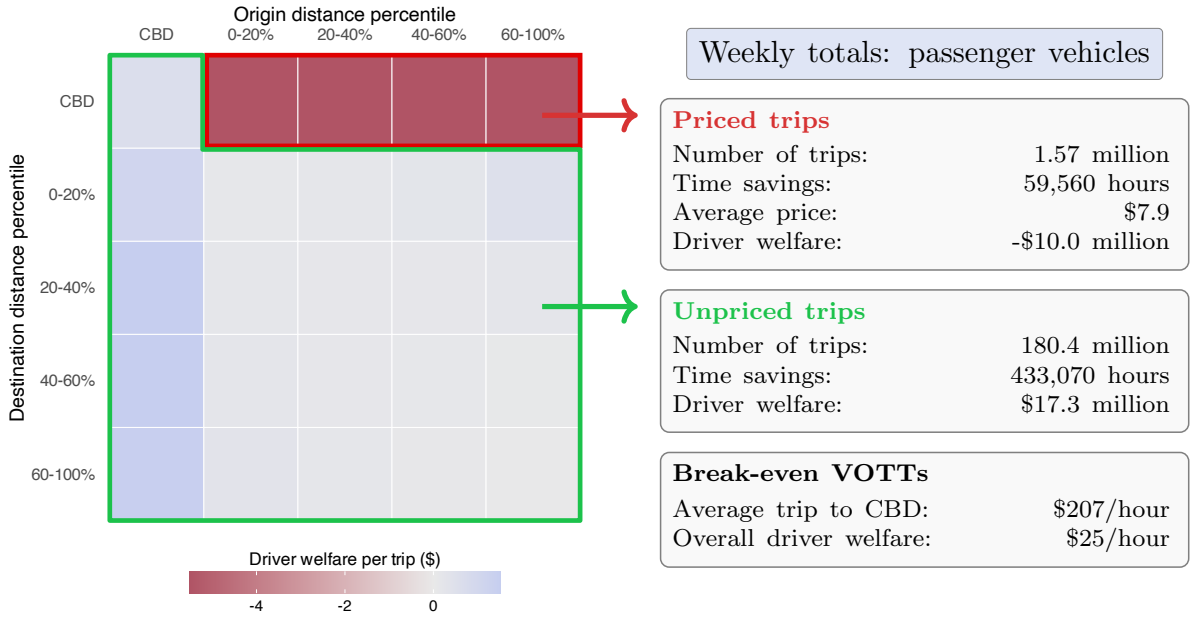
Table 1 summarizes the welfare effects.[25] In total, the estimates suggest that introducing congestion pricing increased welfare—even before any revenue recycling—by at least $7.5 million per week for passenger vehicle drivers and taxi/FHV passengers. While our main estimates assume a VOTT of $40/hour, pre-revenue welfare is positive so long as drivers and taxi/FHV passengers value their time at $26/hour or more. The second column of Table 1 reports welfare under a "pessimistic" scenario, which uses the lower bound of the 95% confidence intervals for each ATT. This increases the break-even VOTT to $41/hour and, at a VOTT of $40/hour, leads to a slightly negative pre-revenue welfare estimate (-$0.2 million). However, the program is also raising substantial revenue for the city, which is earmarked for public transit capital improvements. The MTA reported toll revenues of $1.02 billion for January 2025 through May 2026 and operating expenses of $186.4 million, i.e., about $14.0 million in gross revenue and $11.4 million in net revenue per week (MTA, 2026*c*). If the revenue is fully redistributed among drivers, then the break-even VOTT in our baseline case falls to just $5/hour.

The final column of Table 1 highlights the importance of accounting for the spillovers outside

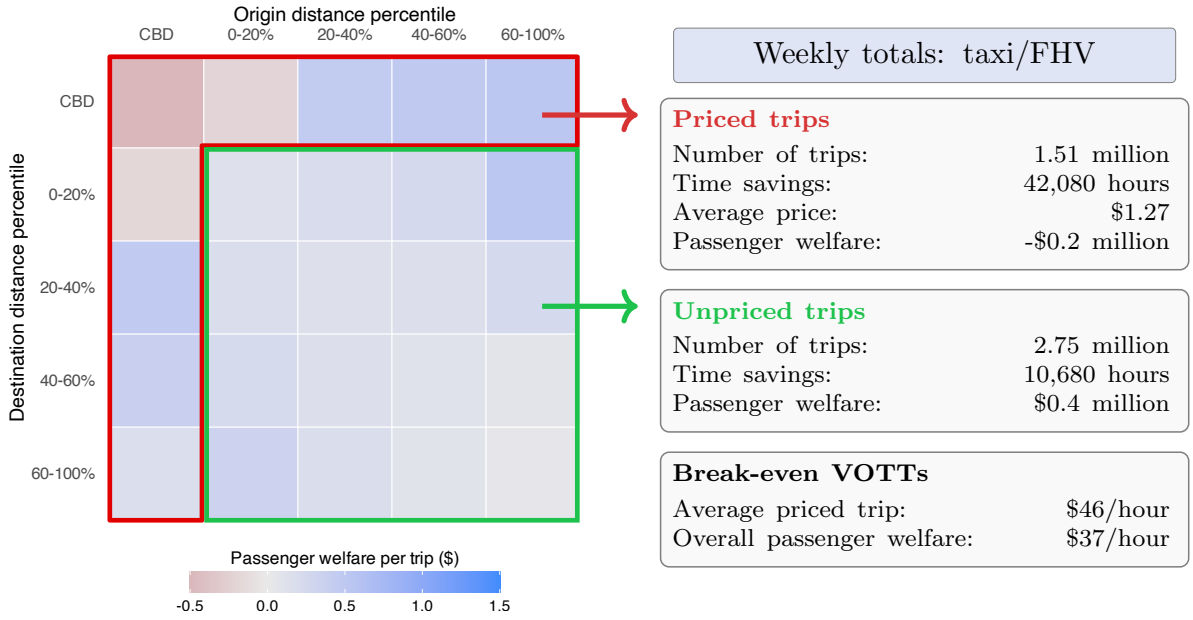
[25] In Appendix B.4, we document all of the input values for our baseline welfare calculations so that readers can explore the results under alternative inputs.

Figure 6: Driver welfare estimates

(a) Private passenger vehicles



(b) Taxi/FHV trips



Notes: Each square is shaded by the per-trip driver welfare, which is computed using the ATTs on trip speeds, the average pre-period duration, and the prices. Each ATT is significantly different from zero at the 95% level, with standard errors clustered at the city-level. The number of trips is the average tract-to-tract passenger car trips or taxi/FHV trips estimated by Replica for a typical week 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, 2026*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) based on the share of trips in each. The welfare estimates are computed for each cell separately, assume a Value of Travel Time (VOTT) of $40/hour, and exclude any revenue recycling.

the CBD by reporting welfare using only the effects on trips that are to, from, or within the CBD. Setting aside the spillovers outside the CBD flips the sign on pre-revenue weekly welfare. The estimated driver welfare falls to -$7.9 million, and the break-even VOTT increases to $89/hour. However, the total welfare effects remain positive ($3.5 million) if the revenue is fully redistributed, although the break-even VOTT (including revenue) increases to $18/hour.

Many excluded welfare channels would likely further increase total welfare gains, including reduced travel time variance, improved fuel economy, faster speeds for buses and commercial vehicles, faster speeds during off-peak hours, and potential environmental benefits. Longer-run adjustments may also affect the final welfare impacts of congestion pricing in NYC, although past work suggests these adjustments are a small share of the total welfare impact. For example, Hierons (2024) simulates the general equilibrium effects of a potential cordon toll in NYC (similar to the policy eventually implemented) and finds that about 10% of the welfare gains stem from worker and firm sorting.

Table 1: Summarizing the Welfare Effects

	Welfare per week (millions of $)		
	Baseline ($\widehat{ATT}$)	Pessimistic (Bottom of 95% CI)	Only trips to/from CBD
Private autos	7.35	-0.01	-7.62
Taxi/FHV passengers	0.18	-0.24	-0.25
Net revenue	11.40	11.40	11.40
Total	18.94	11.15	3.53
Break-even VOTTs			
Excluding revenue	$26/hour	$41/hour	$89/hour
Including revenue	$5/hour	$8/hour	$18/hour

Notes: This table summarizes the estimated welfare effects. The "baseline" and "to/from CBD" scenarios use the ATTs on trip speeds shown in Figure 4b, while the "pessimistic" scenario uses the bottom end of the 95% confidence interval for each ATT. The "to/from CBD" scenario also includes trips entirely within the CBD. The main welfare numbers assume a VOTT of $40/hour. The "break-even" VOTTs are computed as the VOTT for which driver welfare would be exactly zero, either excluding or including revenue (assuming revenue is fully redistributed). Revenue is based on MTA data and is net of reported operating expenses (MTA, 2026*c*).

7 Mechanisms

The effects of introducing congestion pricing on speeds and travel times depend on three underlying ingredients: 1) how pricing affects volumes throughout the network; 2) the exposure of other drivers to changes in these volumes (e.g., via the co-occurrence of segments they traverse with the CBD); and 3) the relationship between volumes and speeds on different road segments. While we cannot speak to how volumes respond to pricing in each city, we evaluate how the

latter two ingredients contributed to the policy's effectiveness in NYC, and whether similar effects may be expected in other cities.

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 of congestion pricing. Consider two extreme cases: 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 may improve, we would expect minimal spillovers.

We combine our segment-level measures of co-occurrence with data on trips between origin-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 road segment type. Using the same matrix of ODs as in Figure 4b to define trips, we compute the average co-occurrence of segments traversed as

$$\bar{e}_{od} = \frac{1}{|R_{od}|} \sum_{i \in R_{od}} \sum_{s \in R_i} \frac{t_{is} \times c_s}{t_i} \tag{15}$$

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 $R_i = \{s_1, s_2, \ldots, s_J\}$, c_s is the co-occurrence of segment s, and t_{is} and t_i are, respectively, the duration on segment s and the total trip duration.

We further aggregate up to trip types $l \in \{\text{To CBD}, \text{From CBD}, \text{Outside CBD}\}$, weighting by the average weekly pre-period volumes between each OD pair (N_{od}^{PV}):[26]

$$\bar{e}_l = \frac{\sum_{o,d} \bar{e}_{od} \times N_{od}^{\text{PV}}}{\sum_{o,d} N_{od}^{\text{PV}}} \tag{16}$$

The first three columns of Table 2 document this measure of exposure for NYC and each of our control cities. In all cities, trips to and from the CBD are naturally more exposed to the policy than trips entirely outside of the CBD. Including the part of the trip within the CBD, the weighted-average segment-level co-occurrence for trips to/from the CBD ranges from about 28% in Philadelphia to nearly 55% in Chicago. NYC drivers to/from the CBD have some of the largest exposure to other CBD drivers, second only to Chicago. For trips to or from the CBD in NYC, the average exposure to other CBD drivers is just over 50%. Outside the CBD, however, NYC is tied for the lowest average exposure alongside Philadelphia, with exposure outside the CBD of less than half that of Chicago, Boston, and Atlanta. Cities with high average exposure

[26] We focus just on passenger vehicle trips for this exercise, as taxis are a small share of the non-NYC markets. The results look similar when taxi/FHV trip counts are included.

on trips from and outside the CBD, such as Chicago, have travel patterns that suggest greater potential for positive spillovers on the speeds of unpriced trips.

Congestion functions. The next key ingredient is the local relationship between volumes and speeds. If all roads are already operating under free-flow conditions, then further decreases in CBD volumes would have no effect on travel times. 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 referred to as the "congestion function" of a road. Congestion functions vary across roads based on their physical characteristics (e.g., width, curvature, and lane structure) and typical driving behaviors (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 lane-mile). Our data cover only a sample of cars on the road, so we cannot directly observe total volumes to measure road density exactly. Instead, we infer density by assuming a constant scaling factor on observed traversals and normalizing relative to the 97.5th percentile of observed densities for a group of road segments. We then fit the following functional form, taken from the Bureau of Public Roads:

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

In this specification, v_{FF} denotes the free-flow speed, ρ_g is the inferred density of vehicles, κ is the capacity (the density at which congestion has a significant effect on speeds), α is a scaling parameter, and β governs the curvature.

We estimate separate congestion functions for each combination of road type and co-occurrence bin in each of our cities.[27] Figures 7a and 7b plot the congestion functions for the CBD and 60-80% co-occurrence bins by road type in NYC; the other NYC congestion functions are shown in Figure C.6. The circles along each congestion function curve represent peak-hour speeds on the average day prior to the policy's launch, and the triangles represent post-policy speeds, computed as the pre-period speeds multiplied by 1 plus the estimated ATT for this road type and co-occurrence bin.

This sample of congestion functions illustrates a few key points that hold across most cities

[27]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, similar to the approaches in Kreindler (2024) and Ater et al. (2026). 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.

and co-occurrence bins. First, congestion functions for local and arterial roads have lower 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 changes in density depend on the slope at the relevant density levels. A highway that is already operating near the speed limit will see little gain from further reductions in density, while a highway operating at a steep point along its congestion function may see large gains from relatively small reductions in density.

Table 2: Potential for Effects in Other Cities

City	Avg. exposure ($\bar{e}_l$)			Avg. elasticity of CF ($\bar{\eta}_l$).		
	To CBD	From CBD	Outside CBD	To CBD	From CBD	Outside CBD
NYC	51.4%	51.0%	3.0%	-0.84	-0.83	-0.60
Philadelphia	27.9%	27.2%	3.0%	-0.46	-0.45	-0.38
Chicago	54.4%	54.4%	6.8%	-1.04	-0.91	-0.64
Boston	36.6%	35.7%	8.1%	-0.71	-0.68	-0.47
Atlanta	43.3%	43.6%	7.4%	-0.71	-0.68	-0.53
Baltimore	34.7%	34.4%	4.3%	-0.40	-0.39	-0.41

Notes: The first three columns document the average exposure to the policy for trips to, from, and outside the CBD, measured as the weighted average co-occurrence of segments traversed with weights corresponding to the average duration on each segment. The latter three columns document the average elasticity of the congestion functions (CF) of roads traversed on each trip type, again weighting by the average duration spent on a given segment type.

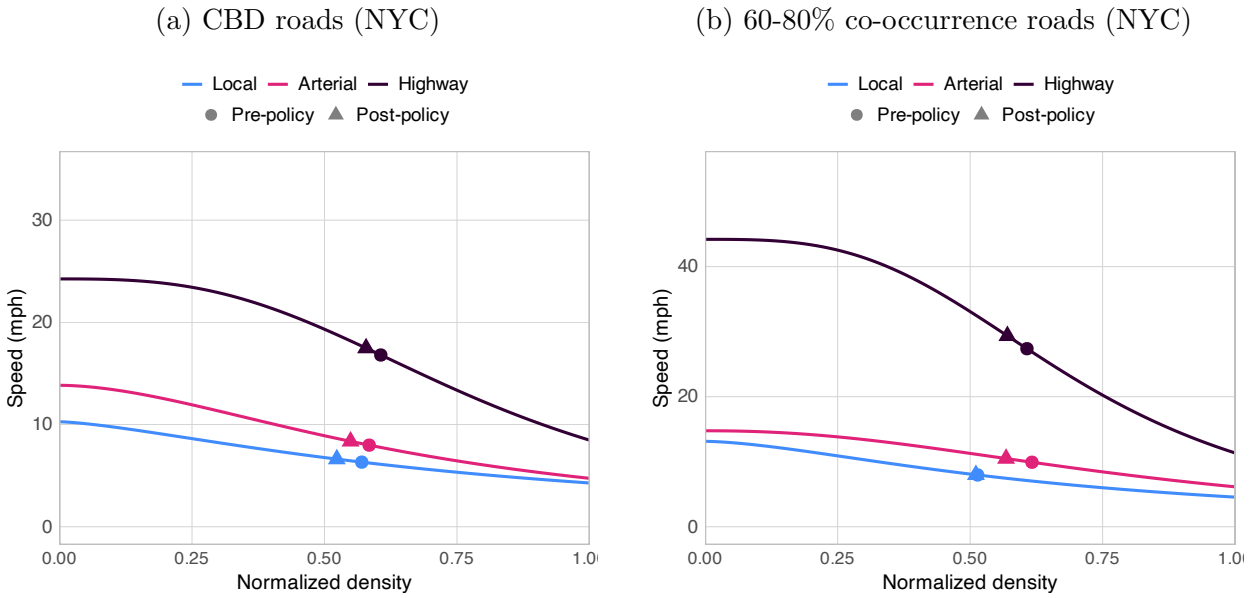
To compare the potential for volume changes to affect speeds in each city, we measure the elasticity of congestion functions on low- and high-co-occurrence roads, evaluated at the average pre-period densities. To further aggregate trips along these segments, we average across origin-destination pairs and weight by the average duration on roads corresponding to a given congestion function. More concretely, for a given OD pair we estimate the average local elasticity of the segment types that trips in that OD traverse as:

$$\bar{\eta}_{od} = \frac{\sum_c \sum_r \bar{t}_{od}^{cr} \left(v'_{cr}(\rho_{cr}) \frac{\rho_{cr}}{v_{cr}(\rho_{cr})} \right)}{\sum_c \sum_r \bar{t}_{od}^{cr}}, \tag{18}$$

where $v_{cr}(\cdot)$ is the estimated congestion function for co-occurrence bin c and road type r, ρ_{cr} is the corresponding pre-period average density, and $\bar{t}_{od}^{cr}$ is the average pre-period duration spent on this co-occurrence bin and road type combination for trips in a given OD pair. As in Equation (16), we further aggregate up for each trip type $l \in \{\text{To CBD}, \text{From CBD}, \text{Outside CBD}\}$ with weights corresponding to the number of pre-period passenger vehicle trips in a given OD pair (N_{od}^{PV}).

The last three columns of Table 2 document $\bar{\eta}_l$ for each city and trip type. Prior to the

Figure 7: Example Congestion Functions



Notes: This figure plots estimated congestion functions for roads in NYC. The congestion functions are estimated using observations of contemporaneous speeds and normalized densities, as shown in 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 using the estimated ATT on speeds.

implementation of congestion pricing, NYC roads traversed by drivers on trips to, from, and outside the CBD were operating at relatively steep points on their congestion functions. On trips to and from the CBD, NYC roads were about 18% more locally elastic than in Boston or Atlanta, the next-ranked cities, and over 80% more elastic than in Philadelphia and Baltimore, the least elastic. The one exception is Chicago, where roads on trips to and from the CBD were about 24% more elastic than in NYC. Roads traversed on trips outside the CBD were likewise among the most elastic in NYC, again second only to Chicago.

Read alongside the exposure measures, these elasticities suggest that NYC realized large spillovers on trip speeds outside the CBD despite relatively low exposure because its congestion functions were steep enough to turn small volume reductions into faster speeds, even on lightly exposed roads. For these unpriced trips entirely outside of the CBD—which contribute substantially to overall welfare—the average exposure in NYC was half that of some other cities. While speeds on trips to/from the CBD in other cities would likely be less affected by changes in volumes—they have lower elasticities and lower average exposure—spillovers on trips outside their respective CBDs may be larger if the higher exposure offsets lower elasticities. A notable exception is Chicago, which ranks above NYC on both margins—higher exposure to CBD trips and steeper congestion functions on every trip type—so a comparable cordon toll there could plausibly generate even larger speed gains throughout the metro region.

Finally, the influence of exposure to CBD trips and congestion function elasticity depends on the nature of the demand response to congestion pricing. Although we cannot observe this

directly in New York, our results are consistent with a large extensive margin response in which drivers to the CBD substituted to other modes of transportation rather than driving to other nearby locations. If drivers in other cities were to respond in other ways—for instance, because alternative modes of transit are less appealing—then the nature of both the direct effects and the spillover effects of cordon-based congestion pricing might be quite different.

8 Conclusion

The introduction of congestion pricing in New York City offers new evidence on how cordon-based pricing can reshape urban mobility beyond the tolled area itself. Inside the cordon zone, introducing pricing raised average speeds while producing no detectable effect on air quality, consumer spending, or foot traffic. Speeds also rose outside the cordon area, with gains that scale with each road's exposure to CBD-bound traffic. Many areas often predicted to absorb displaced traffic, including the Bronx and the New Jersey counties facing Manhattan, instead experienced some of the largest speed gains.

Accounting for these spillovers is first-order for measuring the policy's welfare effects, and evaluations limited to the cordon itself would substantially understate the welfare gains. Before accounting for revenue, we estimate that drivers on trips to the CBD are only better off from the policy if their value of travel time is at least \$207/hour (or \$109/hour if they are on a round-trip journey). Across all trips, however, the break-even value of time falls to \$26/hour. For a value of time of \$40/hour, the policy increases driver welfare by \$7.5 million in addition to \$11.4 million in net revenue. If this revenue is fully re-distributed, then the break-even value of time falls to just \$5/hour.

These welfare estimates also underscore a tension central to the political economy of congestion pricing. Costs are concentrated on a visible group that pays every day, while the gains are spread thinly across many unpriced trips. This is the classic configuration in which efficient policies struggle to survive (Olson, 1965; Wilson, 1980), perhaps consistent with the program's near-cancellation in 2024 and the scarcity of congestion pricing programs in other cities.

Whether other cities should expect similar results depends, in part, on primitives that can be measured before adoption. New York's roads operated on unusually steep regions of their congestion functions, so small volume reductions produced large speed gains even though trips outside its CBD were the least exposed to CBD traffic in our sample. Chicago ranks above New York on both exposure and congestion function elasticity, while Philadelphia and Baltimore rank below on both. However, whether a cordon-based congestion price would generate a similar demand response in other cities as in NYC remains an open question. New York's vast transit system offers a better substitute to driving than exists in most cities, and a cordon toll where drivers have no practical alternative to driving may deliver more diversion and less relief.

References

Adao, Rodrigo, Costas Arkolakis, and Federico Esposito. 2025. "General Equilibrium Effects in Space: Theory and Measurement." *American Economic Journal: Macroeconomics*. forthcoming.

Akbar, Prottoy A, Victor Couture, Gilles Duranton, and Adam Storeygard. 2023. "The fast, the slow, and the congested: Urban transportation in rich and poor countries." National Bureau of Economic Research.

Allen, Treb, and Costas Arkolakis. 2022. "The welfare effects of transportation infrastructure improvements." *The Review of Economic Studies*, 89(6): 2911–2957.

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel G. Hickok, and Tobias Salz. 2025. "Optimal Urban Transportation Policy: Evidence from Chicago." NBER Working Paper No. 32185.

Anderson, Michael L. 2014. "Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion." *American Economic Review*, 104(9): 2763–2796.

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, and Shoshana Vasserman. 2026. "Road Pricing under Heterogeneity."

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 Waxman. 2024. "The Value of Urgency: Evidence from Real-Time Congestion Pricing." *Journal of Political Economy Microeconomics*, 2(4): 786–851. Publisher: The University of Chicago Press.

Börjesson, Maria, Jonas Eliasson, Muriel B Hugosson, and Karin Brundell-Freij. 2012. "The Stockholm congestion charges—5 years on. Effects, acceptability and lessons learnt." *Transport Policy*, 20: 1–12.

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*.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz. 2025. "Personalized pricing and the value of time: Evidence from auctioned cab rides." *Econometrica*, 93(3): 929–958.

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.

Busso, Matias, Jesse Gregory, and Patrick Kline. 2013. "Assessing the Incidence and Efficiency of a Prominent Place Based Policy." *American Economic Review*, 103(2): 897–947.

Butts, Kyle. 2021. "Difference-in-differences estimation with spatial spillovers." *arXiv preprint arXiv:2105.03737*.

Choudhury, Shushman, Abdul Rahman Kreidieh, Iveel Tsogsuren, Neha Arora, Carolina Osorio, and Alexandre M Bayen. 2024. "Scalable Learning of Segment-Level Traffic Congestion Functions." In *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*. 1917–1924, IEEE.

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*.

Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2018. "Speed." *The Review of Economics and Statistics*, 100(4): 725–739.

Criscuolo, Chiara, Ralf Martin, Henry G Overman, and John Van Reenen. 2019. "Some causal effects of an industrial policy." *American Economic Review*, 109(1): 48–85.

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.

Diamond, Peter A. 1973. "Consumption externalities and imperfect corrective pricing." *The Bell Journal of Economics and Management Science*, 526–538.

Downs, Anthony. 1962. "The law of peak-hour expressway congestion."

Dreyfuss, Bnaya. 2025. "Equilibrium Neglect and Political Feasibility."

Duranton, Gilles, and Matthew A. Turner. 2011. "The Fundamental Law of Road Congestion: Evidence from US Cities." *American Economic Review*, 101(6): 2616–2652.

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.

Fajgelbaum, Pablo D, and Edouard Schaal. 2020. "Optimal Transport Networks in Spatial Equilibrium." *Econometrica*, 88(4): 1411–1452.

Fogel, Robert William. 1964. *Railroads and American economic growth.* Johns Hopkins Press Baltimore.

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.

Goldszmidt, Ariel, John A List, Robert D Metcalfe, Ian Muir, V Kerry Smith, and Jenny Wang. 2020. "The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments."

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.

Gu, Yizhen, Chang Jiang, Junfu Zhang, and Ben Zou. 2021. "Subways and road congestion." *American Economic Journal: Applied Economics*, 13(2): 83–115.

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.

He, Meilu, Nueraili Kuerbanjiang, and Suresh Dhaniyala. 2019. "Performance characteristics of the low-cost Plantower PMS optical sensor." *Aerosol Science and Technology*, 54(2): 232–241.

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*.

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.

Kline, Patrick, and Enrico Moretti. 2014. "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority." *The Quarterly journal of economics*, 129(1): 275–331.

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.

Kreindler, Gabriel. 2024. "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications." *Econometrica*, 92(4): 1233–1268. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA18422.

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*.

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.

Miguel, Edward, and Michael Kremer. 2004. "Worms: identifying impacts on education and health in the presence of treatment externalities." *Econometrica*, 72(1): 159–217.

MTA. 2025. "Reduction in Vehicle Entries to the CBD in 2025." https://metrics.mta.info/?cbdtp/vehiclereductions.

MTA. 2026*a*. "Congestion Relief Zone Tolling: First Evaluation Report." https://www.mta.info/document/195631.

MTA. 2026*b*. "MTA Congestion Relief Zone Vehicle Entries: Beginning 2025." https://data.ny.gov/Transportation/MTA-Congestion-Relief-Zone-Vehicle-Entries-Beginni/t6yz-b64h/about_data.

MTA. 2026*c*. "MTA Statement of Operations: Beginning 2019." https://data.ny.gov/Transportation/MTA-Statement-of-Operations-Beginning-2019/yg77-3tkj/about_data.

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.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Vol. 124, Cambridge, MA:Harvard University Press.

Ostrovsky, Michael, and Frank Yang. 2025. "Effective and equitable congestion pricing: New York City and beyond."

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

Rozner, Lisa. 2024. "Congestion pricing could cause more traffic, higher pollution in the Tri-State Area. Here's where." *https://www.cbsnews.com/newyork/news/congestion-pricing-expected-to-cause-more-traffic-higher-pollution-in-some-parts-of-* Accessed: July 15, 2026.

Rubin, Donald B. 1980. "Randomization analysis of experimental data: The Fisher randomization test comment." *Journal of the American statistical association*, 75(371): 591–593.

Sandmo, Agnar. 1975. "Optimal taxation in the presence of externalities." *The Swedish Journal of Economics*, 86–98.

Schaller, Bruce. 2010. "New York City's congestion pricing experience and implications for road pricing acceptance in the United States." *Transport Policy*, 17(4): 266–273.

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.

Tonne, Cathryn, Sean Beevers, Ben Armstrong, Frank Kelly, and Paul Wilkinson. 2008. "Air pollution and mortality benefits of the London Congestion Charge: spatial and socioeconomic inequalities." *Occupational and Environmental Medicine*, 65(9): 620–627.

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.

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.

Wilson, James Q. 1980. *The Politics of Regulation.* New York, NY:Basic Books.

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

Yang, Natalie J. 2025. "Essays in the Economics of Transportation." PhD diss. Columbia University.

Supplementary Appendices

A Data Appendix **1**
A.1 CBD Definitions 1
A.2 Estimating Fuel Efficiency 2
A.3 PurpleAir Pollution Measures 2
A.4 Consumer Spending Data from MBHS3 4
A.5 Foot Traffic Data from Advan 4
A.6 MTA Counts of CBD Entries 6

B Results Appendix **8**
B.1 Robustness Tests 8
B.2 Origin-Destination Pre-Period Volumes 9
B.3 Illustrative Example 12
B.4 Welfare Calculation Inputs 13

C Supplementary Tables and Figures **14**

References for appendices **21**

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	Baltimore Police Department	The "Central District"
🔗	Boston	Boston Air Pollution Control Commission	The outline of the red Boston Proper parking zone
🔗	Chicago	City of Chicago	The "Downtown Zone Area".
🔗	New York City	Metropolitan Transportation Authority	The outline of the "Congestion Relief Zone".
🔗	Philadelphia	Philadelphia City Planning Commission	The "CTR Center City Overlay"

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.

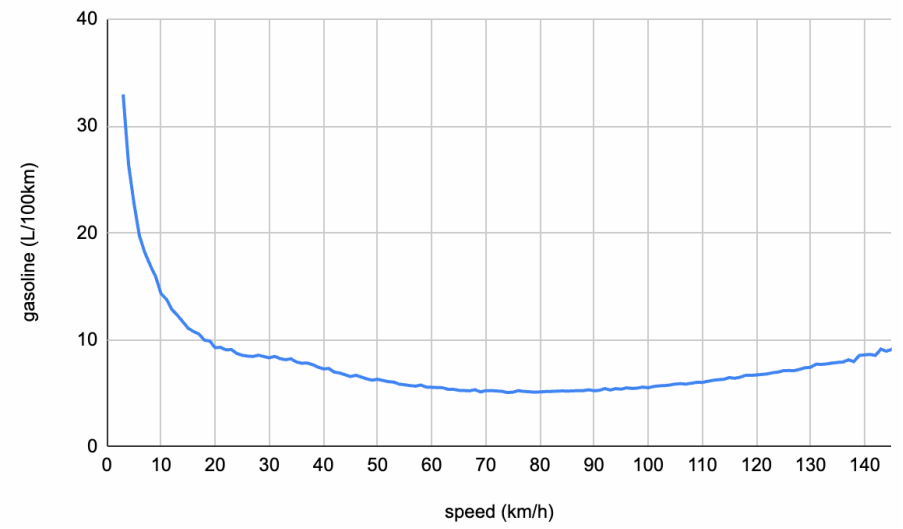
Fuel consumption modeling has been widely studied in the literature (Faris et al., 2011). Models of this form roughly fall into two categories: principled models (Faris et al., 2011), which aim to model the physics underlying energy usage, and empirical models (Department of Energy , DOE; Ersal et al., 2012), which fit non-parametric models to ground-truth fuel consumption data. For our purposes, we integrate models from the *National Renewable Energy Laboratory* (NREL), 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 computational 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 PurpleAir sensors in the corresponding CBSA. For redundancy, each sensor includes two independent

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.

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 is 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. We compute the quantities

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

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

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

given sensor as

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

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

4. Finally, we compute the daily adjusted $\text{PM}_{2.5}$ averages, dropping any days with 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 $\text{PM}_{2.5}$. Some sensors are missing information for one or more days. We impute missing sensor-date $\text{PM}_{2.5}$ 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, and the data exclude online transactions. We define a zipcode as lying within a city's CBD if its centroid is contained by the CBD boundaries, and aggregate the data to the CBD-date level for estimation. For NYC, we manually include zipcode 10004. This zipcode includes Governors Island, which pulls its centroid out into the river, but most of its merchants are in the southern portion of the Financial District. The transaction times are only available at the date-level, so the data include both peak and off-peak transactions.

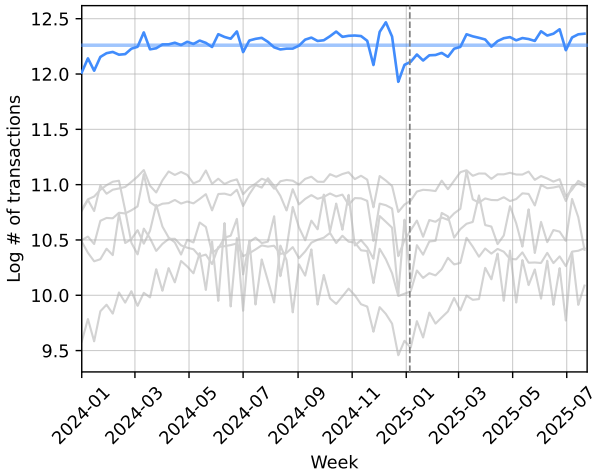
Table A.2 documents the total number of transactions and aggregate transaction amount by category. The final three columns subset to zipcodes within the CBD. The average transaction size in both restaurant and retail categories tends to be higher for merchants in the CBD than in the overall city. Figure A.2 plots the number of transactions and total spend by category over time.

A.5 Foot Traffic Data from Advan

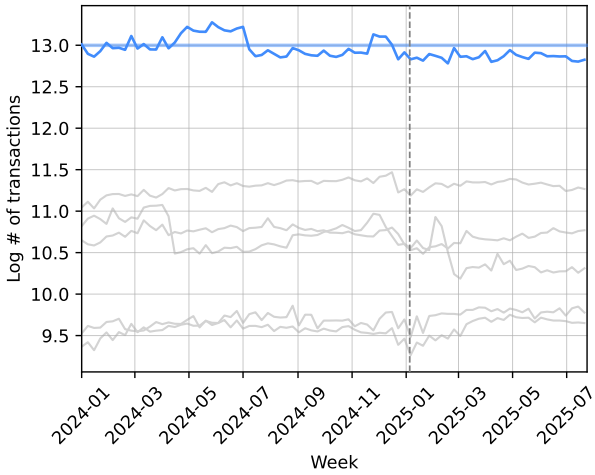
The raw data from Advan Neighborhood Patterns include the number of stops by GPS devices at the hour-block group level. We restrict to peak hours and combine across all block groups within the same Census tract.

Figure A.2: Time Series of Spending

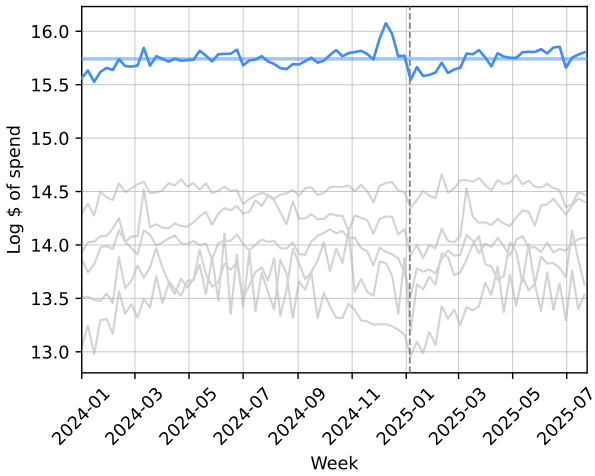
(a) Restaurants – log # transactions



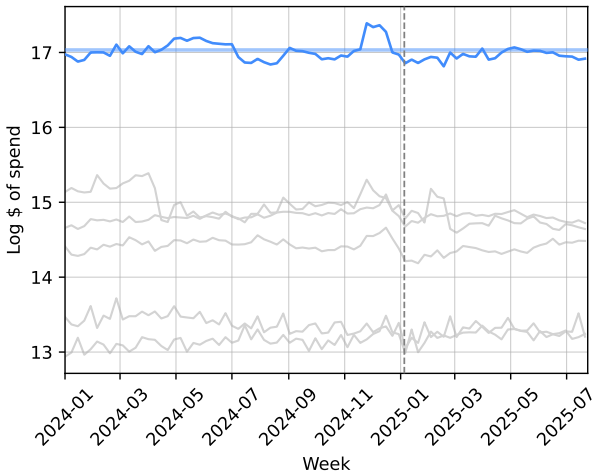
(b) Retail – log # transactions



(c) Restaurants – log $ of spend



(d) Retail – log $ of spend



Notes: These figures document the weekly average number of transactions and total spending in the MBHS3 data for zipcodes within each of the sample city CBDs between January 2024 and July 2025. The blue horizontal line in each panel is the pre-policy average for the NYC CBD.

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

		All			CBD-only		
City	Category	Count (millions)	Amount (millions $)	Avg. size ($/trans.)	Count (millions)	Amount (millions $)	Avg. size ($/trans.)
NYC	Restaurant	129.8	3561.5	27.4	17.6	566.6	32.1
	Retail	301.8	14094.4	46.7	35.0	2020.1	57.8
ATL	Restaurant	126.4	3010.7	23.8	3.2	77.4	24.3
	Retail	281.4	12140.0	43.1	1.4	43.7	32.2
BAL	Restaurant	47.6	881.7	18.5	1.9	59.4	30.8
	Retail	74.8	3294.9	44.0	1.2	52.5	43.6
BOS	Restaurant	41.9	1031.7	24.6	2.8	98.5	34.7
	Retail	66.4	3576.2	53.9	3.5	255.3	73.8
CHI	Restaurant	76.3	1920.3	25.2	4.2	124.0	29.5
	Retail	154.2	7769.9	50.4	3.8	150.9	40.1
PHL	Restaurant	72.0	1896.1	26.3	5.0	165.7	32.9
	Retail	215.2	9049.8	42.0	6.7	224.2	33.6
Total	Restaurant	494.0	12302.1	24.9	34.8	1091.6	31.3
	Retail	1093.9	49925.2	45.6	51.4	2746.7	53.4

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 July 2025.

Figure A.3 plots the log average weekly foot traffic in the CBDs of our sample cities. The horizontal blue line corresponds to the pre-policy average for NYC. CBD foot traffic increases around the holiday period each year. In general, the patterns and levels of foot traffic in 2025 are similar to those in 2024, before the policy launched. The second panel plots a head-to-head of NYC, Chicago, and Philadelphia. NYC and Philadelphia foot traffic measures follow each other closely, while Chicago is more likely to diverge from the other two cities.

A.6 MTA Counts of CBD Entries

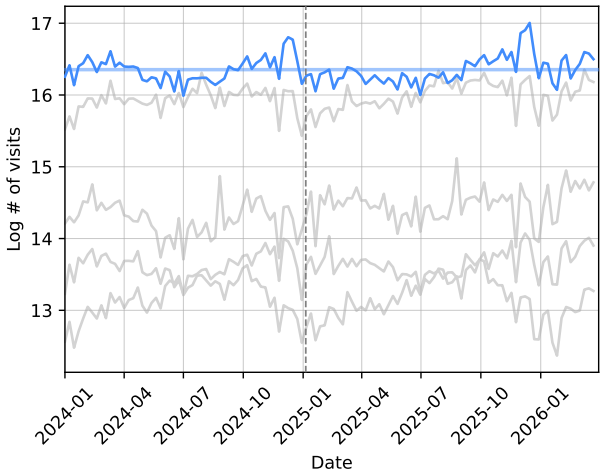
The MTA provides data on all CBD entries since congestion pricing launched, split by entry point, 10-minute interval, and vehicle class. Panel a) of Table A.3 documents the count and share of peak-hour entries by class for the average day between January 2025 and June 2026. The majority of entries are taxi/FHVs ("TLC Taxi/FHV") and passenger car ("Cars, Pickups and Vans") entries.

Panel b) documents how the shares have changed over time. There has been a small shift toward taxi/FHV trips (33-34% in first two quarters of 2025 to just over 35% in the same period in 2026). This may reflect the relatively cheaper congestion fees for taxi/FHVs, consistent with the hypothesis in Ostrovsky and Yang (2025).

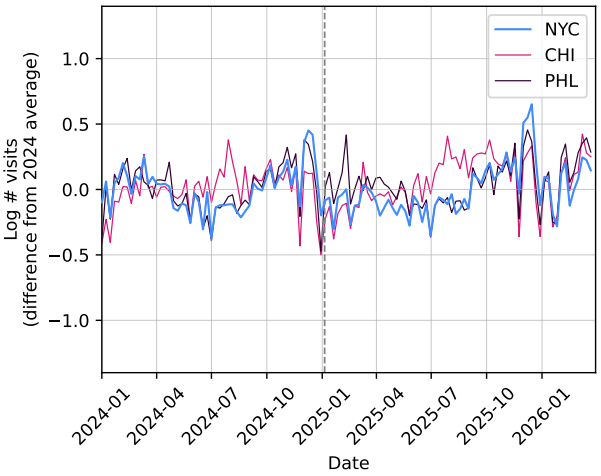
Figure A.4 plots the average number of entries per week by time of day for weekdays. There

Figure A.3: CBD Foot Traffic

(a) Log average stops per day



(b) Log stops relative to pre-policy average



Notes: These figures plot weekly CBD foot traffic, restricted to peak hours. The first panel plots the log total number of visits. The blue horizontal line is the pre-policy 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 pre-policy average.

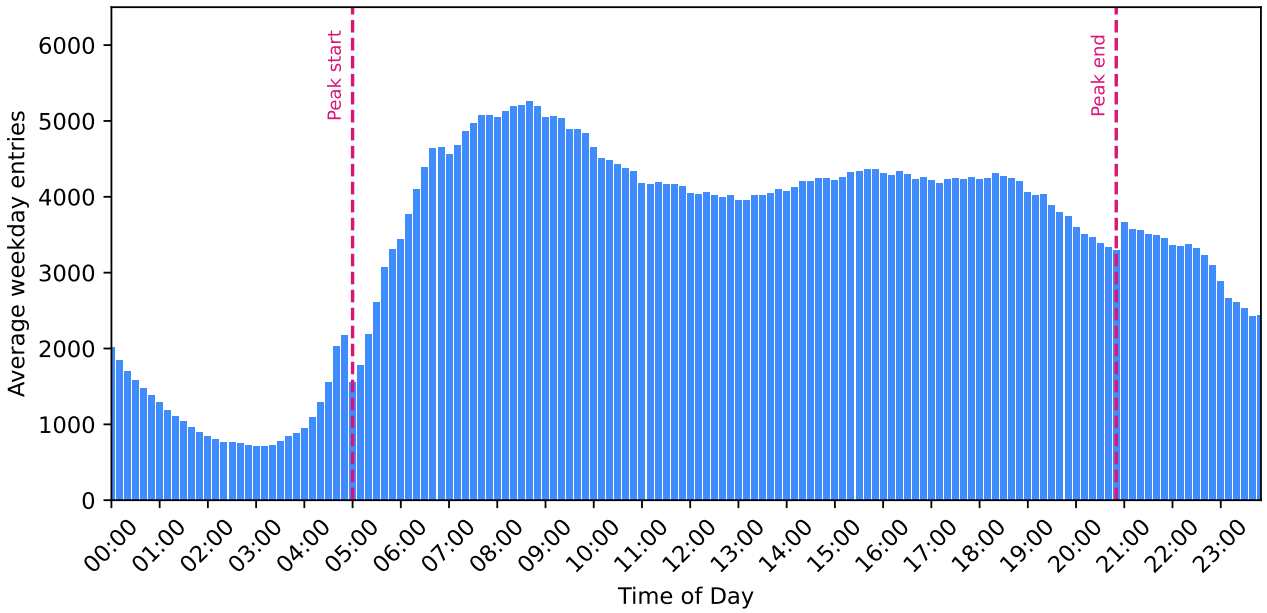
Table A.3: CBD entries by vehicle class

Panel A) Daily average	Avg. entries	Frac. of all entries
Cars, Pickups and Vans	220,346	0.593
Taxi/FHV	127,518	0.343
Single-Unit Trucks	14,689	0.040
Buses	6,907	0.019
Motorcycles	1,279	0.003
Multi-Unit Trucks	716	0.002

Panel B) Over time	Cars/Pickups/Vans	Taxi/FHV	Other
2025Q1	0.590	0.340	0.069
2025Q2	0.605	0.331	0.064
2025Q3	0.601	0.335	0.064
2025Q4	0.591	0.351	0.058
2026Q1	0.587	0.352	0.061
2026Q2	0.584	0.352	0.064

Notes: This table documents the average entries into the CBD for peak hours between Jan 5, 2025 to June 30, 2026. The second panel shows how the share of vehicle classes has changed over time. Data are from the "MTA Congestion Relief Zone Vehicle Entries" dataset made public by the MTA (MTA, 2026).

Figure A.4: MTA-reported CBD entries by time of day



Notes: This figure plots the average number of entries per week by 10-minute bins of time of day. The data are subset to just weekdays and cover entries from January 2025 through June 2026.

is noticeable bunching before peak-hour pricing begins and after it ends, suggesting that some drivers respond to the price by changing when they enter the CBD.

B Results Appendix

B.1 Robustness Tests

We assess the robustness of our primary estimates in two ways, both of which are reported in Figure B.1. First, we re-estimate the treatment effects while dropping one control city at a time, to check that no single comparison city is driving the results (Panel a). Across the six main speed outcomes, the estimates are stable in sign and broadly similar in magnitude regardless of which city is excluded. The largest movements come from excluding Boston, which increases the estimated ATT relative to the baseline specification for a handful of outcomes.

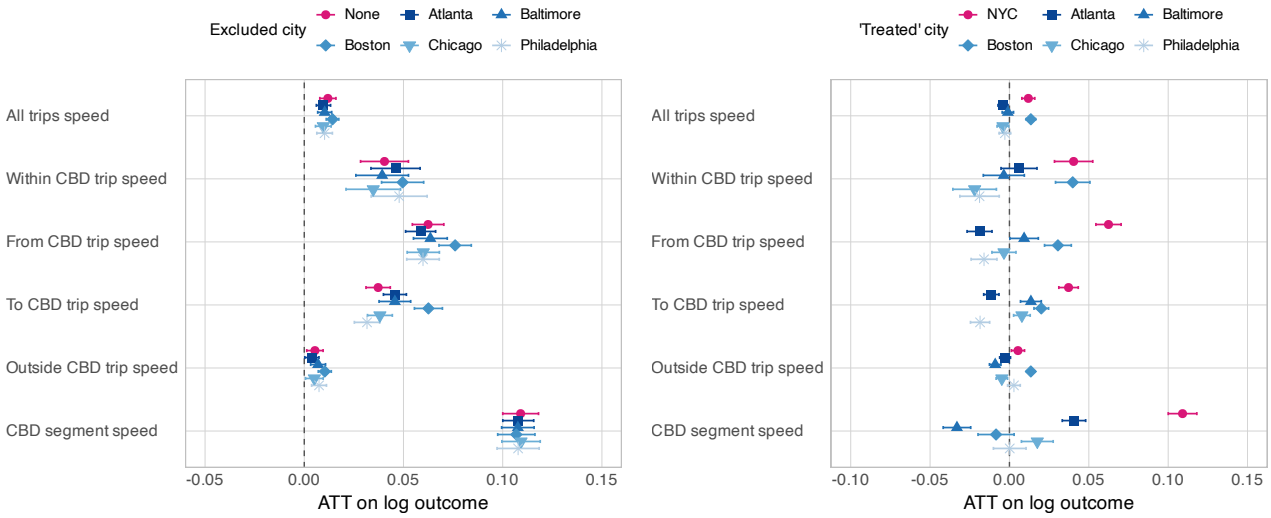
Second, we conduct a placebo exercise in which we set NYC aside and estimate "treatment" effects where we act as if each control city adopted congestion pricing, using the remaining cities as controls (Panel b). Large placebo effects would indicate that the design is picking up city-specific shocks rather than the policy itself. While some of the placebo estimates are statistically significant, the estimates for any given city do not show the same systematic increase in speed across outcomes that we see in NYC. Instead, the signs vary across outcomes within a city.

The exception is Boston, whose placebo estimates are positive across nearly every outcome.[28] However, including a control city that has a spurious positive effect will bias our NYC estimates *downward*, since the synthetic control it helps construct overstates counterfactual speeds in the absence of the policy. This is the same direction implied by Panel a), where dropping Boston raises the estimated ATT for several outcomes.

Figure B.1: Robustness for Primary Estimates

(a) Alternate control cities

(b) Placebo treatment cities



Notes: This figure documents the robustness of our primary speed estimates. Panel a) re-estimates the ATT on each outcome while excluding one control city at a time from the donor pool. Panel b) reports placebo estimates in which NYC is set aside and each control city is treated in turn as if it had adopted congestion pricing, using the remaining cities as its synthetic control. Each point is the aggregate ATT on the corresponding log outcome over the full post-treatment period. Horizontal lines denote 95% confidence intervals. Standard errors are clustered at the city level.

B.2 Origin-Destination Pre-Period Volumes

For computing the aggregate welfare effects in Section 6, we need the number of passenger vehicle and taxi/FHV trips between each set of origins and destinations. One limitation of the data from Google Maps is that we cannot reliably observe volumes, much less changes in volumes over time. We instead estimate these values for the pre-period (September to December 2024) using a combination of data from Replica and the MTA.

We start with estimated tract-to-tract flows for a 'typical' Thursday and Saturday in the fourth quarter of 2024 from Replica, an urban data platform company that spun out of Google in 2019. Replica ingests data from dozens of sources, including location data from GPS devices,

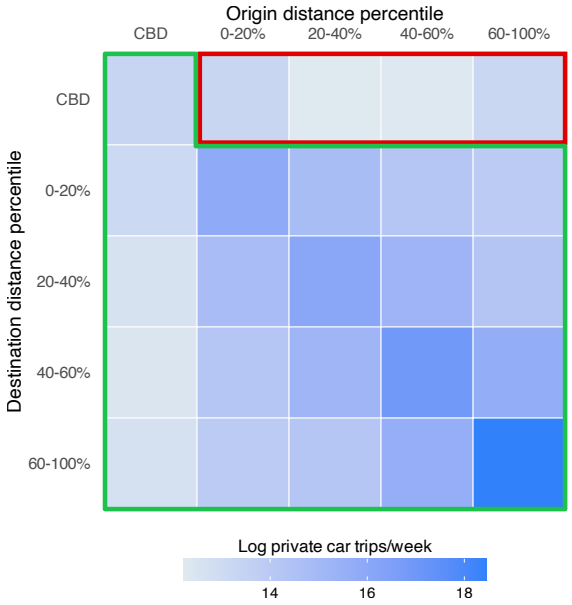
[28] There are two events in Boston that may drive this positive effect. First, the Massachusetts Bay Transportation Authority completed a large track improvement program in December 2024, which substantially improved transit speeds and may have led to some substitution to public transit. Second, the Sumner Tunnel was closed for restoration during some weekends in Fall 2024, which likely worsened congestion during our pre-treatment period.

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 full-population 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 aggregate to typical weekday and weekend flows to the weekly level by multiplying the weekday flows by five and the weekend flows by two.

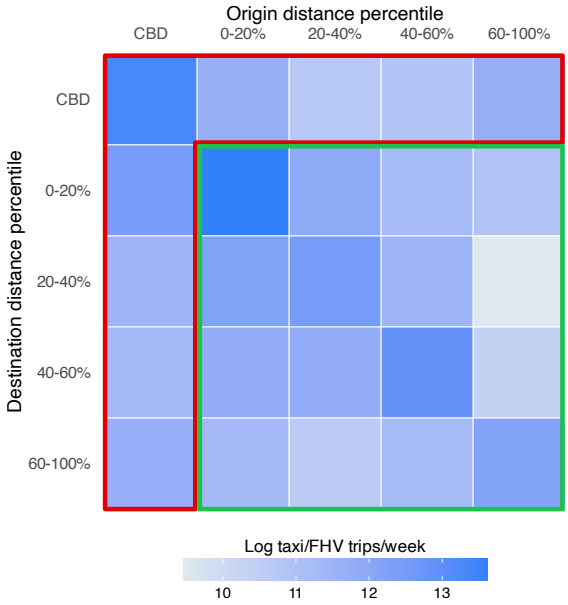
We supplement data from Replica with data from the MTA on CBD entries. CBD entries are an especially important source of welfare in our setting, and the MTA provides realized counts (Table A.3). First, we scale the estimated entries by Replica to exactly match the weekly passenger vehicle and taxi/FHV counts reported by the MTA for our main sample period. 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 change from pre-period estimated CBD entries reported by MTA (2026). This approach combines Replica data on the relative distribution of origin tracts for CBD trips with MTA data on the overall levels. Note that the MTA entries data include priced trips that pass through the CBD en route to other destinations, which may explain some of the discrepancy in CBD trip levels between the MTA and Replica counts.

Figure B.2: Estimated Weekly Trip Counts

(a) Passenger vehicles



(b) Taxi/FHVs

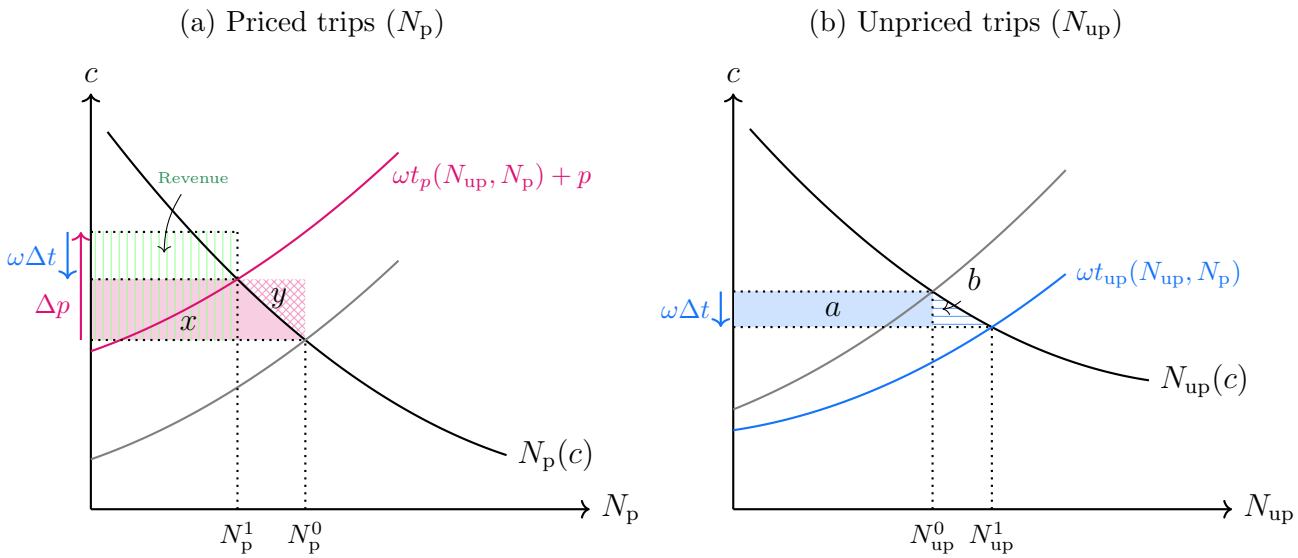


Notes: These figures plot the log weekly estimated trip counts between bins of origins and destinations. The trip counts are estimated using a combination of data from Replica and the MTA and reflect pre-congestion pricing flows. The red polygons indicate priced trips and the green polygons indicate unpriced trips.

B.3 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 $\omega t(N_{\text{up}}, N_{\text{p}})$ and the congestion fee p, setting aside all other costs. As drawn, travel times increase convexly in the number of trips. The average private cost is described by $c = \omega t(N_{\text{up}}, N_{\text{p}}) + p$, where $p = 0$ for unpriced trips. Note that the costs depend on volumes for both types of trips, capturing that there are spillovers from priced to unpriced. The demand curves $N(c)$ can be thought of as the CDF of outside option values (ε_i) across the population of prospective drivers for a representative priced or unpriced trip. Here we assume that the outside option value for each driver is fixed such that demand is unchanged pre- and post-policy.

Figure B.3: Illustration of Driver Welfare Calculations



For priced trips (Figure B.3a), prices increase by the amount of the toll, which reduces the total number of trips. However, at the new equilibrium volumes, the difference in cost is offset by improved travel times ($\omega\Delta t$). The true change in driver welfare is the area x. Here, by using pre-period volumes, our welfare estimate of $\widehat{\Delta W}_{\text{p}} = -x - y$ overstates the true welfare loss by the area y. The green hatched area denotes the revenue with the new volumes. For unpriced trips (Figure B.3b), introducing congestion pricing reduces travel time costs by $\omega\Delta t$, because travel time costs depend on the number of priced and unpriced trips. With elastic demand, the faster travel times encourage additional driving, shifting the equilibrium volume from N_{up}^0 to N_{up}^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 $\widehat{\Delta W}_{\text{up}} = a$ for unpriced trips, which understates the true welfare gains by the area of b.

B.4 Welfare Calculation Inputs

Table B.1 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.

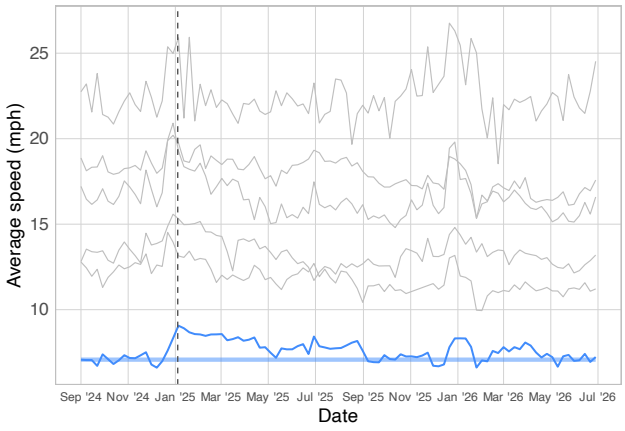
Table B.1: Welfare estimate inputs

				Passenger Vehicles		Taxi/FHVs	
Origin	Destination	$\widehat{\text{ATT}}_{od}$	$\bar{t}_{od}$	N^{PV}_{od}	p^{PV}_{od}	N^{FHV}_{od}	p^{FHV}_{od}
CBD	CBD	0.040	23.2	643104	\$0	607838	\$1.27
CBD	0-20%	0.063	25.4	501032	\$0	262957	\$1.27
CBD	20-40%	0.055	46.6	295583	\$0	100202	\$1.27
CBD	40-60%	0.045	53.8	247297	\$0	83728	\$1.27
CBD	60-100%	0.034	62.2	317983	\$0	123129	\$1.27
0-20%	CBD	0.063	24.8	594195	\$7.87	119394	\$1.27
0-20%	0-20%	0.015	12.6	7738585	\$0	812093	\$0
0-20%	20-40%	0.019	15.0	2627617	\$0	203010	\$0
0-20%	40-60%	0.015	23.3	1508345	\$0	129397	\$0
0-20%	60-100%	0.012	41.1	1122931	\$0	82589	\$0
20-40%	CBD	0.054	47.1	202162	\$7.87	44761	\$1.27
20-40%	0-20%	0.019	15.0	2562395	\$0	148876	\$0
20-40%	20-40%	0.016	13.5	9768322	\$0	265182	\$0
20-40%	40-60%	0.015	14.9	4112082	\$0	133374	\$0
20-40%	60-100%	0.011	26.3	1469808	\$0	40717	\$0
40-60%	CBD	0.047	54.5	215322	\$7.87	53862	\$1.27
40-60%	0-20%	0.016	23.3	1479622	\$0	78476	\$0
40-60%	20-40%	0.015	14.8	4248147	\$0	98916	\$0
40-60%	40-60%	0.013	13.6	22665226	\$0	385637	\$0
40-60%	60-100%	0.009	15.7	6202699	\$0	78086	\$0
60-100%	CBD	0.043	61.5	557136	\$7.87	118015	\$1.27
60-100%	0-20%	0.019	40.9	1085775	\$0	56038	\$0
60-100%	20-40%	0.014	26.0	1562846	\$0	12879	\$0
60-100%	40-60%	0.008	15.6	6613899	\$0	31220	\$0
60-100%	60-100%	0.004	14.1	103590251	\$0	197165	\$0

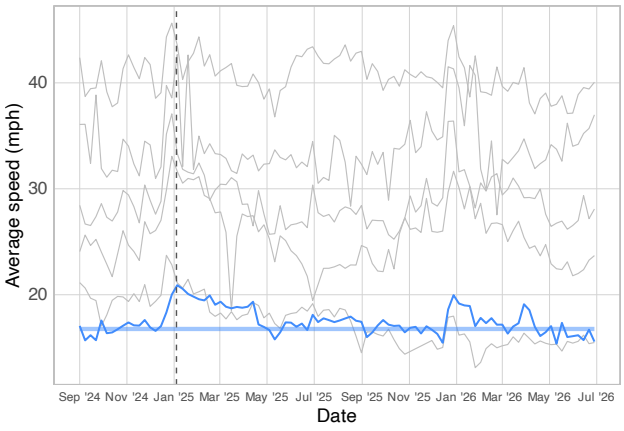
C Supplementary Tables and Figures

Figure C.1: Average speeds on CBD road segments

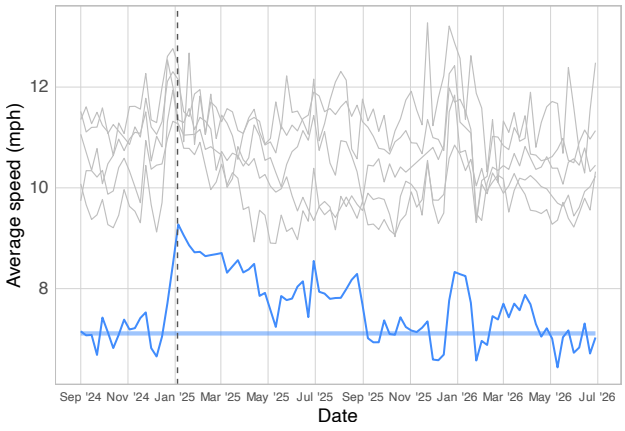
(a) All roads



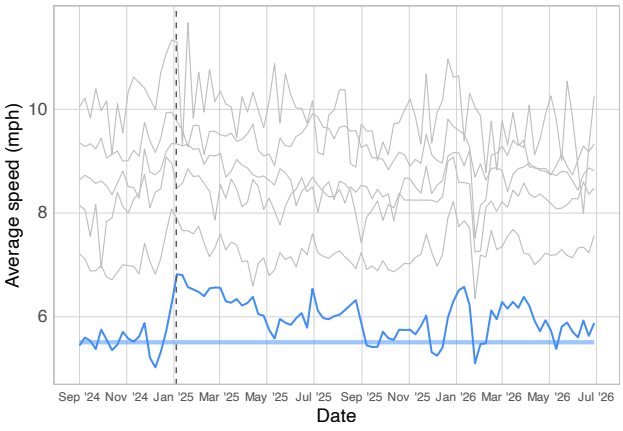
(b) Highways



(c) Arterial roads



(d) Local roads



Notes: This figure documents the traversals-weighted average weekly speed on road segments in the CBD of each city from September 2024 through June 2026. NYC is plotted in blue, and the horizontal blue line represents the average speed in NYC between September 1 and December 15, 2024.

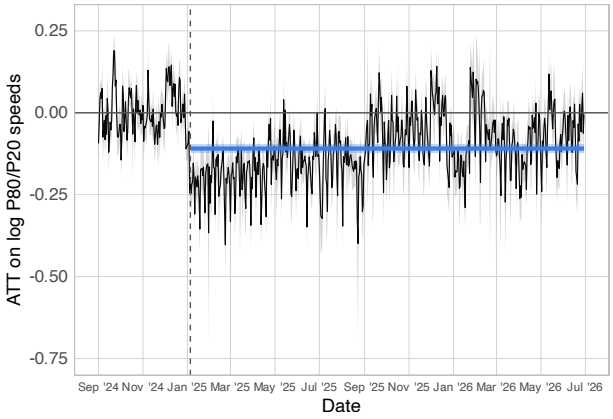
Table C.1: Average CBD speeds before and after congestion pricing

		Average speeds (mph)	
City	Road type	Before Jan 5, 2025	After Jan 5, 2025
Atlanta	All	22.37	22.38
	Highway	33.86	33.75
	Arterial	11.51	11.45
	Local	10.22	9.84
Baltimore	All	12.54	11.69
	Highway	19.86	17.0
	Arterial	10.92	10.66
	Local	8.16	8.3
Boston	All	17.04	16.36
	Highway	28.72	28.78
	Arterial	9.75	9.78
	Local	7.04	7.25
Chicago	All	13.57	13.28
	Highway	25.21	25.39
	Arterial	10.32	10.07
	Local	8.69	8.72
New York	All	7.14	7.69
	Highway	16.95	17.6
	Arterial	7.19	7.69
	Local	5.54	5.99
Philadelphia	All	18.5	17.81
	Highway	40.67	40.65
	Arterial	11.3	11.14
	Local	9.31	9.25

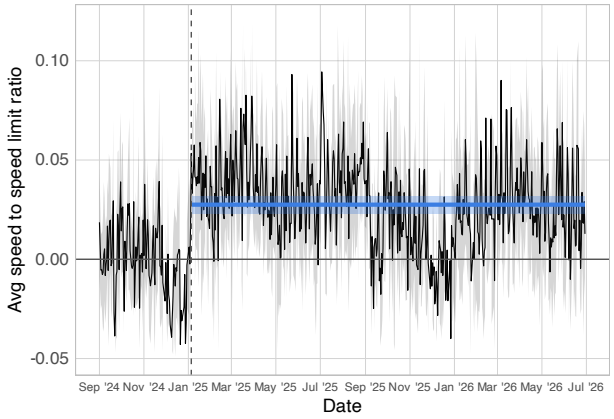
Notes: This table reports the average volume-weighted speeds on segments in each of the cities used throughout our analyses for weekday peak hours. The "before" period data starts September 1st, 2024 and the "after" period data goes through June 2026.

Figure C.2: Event Study Plots for Additional Outcomes

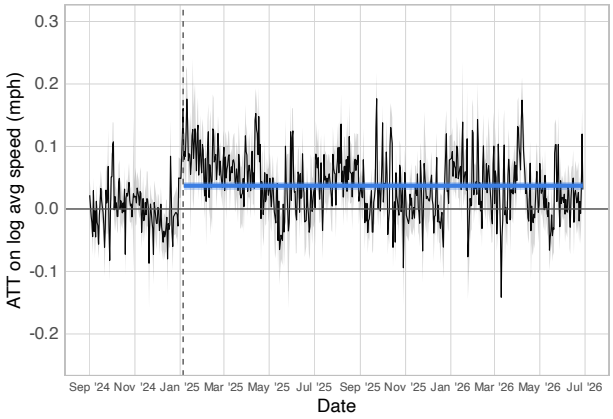
(a) p80/p20 CBD speeds



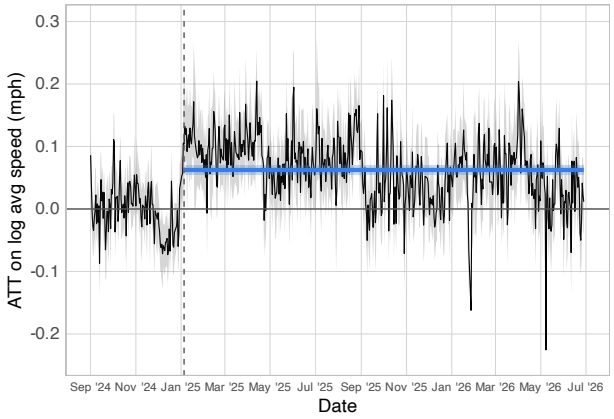
(b) Ratio of speed to speed limit



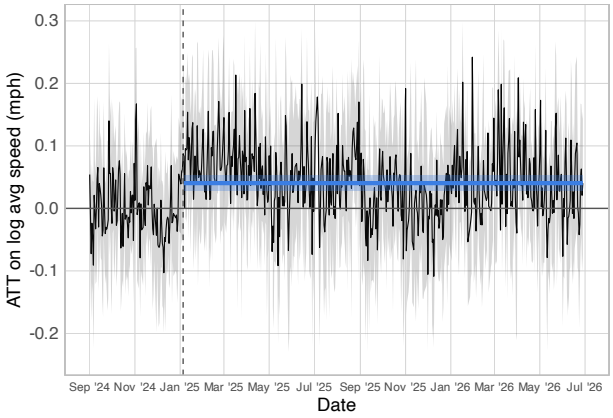
(c) Trips to the CBD



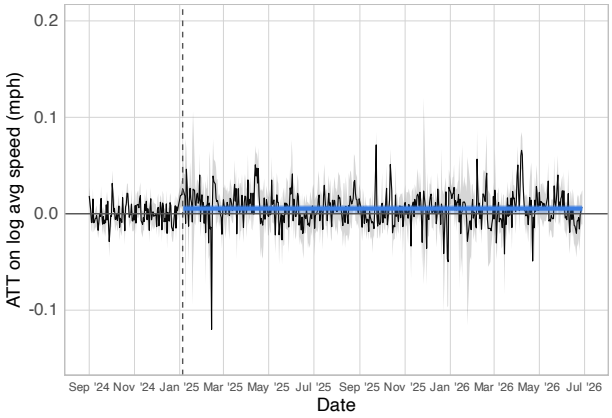
(d) Trips from the CBD



(e) Average speed on within-CBD trips



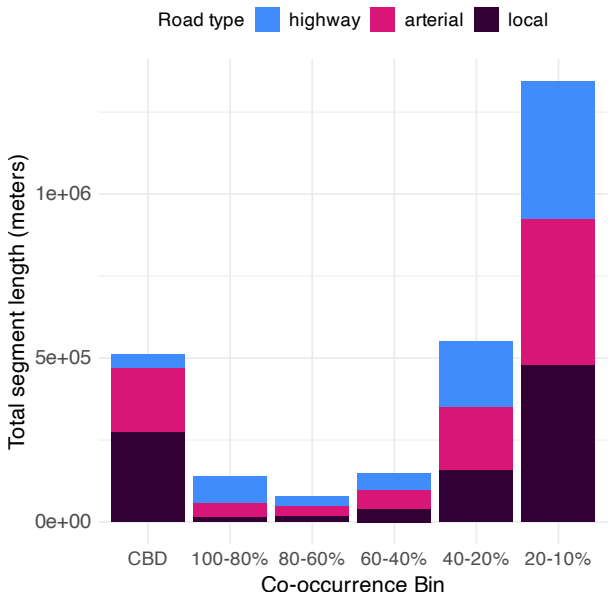
(f) Average speed on outside-CBD trips



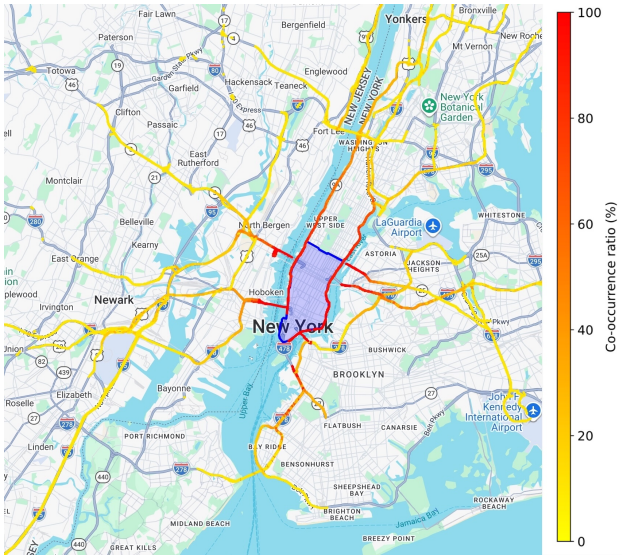
Notes: This figure replicates Figure 1a for additional traffic outcomes. Each figure plots day-level ATTs of congestion pricing over time. The horizontal blue line is the aggregate ATT for all post-treatment periods. Shading denotes 95% confidence intervals. Standard errors are clustered at the city-level.

Figure C.3: Segments by co-occurrence

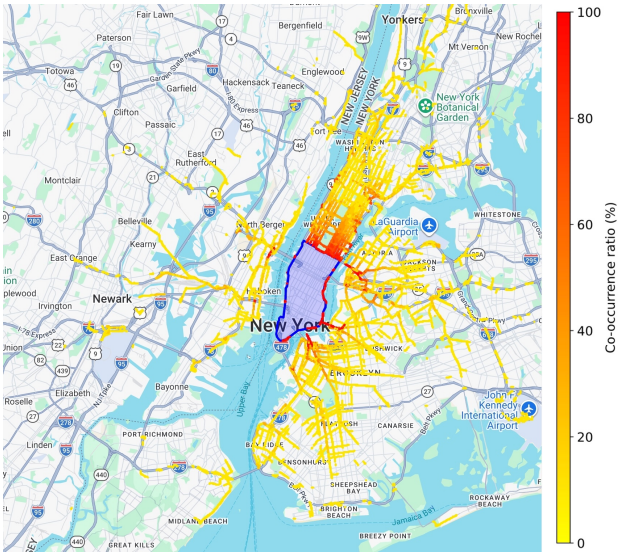
(a) Bin compositions



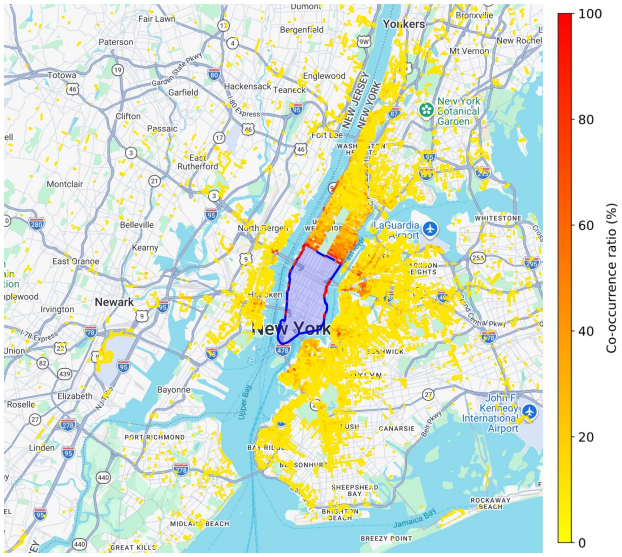
(b) Highways



(c) Arterials

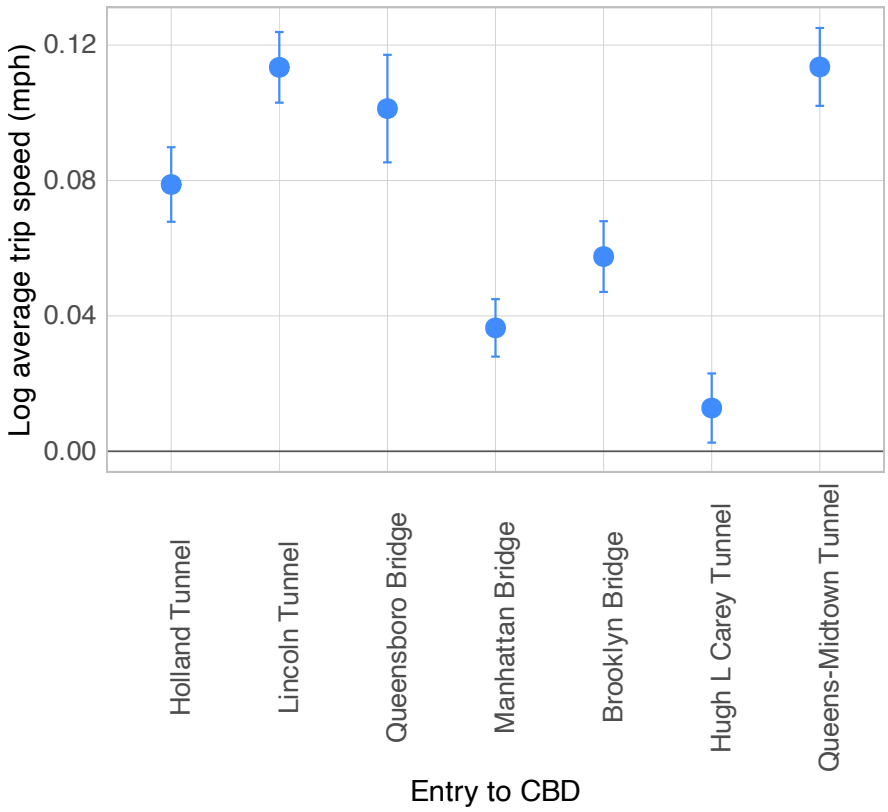


(d) Local roads



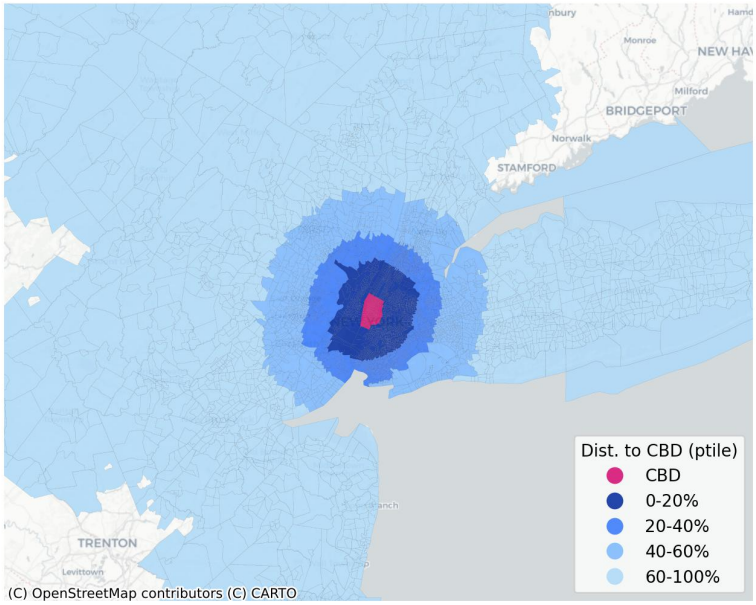
Notes: These figures map the road segments by their level of co-occurrence with the CBD, which is defined as the share of trips crossing the road segment that also cross some segment in the CBD. For the maps, we omit roads with co-occurrence less than 5% as they are too numerous to display easily. The first panel plots the distribution of road types by co-occurrence.

Figure C.4: 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 includes two-hour aggregates of trip speeds on trips to the control city CBDs from any origin. Vertical lines represent 95% confidence intervals. Standard errors are clustered at the city-level.

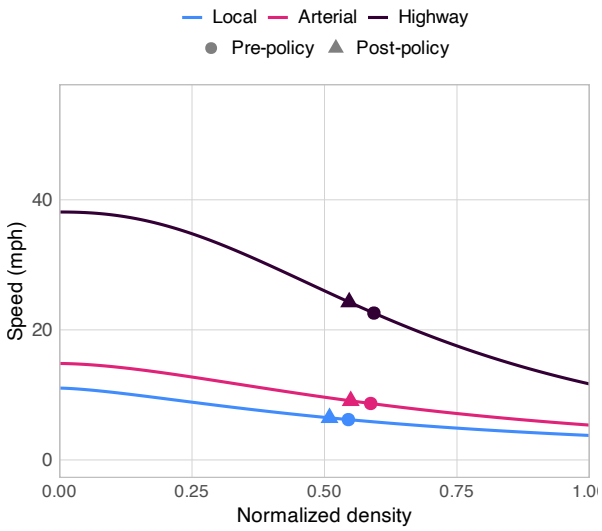
Figure C.5: Bins of distance to CBD



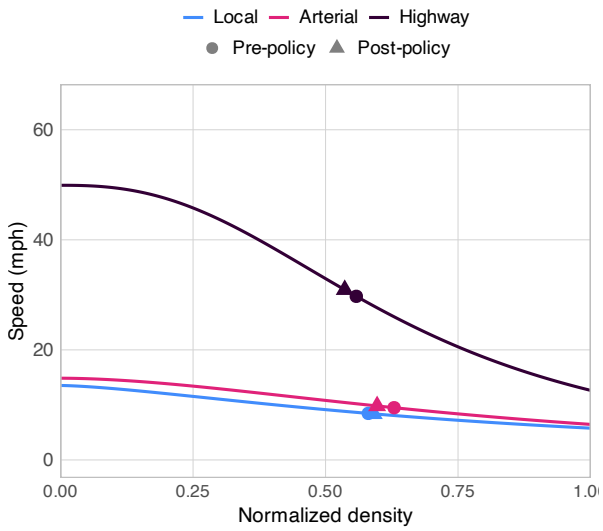
Notes: The boundaries are based on the Core Based Statistical Area boundaries. The map is zoomed in slightly to make the CBD more visible; the CBSA extends further in each direction, especially east down Long Island. Distance to CBD is based on the distance between centroid and the closest point of the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

Figure C.6: Other congestion functions: NYC

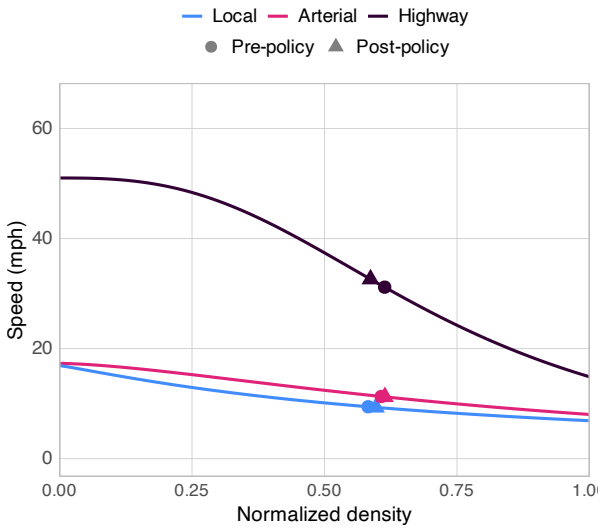
(a) 80-100% co-occurrence roads



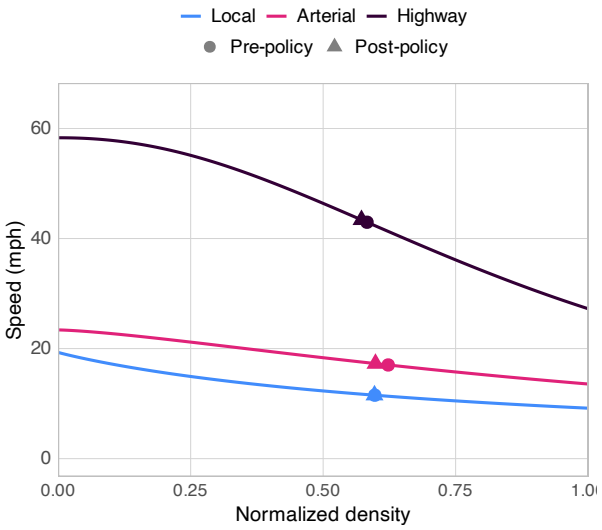
(b) 40-60% co-occurrence roads



(c) 20-40% co-occurrence roads



(d) 0-20% co-occurrence roads



Notes: This figure plots the remaining estimated congestion functions for roads in NYC; see Figure 7 for the CBD and 60–80% 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 using the estimated ATT on speeds.

References

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.

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*.

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

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

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.

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).

MTA. 2026. "MTA Congestion Relief Zone Vehicle Entries: Beginning 2025." https://data.ny.gov/Transportation/MTA-Congestion-Relief-Zone-Vehicle-Entries-Beginni/t6yz-b64h/about_data.

Ostrovsky, Michael, and Frank Yang. 2025. "Effective and equitable congestion pricing: New York City and beyond."